Vol 7 (2026)

The Market

International Journal of Business

The Market: International Journal of Business is a scholarly, open-access peer-reviewed research journal published annually by the Cyprus Centre for Business Research at CIM-Cyprus Business School.

We seek to promote new and productive interaction between various business disciplines and fields. We consider articles that express new and innovative ideas in business, paying particular attention to developments in Cyprus and the broader Eastern Mediterranean area.

Published by:
Cyprus Centre for Business Research
25 Zannetos Street, Ayios Andreas, Nicosia, Cyprus.
E: market@cim.ac.cy

Founding Editor: Prof. John Kaminarides

Editor-in-Chief: Dr Constantinos Constantinou, CIM

Copyeditor: Brian Saunders

Scientific Board:
Prof. Valerio Mancini, *Rome Business School*
Prof. Almuth McDowall, *Birkbeck University of London*
Dr Derek Watson, *University of Sunderland*
Dr Suresh Gamlath, *University of West London*
Dr Serhii Lehenchuk, *Zhytomyr Polytechnic State University, Ukraine*
Dr Myria Kkali, *CIM-Cyprus Business School*
Dr Maria Socratous, *CIM-Cyprus Business School*
Dr George Afxentiou, *CTL Eurocollege*
Dr Katerina Pavlou, *CIM-Cyprus Business School*
Dr Christiana Charalambidou, *CIM-Cyprus Business School*
Dr Roman Puchkov
Dr Louis Karaolis, *CIM-Cyprus Business School*
Dr Isabella Corvino, *Università degli studi di Perugia*
Yotam Werzansky Orland, Esq., *CIM-Cyprus Business School*

Design: Spyros Karkotis

Cyprus Centre for Business Research

Submitting to *The Market: International Journal of Business*
For more information about and guidance on how to submit, please see www.cim.ac.cy/market
© 2026. All rights reserved.

Legally responsible person: Dr (hc) Yangos Hadjiyannis

The Market and its editors, publishers, owners and anyone else invoved in creating, producing or delivering *The Market* do not assume any liability or responsibility for the accuracy, completeness, topicality or value of its contents, nor shall they be liable for any direct, indirect, incidental, special, consequential or punitive damages arising out of its contents.

The Market's policy on the use of AI

Note: The authors are fully responsible for the content of their manuscript, including those parts that have been produced by an artificial intelligence (AI) tool—e.g., ChatGPT, or any Large Language Models (LLMs) programs—and, therefore, are liable for any breach of publication ethics. The authors are expected to acknowledge the use of any AI tool in their writing process, ensuring that they also state how they have used it.

The Market is published by the Cyprus Centre for Business Research, which is part of CIM-Cyprus Business School

ISSN 2547-9202 (print) / ISSN 2547-9210 (online)

CONTENTS

PAGE

EDITORIAL 1

FOREWORD 2

C. Charalambidou

ARTIFICIAL INTELLIGENCE PROCESSES AND COPYRIGHT 3

AI Generated Content and Copyright Post 2024: Human Authorship, AI Training, Memorisation and Doctrinal Continuity 4 - 22

Y. Werzansky-Orland, Esq.

ARTIFICIAL INTELLIGENCE AND ETHICAL LEADERSHIP 23

When the Machine Keeps Its Hands Clean: AI, The Illusion of Purity and the Problem of Dirty Hands 24 - 28

D. Tillyris

ARTIFICIAL INTELLIGENCE AND SMEs 29

Leave AI, End Your Love-Hate Relationship! 30 - 34

M. K. Manesh

ARTIFICIAL INTELLIGENCE AND SOVEREIGNTY COMMERCIALISATION 35

Between dependence and autonomy: Commercialising AI Sovereignty in the European Union, United Kingdom, Canada and India, amid US–China technological competition 36 - 47

I. Sidiropoulos

HUMAN RESOURCE MANAGEMENT 48

Reverse Mentoring as a Catalyst for Authentic Reciprocal Learning in Multigenerational Workforces: An Empirical Research Study in Cyprus 49 - 60

E. Constantinou

INTERNATIONAL MARKET ENTRY STRATEGIES 61

Foreign Market Entry: Traditional and Contemporary Developments. 62 - 73

M. Georgiou

MANAGEMENT 74

Performance Management: The good, the bad and the ugly 75 - 83

Z. Tyrannis

NATION BRANDING, DEFENCE AND GEOPOLITICS 84

Nation Branding, Geopolitics and the Defence Industry: The Case of the Republic of Cyprus. 85 - 92

P. Hadjipavlis and C. Constantinou

ORGANISATIONAL CHANGE MANAGEMENT 93

Change Management Needs Surgery, Not Therapy: Why Radical Consistency Beats Gradual Transformation 94 - 103

S. Meyer

POLITICAL MARKETING 104

The CIS Strategy as a Driver of Self-Esteem and Political Popularity: Posting Selfies to Promote Needs and Values 105 - 111

N. Antoniades

SHIPPING 112

Safety Campaigns in the Maritime Industry 113 - 124

V. Mavrinac

AUTHORS 125 - 131

Disclaimer: The contributing authors hold any potential liability for publishing an identical article to that of another scientific journal, without having first received the explicit consent of the publisher in question. The Editorial Board of 'The Market', the Cyprus Centre for Business Research and CIM-Cyprus Business School may not be held accountable in any way for such or similar acts.

EDITORIAL

Dear Readers,

We are delighted to welcome you to the seventh issue of *The Market: International Journal of Business.*

The Market is a scholarly peer-reviewed and open-access research journal, published annually by the Cyprus Centre for Business Research, the independent research centre housed at CIM-Cyprus Business School. The journal publishes articles that express new and innovative ideas in business, paying particular attention to developments in Cyprus and the broader Eastern Mediterranean region. It publishes the results of research endeavours that show strong future prospects and articles that address both the betterment of human life and of business practices. The expert, dedicated editorial team of *The Market* continues to solicit and welcome articles from scholars at all career stages and on topics related to all fields of business – whether with a local or international outlook.

Thanks to the good work of our team, and as a result of the quality of submissions, *The Market* is hosted on ProQuest and EBSCO, two of the world's most-used databases for scientific journals. Beyond any shadow of doubt, the milestones we celebrate today are thanks to your dedication, invaluable collaboration and genuine support throughout these years. Let me therefore take this opportunity to personally thank each and every one of you for your efforts and commitment.

The Market aims to provide opportunities for the promotion of new dynamic business ideas to enhance research in all business fields – from management, strategy, accounting and finance to HR, energy, cyber security, marketing and shipping. In light of its deliberately broad appeal, *The Market* also welcomes contributions on emerging fields – for example, political marketing, a relatively new area that is increasingly gaining traction worldwide. The goal of *The Market* continues to be the increased awareness of important scholarly achievements and the hosting of break-through research.

In our seventh issue, we have once more ensured coverage of a range of contemporary and pressing issues and topics. These include articles linking artificial intelligence with SMEs, copyright and sovereignty commercialisation, alongside ethical leadership. We also address international market entry strategies, management, including human resource management, and organisational change management. Furthermore, there are articles that provide valuable insights into shipping (safety campaigns), political marketing and nation branding, defence and geopolitics.

In addition to these topics, and demonstrating our commitment to a new generation of scholarship, *The Market* continues to publish the top dissertations of CIM-Cyprus Business School Master's students.

In line with the successful ninth Annual Academic Conference organised by the Centre last September, this seventh issue of the journal runs under the umbrella theme *"Unveiling New Business Strategies: Bringing Human and Artificial Intelligence Together".*

We hope that you will enjoy reading *The Market* as much as we have enjoyed assembling it. We thank you for being with us on this journey, and we look forward to reading your research paper and/or research note next time!

Dr Constantinos Constantinou
Editor-in-Chief

FOREWORD

It is a great pleasure to contribute the Foreword to this year's edition of *The Market*, which once again bridges research and insights on important developments shaping businesses, organisations and economies.

Artificial intelligence (AI) has rapidly shifted from being viewed as an emerging technology to becoming part of everyday business practice. The key question for businesses is therefore no longer whether AI will affect them, but how they can use it effectively, responsibly and strategically.

Although much of the discussion around AI in recent years has focused on its potential to act as a substitute for human tasks, attention is increasingly shifting towards its role as a complement to human intelligence rather than a replacement for it. AI can process vast amounts of information, identify patterns and automate complex tasks, while people contribute judgement, creativity, experience, empathy and an understanding of context. The greatest value is likely to come from combining these strengths rather than treating the two as competing forces.

As AI tools become increasingly accessible, competitive advantage will not only come from simply having access to technology, but from how effectively organisations manage to integrate it into their processes, invest in the right skills and manage the organisational change that follows. This is especially relevant for SMEs, which can now access capabilities that once required significant financial and human resources.

At the same time, AI brings wider challenges. Businesses must learn to combine technological capability with human intelligence, sound judgement and responsible leadership, while also addressing increasingly important questions surrounding copyright, intellectual property, governance, accountability and sovereignty, as well as the implications for employees, skills and the nature of work. Those that successfully strike this balance will be better placed not only to adapt to change, but to help shape it.

The breadth of contributions in this edition of *The Market* reflects this wider business environment. The articles examine among others AI in relation to SMEs, copyright and commercialisation, ethical leadership, while also addressing international market entry strategies, management, human resource management, organisational change, shipping, political marketing, nation branding, defence and geopolitics.

I am confident that the ideas and research presented in this edition will contribute meaningfully to the ongoing discussion on "New Business Strategies: Bringing Human and Artificial Intelligence Together," which also served as the umbrella theme of this year's 9th CIM Academic Conference.

Dr Christiana Charalambidou

Associate Professor, MBA dissertation supervisor and Student mentor

AI Generated Content and Copyright Post 2024: Human Authorship, AI Training, Memorisation and Doctrinal Continuity

By

Yotam Werzansky-Orland, Esq.

Head of Innovation and Intellectual Property and Doctoral Fellow at

the Cyprus Centre for Business Research – CIM-Cyprus Business School

Abstract:

Generative artificial intelligence has moved copyright law from theoretical anxiety to judicial practice. This article revisits the core doctrine developed in the author's earlier work and tests it against the first major wave of post 2024 cases, legislative measures and regulatory instruments concerning AI-generated works, AI training, memorisation, outputs and platform liability. It argues that copyright law has not become obsolete. Courts and regulators have largely reaffirmed its foundational commitments: human authorship remains central; influence is not expression; learning is not automatically copying; and liability still depends on control, causation and legally cognisable acts of use.

Yet the emerging case law also exposes areas where doctrine remains underdeveloped. The legal characterisation of model training, the threshold at which memorisation becomes reproduction, the allocation of responsibility among developers, deployers and users and the evidentiary burden of proving infringement all remain unsettled. The article therefore rejects both extremes: neither wholesale immunity for AI development nor a categorical reconstruction of copyright law is justified. The better path is doctrinal continuity with sharper operational tests. Copyright's existing architecture can still govern generative AI, but only if courts clarify how its traditional concepts apply inside the technical realities of machine learning systems.

Keywords: generative AI; copyright law; human authorship; AI training; reproduction right; text and data mining; fair use; memorisation; AI outputs; platform liability.

Contents

I. Introduction – From Doctrine to Judicial Practice

II. The Original Doctrine Revisited
- 2.1 Human Authorship as the Foundation of Copyright
- 2.2 AI Training and the Distinction Between Learning and Copying
- 2.3 Conscious and Willful Creative Involvement

III. Judicial Practice Since 2024
- 3.1 Human Authorship
- 3.2 AI Training: Learning, Copying and the Emerging Judicial Consensus
- 3.3 AI-Generated Outputs, Memorisation and Liability

IV. Where Judicial Practice Confirms the Original Theory
- 4.1 Human Authorship Remains the Foundation of Copyright
- 4.2 Existing Copyright Doctrine Remains Structurally Applicable
- 4.3 AI Training Is Not Automatically Treated as Copyright Infringement
- 4.4 The Debate Has Shifted from Doctrine to Application

V. Where Judicial Practice Still Falls Short
- 5.1 The Continuing Uncertainty Surrounding AI Training
- 5.2 Memorisation Remains the Principal Doctrinal Challenge
- 5.3 Market Substitution and the Changing Nature of Copyright Harm
- 5.4 Allocation of Responsibility Within the AI Ecosystem
- 5.5 Transparency, Evidence and the Practical Administration of Copyright
- 5.6 Refinement Rather Than Replacement

VI. What Courts Still Need to Clarify
- 6.1 Defining the Boundary Between Learning and Reproduction
- 6.2 Developing a Coherent Doctrine of Memorisation
- 6.3 Clarifying the Threshold of Human Creative Contribution
- 6.4 Allocating Responsibility Across the AI Ecosystem
- 6.5 Transparency, Evidence and the Administration of Copyright
- 6.6 Refining Existing Doctrine Rather Than Creating New Doctrine

VII. Conclusion

I. Introduction – From Doctrine to Judicial Practice

This article follows the author's 2024 article, entitled "*AI-generated content and the question of copyright*", published in *The Market: International Journal of Business* (Werzansky-Orland, 2024). That article was written at a moment when many of the central copyright questions raised by generative artificial intelligence were still framed as hypotheticals. It argued that the emergence of systems capable of producing text, images, music and other expressive outputs did not require abandonment of copyright doctrine's human-centered foundations. The core claim was that existing principles - authorship, originality, reproduction, infringement, exceptions and liability - remained structurally capable of addressing AI-generated content, provided they were applied with sufficient doctrinal precision (Gervais, 2020; Ginsburg & Budiardjo, 2019; Samuelson, 2023; U.S. Copyright Office, 2025a; Werzansky-Orland, 2024).

That proposition can now be tested against judicial practice. Since 2024, courts and regulators have considered AI authorship, training, memorisation, infringing outputs, platform responsibility, fair use, text and data mining and transparency obligations across a range of jurisdictions (Thaler v. Perlmutter, 2025; Bartz v. Anthropic, 2025; Kadrey v. Meta Platforms, Inc., 2025; Kneschke v. LAION e.V., 2024; GEMA v. OpenAI, 2025; Regulation (EU) 2024/1689; U.S. Copyright Office, 2025b). These developments have not produced a settled global doctrine, but they have moved the debate from abstract theory to operational legal practice.

This article therefore revisits the original doctrine in light of those developments. It does not seek to replace the earlier thesis or propose a new copyright code for AI. Its purpose is to examine where recent judicial and regulatory practice confirms the original argument, where it exposes unresolved uncertainties, and what courts still need to clarify if copyright law is to remain coherent in an environment shaped by probabilistic models, largescale datasets and increasingly ordinary AI-assisted creative workflows.

The argument proceeds in five steps. Section II restates the original doctrinal framework. Section III examines the principal judicial and regulatory developments since 2024. Section IV identifies the respects in which those developments confirm the original theory. Section V explains where judicial practice still falls short. Section VI identifies the issues courts still need to clarify. Section VII concludes that generative AI has not made copyright doctrine obsolete, but has made its disciplined application more important.

II. The Original Doctrine Revisited

The "original doctrine" referred to in this section is the framework developed in the author's earlier article, which rested on three connected propositions: first, that copyright protection remains anchored in human authorship; second, that copyright must distinguish copying protected expression from learning, influence, inspiration or stylistic proximity; and third, that legal responsibility in AI-assisted creation should turn on meaningful human involvement, control and conscious contribution rather than on the mere technical use of an AI system (Werzansky-Orland, 2024).

2.1 Human Authorship as the Foundation of Copyright

The emergence of generative artificial intelligence has reignited debate concerning one of copyright law's most fundamental principles: that copyright protects the product of human intellectual creativity. As AI systems become increasingly capable of generating text, images, music and other forms of expressive content that resemble human created works, questions have arisen as to whether existing copyright doctrine remains adequate or whether a fundamentally new legal framework is required (Hyun, Park & Kim, 2025; Lucchi, 2025; Samuelson, 2023).

The earlier article argued that, notwithstanding these technological developments, copyright law remains conceptually grounded in human authorship. Across common law and civil law jurisdictions alike, copyright protection has consistently depended upon the existence of a human author who exercises creative judgment in producing an original work. Although different legal systems articulate this requirement in different ways—whether through originality, intellectual creation or creative expression—the underlying premise is the same: copyright protects the fruits of human creativity rather than the autonomous operation of machines (Gervais, 2020; Ginsburg & Budiardjo, 2019; Samuelson, 2023; U.S. Copyright Office, 2025a).

Generative AI undoubtedly challenges the practical application of this principle. Modern AI systems are capable of producing sophisticated outputs that may appear creative and, in some instances, may be difficult to distinguish from works produced by humans. However, apparent creativity should not be conflated with legal authorship. Contemporary AI systems do not possess consciousness, intention, legal personality or independent creative will. Rather, they generate outputs through probabilistic inference based upon patterns identified during the training process (Bommasani et al., 2021; Goodfellow, Bengio & Courville, 2016). Consequently, while AI may function as a powerful creative tool, it does not itself satisfy the traditional legal conception of an author.

Accordingly, the central question is not whether AI-generated outputs may possess aesthetic or commercial value, but whether the legal concept of authorship should extend beyond human creative activity. The earlier article answered this question in the negative, arguing that copyright protection should continue to depend upon meaningful human creative involvement. As demonstrated in the following sections, the first wave of judicial and regulatory developments since 2024 has largely reaffirmed this fundamental premise, while simultaneously exposing important practical questions regarding the degree of human contribution required in AI-assisted creative processes.

2.2 AI Training and the Distinction Between Learning and Copying

One of the principal challenges posed by generative artificial intelligence concerns the legal characterisation of the training process itself. AI systems are developed by analysing vast quantities of existing material in order to identify statistical relationships between words, images, sounds and other forms of expression. This has prompted the question whether AI training should be regarded as an act of copyright infringement or whether it is more appropriately understood as a technologically advanced form of learning (Henderson et al., 2023; Lemley & Casey, 2021; Lucchi, 2025; U.S. Copyright Office, 2025b).

The earlier article argued that AI training is conceptually closer to learning than to copying. Human creators routinely acquire knowledge by reading books, studying works of art, listening to music and observing the creative achievements of others. The influence of prior works on subsequent creativity has never, in itself, constituted copyright infringement. Copyright protects the expression of ideas rather than knowledge, styles, techniques or general artistic influence. Consequently, the process of learning from existing works has traditionally remained outside the scope of copyright protection (Gervais, 2020; Lemley & Casey, 2021).

Generative AI similarly acquires the capacity to produce new outputs by identifying patterns within large datasets rather than by reproducing the underlying works for expressive purposes. The objective of training is not to preserve or disseminate protected expression, but to enable the model to infer linguistic, artistic or structural relationships from the data presented to it (Bommasani et al., 2021; Goodfellow, Bengio & Courville, 2016; Henderson et al., 2023; Lucchi, 2025; U.S. Copyright Office, 2025b). From this perspective, AI training is analogous to the acquisition of knowledge rather than the creation of substitute copies.

This argument develops, rather than merely repeats, the position advanced in the earlier article. The 2024 article treated the AI system's use of prior works as closer to inspiration than copying where the earlier works

are not reproduced "as is" but are used as inputs in the generation of a new work (Werzansky-Orland, 2024). The present article applies that premise to the more specific post 2024 debate concerning training datasets, intermediate reproductions, memorisation and outputs.

This distinction does not suggest that all forms of AI training are necessarily lawful. Copyright concerns may arise where protected works are unlawfully acquired, where statutory exceptions do not apply, or where the training process results in the retention or reproduction of protected expression beyond what copyright law permits. Likewise, evidence that a model has memorised and reproduced substantial portions of copyrighted works may give rise to separate questions concerning infringement (Carlini et al., 2023; Dornis, 2025; Stober & Dornis, 2026). These issues, however, should be distinguished from the broader proposition that the act of learning from existing works is not, in itself, equivalent to unlawful copying.

Accordingly, the legal analysis should focus not on the existence of training alone, but on the manner in which training is conducted, the provenance of the training materials, the applicable statutory exceptions, and the nature of the resulting outputs. As demonstrated in the following sections, recent judicial decisions have largely adopted this more nuanced approach, distinguishing between the process of training, the acquisition of training data, and the generation of infringing outputs.

2.3 Conscious and Willful Creative Involvement

The earlier article proposed that the question of authorship in AI-assisted works should not be determined solely by the presence or absence of artificial intelligence, but by the extent of meaningful human creative involvement in the resulting work. The mere use of AI as a tool should not, in itself, preclude copyright protection. Rather, the decisive consideration should be whether the final work reflects conscious and willful creative choices attributable to a human author.

The earlier article developed this model primarily as a test for copyright eligibility in non-human generated works, asking whether the output resulted from random or autonomous generation, or from conscious and willful creative involvement (Werzansky-Orland, 2024). The present article applies that model more narrowly to AI-assisted human creativity: not to confer authorship on the AI system, but to identify whether the human user's contribution is sufficiently creative, deliberate and controlling to justify copyright protection.

This approach recognises that copyright law has historically adapted to technological innovation without abandoning its underlying principles. The introduction of photography, digital image editing, computer-aided design and other creative technologies did not eliminate human authorship; instead, the law continued to focus on the intellectual contribution made by the individual using those technologies (Gervais, 2020; Samuelson, 2023). Generative AI should be assessed according to the same principle.

Meaningful human involvement may be demonstrated through activities such as selecting training inputs, formulating and refining prompts, exercising creative judgment during iterative generation, selecting among alternative outputs, editing or modifying generated material, or otherwise directing the creative process in a manner that results in the author's own intellectual creation. Conversely, where an AI system generates expressive content without meaningful human creative direction, the resulting output should not qualify for copyright protection merely because a human initiated the process (Fritz, 2025; U.S. Copyright Office, 2025a).

The conscious and willful involvement model therefore seeks to distinguish between AI as an instrument of human creativity and AI as an autonomous generator of expressive content. This distinction preserves the human-centred foundations of copyright law while recognising that creative practice increasingly incorporates sophisticated computational tools. It avoids both extremes of automatically denying protection to all AI-assisted works and automatically recognising copyright in every AI-generated output.

As demonstrated in the following sections, the judicial and regulatory developments since 2024 have largely continued to focus on identifying the nature and extent of human creative contribution, rather than abandoning the traditional requirement of human authorship. Although courts have employed different doctrinal formulations, the underlying inquiry has consistently remained whether the protected expression ultimately reflects the intellectual creation of a human author.

III. Judicial Practice Since 2024

3.1 Human Authorship

Since the publication of the earlier article, courts and regulatory authorities have begun applying established copyright principles to generative AI with increasing frequency. While the factual circumstances differ across jurisdictions, a common theme has emerged: none has departed from the principle that copyright subsists only in works attributable to human creative activity. Rather than recognising AI systems as authors, recent decisions have focused on identifying whether sufficient human intellectual contribution exists to justify copyright protection (Japan Copyright Office, 2024; U.S. Copyright Office, 2025a).

The clearest affirmation of this principle is found in Thaler v. Perlmutter, in which the United States Court of Appeals for the District of Columbia Circuit held that the Copyright Act requires human authorship and therefore does not extend protection to works created solely by artificial intelligence (*Thaler v. Perlmutter*, 2025). Importantly, the Court distinguished between works autonomously generated by AI and works created with the assistance of AI, expressly leaving open the possibility that AI-assisted works may qualify for protection where the requisite human authorship is present. The decision therefore reinforces, rather than alters, the traditional human centred conception of copyright.

A similar approach was adopted by the Prague Municipal Court in S. Š. v. Taubel Legal, where the claimant sought copyright protection for an image allegedly generated using artificial intelligence. The Court rejected the claim, finding that the claimant had failed to demonstrate that the image constituted the product of his own creative intellectual activity. Rather than questioning existing copyright doctrine, the Court applied conventional principles of authorship and evidentiary proof to an AI-assisted work (S. Š. v. Taubel Legal, Municipal Court in Prague, 2023).

Regulatory developments likewise support this approach. The United States Copyright Office has consistently maintained that copyright protects only works embodying human authorship, while recognising that AI-assisted works may qualify for protection where sufficient human creative expression is reflected in the final work (U.S. Copyright Office, 2025a). Similarly, the Japanese Copyright Office distinguishes between AI functioning as a creative tool under human direction and outputs generated without meaningful human creative contribution (Japan Copyright Office, 2024). Italy's 2025 AI law points in the same direction by expressly tying protection for AI-assisted works to the author's human intellectual work (Italy, 2025). Although these authorities differ in terminology and legal framework, they converge on the principle that human intellectual creation remains the indispensable foundation of copyright protection.

Accordingly, the first generation of judicial and regulatory developments has not displaced the theoretical framework advanced in the earlier article. On the contrary, they largely confirm that the central challenge posed by generative AI lies not in redefining authorship, but in determining the threshold of human creative involvement necessary to satisfy established copyright principles. The debate has therefore shifted from whether human authorship remains relevant to how that requirement should be applied in increasingly sophisticated AI-assisted creative processes.

3.2 AI Training: Learning, Copying and the Emerging Judicial Consensus

The legality of AI training has become one of the most contested issues in contemporary copyright law. While the earlier article argued that AI training is conceptually closer to learning than to copying, recent judicial decisions demonstrate that the issue is considerably more nuanced than a simple dichotomy between lawful learning and unlawful reproduction. Rather than adopting categorical positions, courts have increasingly distinguished between the acquisition of training materials, the process of training itself, the applicability of statutory exceptions, and the generation of infringing outputs (Dornis & Lucchi, 2025; Henderson et al., 2023; Lucchi, 2025; U.S. Copyright Office, 2025b).

The United States District Court for the Northern District of California addressed these issues directly in Bartz v. Anthropic. The Court held that the use of copyrighted works to train a large language model was, in principle, highly transformative and capable of constituting fair use, emphasising that the objective of training was not to reproduce expressive content but to enable the model to develop new capabilities (*Bartz v. Anthropic*, 2025). At the same time, however, the Court drew a clear distinction between the training process and the acquisition of the training corpus, holding that the alleged use of pirated copies could not be justified merely because the ultimate purpose of training was transformative. The decision therefore supports the proposition that the legality of AI training cannot be assessed independently of the provenance of the training materials.

Only days later, the same court considered similar issues in Kadrey v. Meta Platforms, Inc. While ultimately granting summary judgment in favour of Meta, the Court rejected the suggestion that transformative use automatically establishes fair use in every AI training case. Instead, it emphasised the importance of examining evidence of market substitution, competitive harm and the economic impact of AI-generated outputs on copyright holders (Kadrey v. Meta Platforms, Inc., 2025). The judgment therefore illustrates that even where AI training may be permissible in principle, the broader commercial consequences of generative AI remain relevant to the copyright analysis.

A different analytical framework has emerged within the European Union. In Kneschke v. LAION e.V., the Regional Court of Hamburg considered whether the creation of AI training datasets fell within the text and data mining exceptions introduced by the DSM Directive. The Court concluded that AI training may constitute text and data mining within the meaning of the Directive and that, in the circumstances of the case, LAION's activities were protected by the scientific research exception under German copyright law (Kneschke v. LAION e.V., LG Hamburg, 2024). At the same time, the judgment recognised that reproductions made during the training process remain acts governed by copyright law and therefore require an applicable statutory exception or other legal justification.

Regulatory developments broadly reflect the same approach. The Japanese Copyright Office has recognised that AI training will generally fall within Article 304 of the Japanese Copyright Act where works are used for non-enjoyment purposes, while acknowledging that different considerations may arise where training is directed towards reproducing particular protected expression or where protected material is obtained unlawfully (Japan Copyright Office, 2024). The European Union's AI Act and the General Purpose AI Code of Practice increasingly emphasise transparency, copyright compliance and respect for rights reservations as part of the broader governance framework for general purpose AI models (Regulation (EU) 2024/1689, arts. 53 and 56; European Commission, 2025a; Peukert & Castets-Renard, 2025). National legislation has also followed a human-centred governance path, including Italy's copyright amendments for AI-assisted works and Japan's promotion-oriented AI Act (Italy, 2025; Japan, 2025; White & Case LLP, 2026).

Taken together, these developments largely support the central proposition advanced in the earlier article that AI training should not automatically be equated with copyright infringement. At the same time, they demonstrate that contemporary copyright analysis has moved beyond the simple question of whether training

is lawful. Judicial and regulatory attention has increasingly shifted towards the legality of dataset acquisition, compliance with statutory exceptions, market effects, transparency obligations and the evidentiary question of whether protected expression is ultimately reproduced by the trained model. The emerging jurisprudence therefore refines, rather than rejects, the theoretical framework advanced in the earlier article.

3.3 AI-Generated Outputs, Memorisation and Liability

While the legality of AI training has attracted considerable attention, recent litigation demonstrates that the more difficult questions increasingly concern the outputs generated by AI systems and the circumstances in which those outputs may infringe copyright. In particular, courts have begun examining whether AI-generated outputs merely reflect statistical learning or whether they reproduce protected expression to such an extent that existing copyright principles are engaged. This shift has moved the debate beyond the training process itself towards questions of memorisation, substantial similarity and the allocation of legal responsibility (Carlini et al., 2023; Dornis, 2025; Foong, 2025; Lucchi, 2025).

One of the most significant developments in this regard is the judgment of the Regional Court of Munich I in GEMA v. OpenAI. The proceedings required the Court to consider whether AI-generated outputs reproducing copyrighted song lyrics constituted actionable copyright infringement. In doing so, the Court examined issues extending beyond conventional questions of training, including the possibility that protected expression may be retained within a model and subsequently reproduced in response to user prompts (GEMA v. OpenAI, LG München I, 2025; Dornis, 2025). Although the judgment does not establish a general doctrine of model memorisation, it illustrates the increasing judicial willingness to examine whether the generation of protected expression may itself constitute copyright infringement, irrespective of the broader legality of the training process.

Chinese courts have likewise begun addressing copyright issues arising from AI-generated outputs. In Shanghai Character License Administrative Co., Ltd. v. AI Company, the Guangzhou Internet Court held that the operator of an AI image generation platform was liable for generating images closely resembling the copyrighted Ultraman character, emphasising the provider's duty to adopt reasonable measures to prevent infringing outputs (Shanghai Character License Administrative Co., Ltd. v. AI Company, Guangzhou Internet Court, 2024). Similarly, in Shanghai [Culture Development Company] v. Hangzhou [AI Technology Company], the Hangzhou Internet Court, subsequently affirmed by the Hangzhou Intermediate People's Court, recognised contributory liability arising from the dissemination of AI models capable of generating infringing works (Shanghai [Culture Development Company] v. Hangzhou [AI Technology Company], Hangzhou Internet Court, 2024; Hangzhou Intermediate People's Court, 2024). Although arising within a different legal framework, both decisions demonstrate an increasing judicial focus on platform responsibility and reasonable preventive measures.

The 2025 judgment in Getty Images (US), Inc. v. Stability AI Ltd likewise illustrates the evolving distinction between training activities and infringing outputs. Although several copyright claims were abandoned before judgment, the case highlighted fundamental questions concerning the relationship between training datasets, model parameters, generated outputs and the evidentiary burden required to establish infringement (Getty Images v. Stability AI, 2025). The case demonstrates that future copyright disputes are likely to focus less on the abstract legality of AI training and more on whether particular outputs reproduce protected expression in a legally significant manner.

Collectively, these developments indicate that copyright litigation involving generative AI is increasingly centred on established infringement principles rather than on the creation of novel legal doctrines. Courts have not abandoned traditional concepts such as reproduction, substantial similarity, causation and authorship. Instead, they are applying those principles to increasingly complex factual circumstances involving probabilistic models, large scale datasets and AI-generated outputs. The principal challenge therefore lies not in identifying

new legal rules, but in determining how existing copyright doctrine should be applied where the relationship between training data, model behaviour and generated expression is technically complex and often difficult to prove.

IV. Where Judicial Practice Confirms the Original Theory

The judicial and regulatory developments examined above demonstrate a remarkable degree of continuity with the theoretical framework advanced in the earlier article. Despite the unprecedented capabilities of generative AI, courts have not responded by abandoning or fundamentally restructuring copyright doctrine. Instead, they have consistently applied established copyright principles to new technological circumstances. While significant uncertainty remains regarding particular aspects of AI training and output generation, the core architecture of copyright law has proven considerably more resilient than many early commentators anticipated (Lucchi, 2025; Samuelson, 2023).

4.1 Human Authorship Remains the Foundation of Copyright

Perhaps the strongest confirmation of the original theory is the continued insistence upon human authorship as a prerequisite for copyright protection. Neither courts nor regulators have recognised AI systems as authors capable of originating copyright. Instead, recent authorities consistently focus on identifying the existence and extent of human intellectual contribution to the work in question.

This position is reflected in Thaler v. Perlmutter, the decision of the Prague Municipal Court in S. Š. v. Taubel Legal, the guidance issued by the United States Copyright Office, and the Japanese Copyright Office's 2024 guidance. Although these authorities arise from different legal systems and employ different doctrinal language, they converge on the same fundamental premise: copyright protects human intellectual creation, while AI remains a technological tool whose legal significance depends upon the nature and extent of human creative contribution (Thaler v. Perlmutter, 2025; S. Š. v. Taubel Legal, 2023; U.S. Copyright Office, 2025a; Japan Copyright Office, 2024).

Accordingly, developments since the publication of the earlier article strongly support its central conclusion that technological sophistication alone does not justify extending authorship to autonomous AI systems. Rather than redefining authorship, courts have reinforced the traditional inquiry into whether the protected expression reflects the intellectual creation of a human author.

4.2 Existing Copyright Doctrine Remains Structurally Applicable

The earlier article argued that generative AI does not require an entirely new body of copyright law. Instead, established doctrines - including authorship, originality, reproduction, infringement, fair use and statutory exceptions - remain capable of addressing the legal questions presented by AI-generated content.

The first generation of AI litigation largely confirms that proposition. Rather than creating AI specific copyright doctrines, courts have consistently relied upon familiar legal concepts when resolving disputes involving generative AI. In the United States, Bartz and Kadrey applied traditional fair use analysis. Within the European Union, Kneschke v. LAION e.V. examined the applicability of the existing text and data mining exceptions introduced by the DSM Directive. In Germany, GEMA v. OpenAI analysed reproduction and infringement through conventional copyright principles. Likewise, the Chinese decisions concerning AI-generated Ultraman images applied established concepts of direct and contributory copyright liability to AI-generated outputs rather than developing AI-specific liability rules (Bartz v. Anthropic, 2025; Kadrey v. Meta Platforms, Inc., 2025; Kneschke v. LAION e.V., 2024; GEMA v. OpenAI, 2025; Shanghai Character License Administrative Co., Ltd. v. AI Company, 2024).

The legal questions have undoubtedly become more technically complex. Nevertheless, the analytical framework employed by courts remains recognisably one of copyright law rather than technology regulation. This suggests that the principal challenge posed by generative AI lies not in the inadequacy of copyright doctrine itself, but in its application to increasingly complex factual and technical circumstances.

4.3 AI Training Is Not Automatically Treated as Copyright Infringement

A central proposition advanced in the earlier article was that AI training should not automatically be equated with unlawful copying. Recent judicial decisions support this proposition, although they also demonstrate that the issue requires a more nuanced legal analysis than was initially apparent.

Both Bartz and Kneschke recognised that AI training may, depending upon the applicable legal framework, be justified through existing copyright doctrines, including fair use and statutory text and data mining exceptions. At the same time, neither court suggested that AI developers enjoy blanket immunity. Instead, both emphasised that separate questions - such as the provenance of training materials, compliance with statutory requirements, the legality of intermediate reproductions and the reproduction of protected expression in outputs - must be examined independently (Bartz v. Anthropic, 2025; Kneschke v. LAION e.V., 2024).

Regulatory developments point in the same direction. The Japanese Copyright Office accepts that AI training may constitute a permissible non enjoyment use under Article 304 of the Japanese Copyright Act, while recognising that unlawful acquisition of works or the reproduction of protected expression may fall outside that protection. Similarly, the European Union's evolving regulatory framework increasingly focuses upon transparency, copyright compliance and rightsholder reservations rather than treating AI training as inherently unlawful (Japan Copyright Office, 2024; Regulation (EU) 2024/1689, art. 53; European Commission, 2025a; Peukert & Castets-Renard, 2025).

Taken together, these developments confirm the central premise advanced in the earlier article: the legality of AI training cannot be determined through categorical assumptions. Instead, it requires a contextual analysis of the applicable copyright doctrines, statutory framework and factual circumstances.

4.4 The Debate Has Shifted from Doctrine to Application

Perhaps the most significant development since the publication of the earlier article is that the principal questions are no longer conceptual but operational. Courts have largely accepted that existing copyright doctrine provides the appropriate legal framework for resolving disputes involving generative AI. The remaining challenges concern its practical application to increasingly sophisticated AI technologies.

Current disputes therefore focus upon questions such as the degree of human creative involvement required for authorship, the lawful acquisition and use of training materials, the significance of memorisation and regurgitation, the allocation of responsibility among developers, deployers and users, the operation of statutory exceptions, and the evidentiary burden required to establish infringement. These are not questions arising from deficiencies in copyright doctrine itself, but from the practical application of long-established legal principles to novel technological contexts (Carlini et al., 2023; Henderson et al., 2023; Lucchi, 2025; U.S. Copyright Office, 2025b).

Taken together, the authorities examined in this article largely validate the central thesis advanced in the earlier article. Existing copyright doctrine has neither been displaced nor rendered obsolete by generative AI. Rather, it continues to provide the conceptual framework within which courts, regulators and policymakers are addressing the legal implications of rapidly evolving AI technologies. The developments since 2024 therefore suggest not the emergence of a new copyright paradigm, but the continued evolution of established doctrine in response to unprecedented technological capabilities.

V. Where Judicial Practice Still Falls Short

Although the judicial and regulatory developments examined above largely validate the central propositions advanced in the earlier article, they also reveal that significant doctrinal and practical uncertainties remain. Importantly, these uncertainties do not suggest that existing copyright doctrine has become obsolete. Rather, they demonstrate that courts are still in the early stages of applying well established copyright principles to technologies that differ fundamentally from traditional methods of creation and reproduction. The principal challenge is therefore not whether copyright law is capable of addressing generative AI, but how its established concepts should be interpreted and applied in increasingly complex technical environments (Lucchi, 2025; Samuelson, 2023).

5.1 The Continuing Uncertainty Surrounding AI Training

Perhaps the most significant unresolved issue concerns the legal characterisation of AI training. Although recent decisions have generally rejected the proposition that training is inherently unlawful, neither have they established a universally applicable framework capable of determining when training falls within permissible copyright use and when it exceeds those limits.

The decisions in Bartz v. Anthropic and Kneschke v. LAION e.V. illustrate this uncertainty from different legal perspectives. In the United States, the court approached the issue through the doctrine of fair use, emphasising the transformative nature of AI training while simultaneously distinguishing the acquisition of pirated materials from the subsequent training process. In contrast, the Regional Court of Hamburg analysed the same activity through the text and data mining exceptions contained in the DSM Directive, recognising that acts of reproduction occur during training but may nevertheless be justified by statutory exceptions (Bartz v. Anthropic, 2025; Kneschke v. LAION e.V., 2024).

These decisions demonstrate substantial doctrinal convergence in one respect: neither court accepted the proposition that AI training is automatically infringing. At the same time, however, neither judgment establishes a comprehensive methodology for determining the legality of future training activities involving different datasets, different model architectures or different methods of data acquisition. Consequently, developers, rights holders and courts remain without clear guidance regarding the precise legal boundaries of permissible AI training.

Academic commentary reflects the same uncertainty. While many scholars regard AI training as functionally analogous to human learning, others caution that the scale, automation and commercial deployment of generative AI distinguish it from ordinary human cognitive processes (Henderson et al., 2023; Lucchi, 2025; Matulionyte, 2024; Samuelson, 2023; Stober & Dornis, 2026; U.S. Copyright Office, 2025b). The resulting debate demonstrates that the legal characterisation of AI training remains one of the most significant unresolved issues in contemporary copyright law.

5.2 Memorisation Remains the Principal Doctrinal Challenge

The earlier article distinguished between learning from existing works and reproducing protected expression. Since its publication, judicial attention has increasingly shifted towards determining whether modern AI systems merely learn statistical relationships or whether they retain copyright protected expression capable of subsequent reproduction.

This issue emerged most prominently in GEMA v. OpenAI, where the Regional Court of Munich I considered allegations that AI-generated outputs reproduced copyrighted song lyrics. Although the judgment contributes significantly to the developing jurisprudence concerning memorisation and AI-generated outputs, it does not establish a universally applicable legal or technical test for determining when memorisation has occurred or

what degree of similarity should suffice to establish infringement (GEMA v. OpenAI, 2025; Dornis, 2025; Stober & Dornis, 2026).

Similarly, *Bartz v. Anthropic* acknowledged allegations concerning memorisation while ultimately resolving the dispute primarily through fair use analysis and the evidentiary record before the court. As a result, current jurisprudence provides only limited guidance regarding one of the central questions raised by generative AI: when does statistical learning become legally significant reproduction?

The technical literature likewise demonstrates that memorisation is neither inevitable nor impossible. Research has shown that large language models may, under particular conditions, reproduce portions of their training data, especially where rare or highly repeated material is involved (Carlini et al., 2023). The legal significance of that phenomenon, however, remains unresolved. Future courts will likely be required to distinguish between isolated regurgitation, statistical pattern learning, substantial similarity and actionable reproduction through increasingly sophisticated technical evidence.

5.3 Market Substitution and the Changing Nature of Copyright Harm

A further development since the publication of the earlier article is the increasing importance attached to economic harm arising from generative AI. Early academic discussion focused primarily upon whether AI training constituted copying. More recent litigation suggests that future disputes may instead concentrate upon whether AI-generated outputs substitute for copyrighted works within existing or potential licensing markets.

This development is particularly evident in *Kadrey v. Meta Platforms, Inc.*, where the court devoted considerable attention to issues of market substitution and competitive harm. Although the plaintiffs failed to establish the evidentiary foundation necessary to defeat summary judgment, the judgment demonstrates that future copyright litigation may increasingly focus upon the commercial consequences of AI-generated outputs rather than solely upon the mechanics of model training (*Kadrey v. Meta Platforms, Inc.*, 2025).

The same concern increasingly appears within policy documents and pleadings. The European Parliament study and the United States Copyright Office have emphasised the potential effect of generative AI upon licensing markets, remuneration structures and incentives for future creative production. Recent publisher complaints, including Encyclopædia Britannica and Merriam Webster v. OpenAI, allege similar substitution and traffic diversion harms (Lucchi, 2025; U.S. Copyright Office, 2025b; Encyclopædia Britannica and Merriam Webster v. OpenAI, 2026). Whether existing copyright doctrines adequately address these forms of economic harm remains an open question that future litigation will inevitably continue to explore.

5.4 Allocation of Responsibility Within the AI Ecosystem

Another issue that remains insufficiently developed concerns the allocation of responsibility among the various participants involved in the creation and deployment of generative AI systems. Unlike traditional copyright disputes, AI-generated outputs may involve dataset providers, foundation model developers, finetuning organisations, platform operators, commercial deployers and end users, each exercising different degrees of influence over the resulting output.

Recent decisions provide only partial guidance regarding the respective responsibilities of these participants. The Chinese Ultraman decisions demonstrate one possible approach by emphasising the obligations of platform operators to implement reasonable preventive measures and exercise appropriate control over AI-generated outputs (Shanghai Character License Administrative Co., Ltd. v. AI Company, Guangzhou Internet Court, 2024; Hangzhou Intermediate People's Court, 2024). By contrast, the American and European decisions have generally resolved liability only with respect to the specific parties before the court, leaving broader questions concerning the allocation of responsibility throughout the AI lifecycle unresolved.

As generative AI systems become increasingly integrated into commercial products and professional services, clearer principles governing responsibility across the AI value chain will become essential. Existing concepts such as control, knowledge, causation and foreseeability remain applicable, but their application within distributed AI systems has yet to be comprehensively articulated.

5.5 Transparency, Evidence and the Practical Administration of Copyright

Perhaps the greatest practical obstacle confronting future AI copyright litigation concerns access to evidence. Existing copyright doctrine can operate effectively only where courts are able to determine how models were trained, what datasets were used, whether copyright reservations were respected, whether protected expression was retained, and how particular outputs were generated. In practice, much of this information remains inaccessible to copyright owners.

Recent regulatory initiatives recognise this problem. The European Union's AI Act and the accompanying General-Purpose AI Code of Practice place increasing emphasis upon transparency, technical documentation and copyright compliance obligations for providers of general purpose AI models (Regulation (EU) 2024/1689; European Commission, 2025a, 2025b; Peukert & Castets-Renard, 2025). While these measures do not themselves resolve substantive copyright questions, they acknowledge that effective enforcement increasingly depends upon improved access to technical information concerning AI systems.

Similarly, the Japanese Copyright Office has emphasised the importance of distinguishing lawful training from circumstances involving unlawful acquisition of works or the reproduction of protected expression (Japan Copyright Office, 2024). These developments demonstrate that future copyright disputes are likely to depend as much upon evidentiary transparency as upon substantive copyright doctrine.

Accordingly, one of the most important tasks facing courts will be the development of evidentiary standards capable of enabling existing copyright principles to operate effectively despite the technical complexity of contemporary AI systems.

5.6 Refinement Rather Than Replacement

Taken together, the developments since the publication of the earlier article suggest that copyright law is not experiencing a doctrinal crisis but an operational one. Courts have consistently relied upon established concepts such as authorship, originality, reproduction, fair use, statutory exceptions and infringement. What remains uncertain is the manner in which those concepts should be applied where creative processes involve probabilistic models trained on unprecedented quantities of copyrighted material.

The first generation of AI copyright litigation therefore does not demonstrate the failure of existing copyright doctrine. Rather, it demonstrates the need for greater doctrinal precision, clearer evidentiary standards and more consistent judicial guidance. The legal framework already exists; the continuing challenge lies in applying it coherently to technologies that continue to evolve more rapidly than judicial doctrine itself. This is not a deficiency of copyright law, but an inevitable consequence of attempting to apply enduring legal principles to an evolving technological landscape.

VI. What Courts Still Need to Clarify

The judicial and regulatory developments examined in this article demonstrate that copyright law has not reached an impasse. On the contrary, courts across multiple jurisdictions have consistently relied upon established copyright principles when resolving disputes involving generative AI. Nevertheless, the first generation of AI copyright litigation also reveals that significant uncertainty remains concerning the practical application

of those principles. The challenge facing courts is therefore no longer whether copyright doctrine is capable of governing generative AI, but how its established concepts should be interpreted in circumstances involving increasingly sophisticated machine learning systems (Dornis & Lucchi, 2025; Foong, 2025; Lucchi, 2025; Samuelson, 2023).

The following issues are likely to define the next stage of judicial development.

6.1 Defining the Boundary Between Learning and Reproduction

Perhaps the most fundamental issue concerns the legal boundary between learning from protected works and reproducing them. The earlier article argued that AI training should generally be understood as analogous to learning rather than copying. Recent judicial decisions have largely rejected the proposition that training is inherently unlawful, yet they have simultaneously demonstrated that the analogy alone is insufficient to resolve every copyright question.

The decisions in Bartz v. Anthropic, Kadrey v. Meta Platforms, Inc., and Kneschke v. LAION e.V. demonstrate that the legality of AI training depends upon multiple considerations extending beyond the mere existence of the training process itself. Courts have distinguished between the acquisition of training materials, the legality of intermediate reproductions, the availability of statutory exceptions, the provenance of datasets and the commercial purpose of the training activity (Bartz v. Anthropic, 2025; Kadrey v. Meta Platforms, Inc., 2025; Kneschke v. LAION e.V., 2024).

Future courts should articulate more clearly where lawful learning ends and legally significant reproduction begins. Such clarification would benefit not only copyright litigation but also licensing practices, compliance strategies and AI governance more generally. Without a clearer doctrinal boundary, both rights holders and AI developers will continue to operate under substantial legal uncertainty (Henderson et al., 2023; Lucchi, 2025; Like Company v. Google Ireland, Case C250/25).

6.2 Developing a Coherent Doctrine of Memorisation

A second issue concerns the growing importance of memorisation within AI copyright litigation. Although academic literature has demonstrated that large language models may, under certain circumstances, reproduce portions of their training data, recent judicial decisions indicate that courts have only begun addressing the legal significance of that phenomenon (Carlini et al., 2023).

The judgment in GEMA v. OpenAI represents one of the first judicial examinations of allegations that AI-generated outputs reproduced copyrighted expression retained within a model. By contrast, Bartz acknowledged arguments concerning memorisation while resolving the dispute primarily through fair use analysis. Neither decision establishes a comprehensive evidentiary framework capable of distinguishing statistical learning from legally significant reproduction (GEMA v. OpenAI, 2025; Bartz v. Anthropic, 2025; Dornis, 2025; Stober & Dornis, 2026).

Future courts will therefore need to address several interrelated questions. What constitutes sufficient evidence that a model has memorised protected expression? Should isolated instances of regurgitation be treated differently from systematic reproduction? What role should expert technical evidence play? How should courts distinguish between memorisation, substantial similarity and independent generation? Until these questions receive consistent judicial treatment, memorisation is likely to remain one of the most contested issues in AI copyright litigation.

6.3 Clarifying the Threshold of Human Creative Contribution

Although recent authorities consistently reaffirm the principle of human authorship, comparatively little guidance has been provided concerning the degree of human involvement necessary to satisfy that requirement.

Decisions such as Thaler v. Perlmutter and S. Š. v. Taubel Legal confirm that copyright does not extend to works generated autonomously by AI. They do not, however, establish clear criteria for determining when human participation in AI-assisted creative processes becomes sufficiently substantial to constitute authorship (Thaler v. Perlmutter, 2025; S. Š. v. Taubel Legal, 2023).

This issue will become increasingly important as AI tools become integrated into ordinary creative practice. Modern workflows frequently involve iterative prompting, selection among numerous generated outputs, repeated refinement, editing, combination of AI-generated and human created material, and extensive post generation modification. Future courts should clarify how these activities should be evaluated when assessing originality and authorship.

Rather than adopting rigid quantitative thresholds, judicial analysis should continue to focus upon whether the protected expression ultimately reflects meaningful human intellectual creation. Greater clarity in this area would increase legal certainty while preserving the human-centred foundations of copyright law (Fritz, 2025; U.S. Copyright Office, 2025a).

6.4 Allocating Responsibility Across the AI Ecosystem

Traditional copyright disputes generally involve relatively straightforward relationships between authors, publishers and alleged infringers. Generative AI has significantly complicated this picture. A single AI-generated output may involve dataset providers, foundation model developers, finetuning organisations, platform operators, commercial deployers and end users, each exercising different degrees of control over the resulting work.

Current case law provides only partial guidance regarding the allocation of responsibility among these participants. The Chinese Ultraman decisions illustrate one approach by emphasising the obligations of platform operators to implement reasonable safeguards against foreseeable infringement (Shanghai Character License Administrative Co., Ltd. v. AI Company, Guangzhou Internet Court, 2024; Hangzhou Intermediate People's Court, 2024). Other courts have focused primarily upon the conduct of the parties before them without attempting to establish comprehensive principles governing responsibility throughout the AI lifecycle.

Future litigation should therefore clarify how traditional copyright concepts such as causation, knowledge, control, foreseeability and contributory liability should operate within increasingly distributed AI systems. The answer is likely to have significant implications for AI developers, commercial users and rightsholders alike.

6.5 Transparency, Evidence and the Administration of Copyright

Perhaps the greatest practical challenge identified by the recent case law concerns evidence. Existing copyright doctrine can operate effectively only where parties are capable of establishing the relevant facts. In AI litigation, however, many of those facts remain inaccessible. Questions concerning the composition of training datasets, compliance with copyright reservations, implementation of technical safeguards, model architecture, fine tuning practices and output generation frequently involve information available only to AI developers.

Recent regulatory initiatives reflect increasing recognition of this evidentiary problem. The European Union's AI Act, together with the General Purpose AI Code of Practice, places growing emphasis upon transparency, technical documentation and copyright compliance obligations (Regulation (EU) 2024/1689; European Commission, 2025a, 2025b, 2025c; Peukert & Castets-Renard, 2025). Broader AI Act transparency obligations, including Article 50 disclosure duties for AI-generated or manipulated content, illustrate the same regulatory

movement towards evidentiary visibility (EU Artificial Intelligence Act, 2026). Similarly, guidance issued by the Japanese Copyright Office recognises the importance of distinguishing lawful training from circumstances involving unlawful acquisition or reproduction of protected works (Japan Copyright Office, 2024).

These developments do not replace copyright doctrine; rather, they facilitate its effective application by improving access to information relevant to copyright disputes. Future judicial development should therefore focus not only upon substantive copyright principles but also upon evidentiary standards capable of enabling courts to determine how AI systems have been trained, why particular outputs have been generated and whether legally relevant reproduction has occurred.

6.6 Refining Existing Doctrine Rather Than Creating New Doctrine

Taken together, the authorities examined throughout this article suggest that copyright law is undergoing refinement rather than transformation. The first generation of AI litigation has not demonstrated the need for a separate body of AI specific copyright law. Instead, it has shown that established principles-including authorship, originality, reproduction, infringement, fair use, statutory exceptions and liability-remain capable of resolving disputes involving generative AI.

What remains necessary is greater doctrinal precision. Courts should continue to clarify the boundary between learning and reproduction, develop coherent approaches to memorisation, articulate clearer standards for human creative contribution, refine principles governing responsibility within complex AI ecosystems and strengthen the evidentiary framework through which copyright disputes involving AI are resolved.

Accordingly, the future development of AI copyright law is likely to consist not of replacing existing doctrine, but of progressively refining its application to increasingly sophisticated technologies. The first wave of judicial and regulatory developments strongly suggests that copyright's fundamental principles remain intact. The task ahead is to ensure that those principles continue to operate coherently, consistently and predictably as generative AI becomes an increasingly pervasive feature of creative, commercial and scientific activity.

VII. Conclusion

The earlier article advanced the proposition that generative artificial intelligence, notwithstanding its unprecedented capabilities, did not require a fundamental reconstruction of copyright law. Rather, it argued that existing doctrine, properly understood and carefully applied, remained capable of addressing the legal questions raised by AI-generated content. The judicial and regulatory developments that have followed neither fully resolve those questions nor render the original framework obsolete. They do, however, provide the first meaningful opportunity to test theory against practice.

Taken together, the emerging jurisprudence reveals a striking degree of doctrinal continuity. Courts across multiple jurisdictions have consistently resisted invitations either to recognise AI systems as authors or to abandon established copyright principles in favour of technology specific rules. Instead, they have applied familiar concepts of authorship, originality, reproduction, infringement, fair use, statutory exceptions and liability to disputes involving increasingly sophisticated AI systems. The result is neither doctrinal revolution nor technological exceptionalism, but the gradual adaptation of established legal principles to novel factual circumstances.

At the same time, the litigation demonstrates that the most difficult questions have shifted. The central issues are no longer whether copyright law is capable of addressing generative AI, but how its existing principles should operate where creative processes involve probabilistic models, vast training datasets and opaque

computational architectures. Questions concerning human authorship, memorisation, reproduction, causation, control, foreseeability and evidentiary traceability have become increasingly significant precisely because they expose the practical limits of applying familiar doctrine to unfamiliar technology. The challenge is therefore one of precision rather than replacement.

This shift carries implications well beyond the courtroom. Legislators face increasing pressure to ensure that regulatory initiatives complement rather than displace copyright doctrine. Courts are called upon to develop coherent evidentiary standards capable of addressing technically complex disputes without sacrificing legal certainty. Intellectual property practitioners must navigate an evolving landscape in which traditional copyright concepts remain applicable but require increasingly sophisticated technical understanding. Rights holders seek meaningful protection against unlawful appropriation while recognising that copyright has never conferred a monopoly over learning or influence. Developers require legal certainty that permits innovation without creating unwarranted immunity from responsibility. Ordinary users, who increasingly rely upon generative AI as an everyday creative and professional tool, likewise benefit from legal rules that are principled, predictable and comprehensible.

The law's task is not to romanticise human creativity, demonise machine learning, or immunise technological actors from responsibility. Nor is it to treat every technological advance as requiring an entirely new legal architecture. Its task is more disciplined than that. It is to distinguish between learning and appropriation, influence and reproduction, autonomous generation and legally relevant control, technological novelty and doctrinal impossibility. Those distinctions have always lain at the heart of copyright law. Generative AI has not displaced them; it has made their careful application more important than ever.

Academics, judges, legislators, intellectual property professionals, technology developers and ordinary users are witnessing more than another stage in the evolution of copyright law. They are observing a profound realignment of the relationship between law and technology—a movement of tectonic plates that is reshaping long held assumptions about creativity, authorship, evidence and responsibility. Whether copyright law emerges from this period with its coherence intact will depend not upon its willingness to abandon established principles, but upon its ability to apply them with intellectual discipline to technologies that challenge conventional methods of creation without rendering those principles obsolete. In that sense, generative AI presents not merely a difficult copyright problem, but a defining test of whether legal systems can absorb technological rupture while remaining faithful to the conceptual foundations that give them legitimacy.

References

Academic and secondary sources

Bird & Bird (2024). Liability of AI Service Providers for Copyright Infringement: Guangzhou Internet Court Reaches World's First Decision. 10 April 2024.

Bommasani, Rishi, et al. (2021). On the Opportunities and Risks of Foundation Models. arXiv:2108.07258.
Carlini, Nicholas, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr and Chiyuan Zhang (2023). Quantifying Memorization Across Neural Language Models. ICLR 2023.

Dornis, Tim W. (2025). Generative AI, Reproductions Inside the Model, and the Making Available to the Public. IIC 56, 909938. https://doi.org/10.1007/s40319025015829.

Dornis, Tim W. and Nicola Lucchi (2025). Generative AI and the Scope of EU Copyright Law: A Doctrinal Analysis in Light of the Referral in Like Company v. Google. IIC. https://doi.org/10.1007/s40319025016497.

EU Artificial Intelligence Act (2026). The EU AI Act's Transparency Rules: A Practical Guide to Article 50. 14 May 2026.

Foong, Cheryl (2025). Generative Artificial Intelligence Models and Copyright Infringement: Doctrinal Challenges and Regulatory Gap-Filling Using Unfair Competition Principles. UNSW Law Journal 48(4), 13611398.

Fritz, Johannes (2025). Understanding Authorship in Artificial Intelligence-Assisted Works. Journal of Intellectual Property Law & Practice 20(5), 354364. https://doi.org/10.1093/jiplp/jpae119.

Gervais, Daniel J. (2020). The Machine as Author. Iowa Law Review 105, 20532106.

Ginsburg, Jane C. and Luke Ali Budiardjo (2019). Authors and Machines. Berkeley Technology Law Journal 34(2).

Goodfellow, Ian, Yoshua Bengio and Aaron Courville (2016). Deep Learning. MIT Press.

Henderson, Peter, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A. Lemley and Percy Liang (2023). Foundation Models and Fair Use. Journal of Machine Learning Research 24(400), 179.

Hyun, Daiwon, Sunho Park and Eutteum Kim (2025). Towards a New Copyright Paradigm for Generative AI: Bridging Human-Machine Creativity Through Legal and Policy Reform. ICCV Workshop / Computer Vision Foundation.

Lemley, Mark A. and Bryan Casey (2021). Fair Learning. Texas Law Review 99(4), 743785.

Lucchi, Nicola (2025). Generative AI and Copyright: Training, Creation, Regulation. European Parliament, Policy Department for Justice, Civil Liberties and Institutional Affairs, study requested by the Committee on Legal Affairs, PE 774.095, 9 July 2025.

Matulionyte, Rita (2024). Reconceptualising the Reproduction Right in the Age of AI. Working paper, 27 August 2024.

Peukert, Alexander and Céline Castets-Renard (2025). The Copyright Chapter of the EU's Code of Practice for General Purpose AI Models: A Commentary. IIC 56, 19501993. https://doi.org/10.1007/s40319025016648.

Samuelson, Pamela (2023). Generative AI Meets Copyright. Science 381(6654), 158161. https://doi.org/10.1126/science.adi0656.

Werzansky-Orland, Yotam (2024). AI-Generated content and the question of copyright. *The Market: International Journal of Business*, Vol. 5, 2–20.

Stober, Sebastian and Tim W. Dornis (2026). Generative AI Training and Copyright Law. Transactions of the International Society for Music Information Retrieval, forthcoming; arXiv:2502.15858v2.

White & Case LLP (2026). Japan's First AI Legislation Becomes Law Focus Is on Promoting Research and Development; No Monetary Penalties. 14 April 2026.

Cases and judicial materials

Bartz v. Anthropic PBC, No. C 2405417 WHA, Order on Fair Use (United States District Court for the Northern District of California, 23 June 2025).

Encyclopædia Britannica, Inc. and Merriam Webster, Inc. v. OpenAI, Inc. et al., Complaint, No. 1:26cv02097UA (United States District Court for the Southern District of New York, filed 13 March 2026).

Getty Images (US), Inc. v. Stability AI Ltd [2025] EWHC 2863 (Ch) (High Court of Justice, Business and Property Courts of England and Wales, 4 November 2025).

Shanghai Character License Administrative Co., Ltd. v. AI Company / 上海新创华文化发展有限公司与AI公司网络侵权责任纠纷案, Civil Judgment, (2024) Yue 0192 Min Chu No. 113 / (2024) 粤0192民初113号 (Guangzhou Internet Court, 2024).

Shanghai [Culture Development Company] v. Hangzhou [AI Technology Company] / 上海某文化发展有限公司诉杭州某智能科技有限公司著作权侵权及不正当竞争纠纷案, Civil Judgment, (2024) Zhe 0192 Min Chu No. 1587 / (2024) 浙0192民初1587号 (Hangzhou Internet Court, 2024)

Shanghai [Culture Development Company] v. Hangzhou [AI Technology Company], Civil Judgment, (2024) Zhe 01 Min Zhong No. 10332 / (2024) 浙01民终10332号 (Hangzhou Intermediate People's Court, 2024).

Kadrey et al. v. Meta Platforms, Inc., No. 23cv03417VC, Order Denying Plaintiffs' Motion for Partial Summary Judgment and Granting Meta's Cross-Motion for Partial Summary Judgment (United States District Court for the Northern District of California, 25 June 2025).

Kneschke v. LAION e.V., LG Hamburg, Judgment of 27 September 2024, 310 O 227/23, openJur 2024, 9199.
GEMA v. OpenAI, LG München I, Endurteil v. 11 November 2025, 42 O 14139/24, BeckRS 2025, 30204; GRUR 2025, 1917; NJW 2026, 183.

Like Company v. Google Ireland Ltd, Case C250/25, request for a preliminary ruling lodged before the Court of Justice of the European Union on 3 April 2025.

S. Š. v. TAUBEL LEGAL, Municipal Court in Prague, No. 10 C 13/202316, 11 October 2023.
Thaler v. Perlmutter, No. 235233 (United States Court of Appeals for the District of Columbia Circuit, 18 March 2025).

Legislation, regulatory materials, and official reports

Agency for Cultural Affairs / Japan Copyright Office (2024). General Understanding on AI and Copyright in Japan: Overview. May 2024.

European Commission (2025a). Code of Practice for General Purpose AI Models: Copyright Chapter.

European Commission (2025b). Code of Practice for General Purpose AI Models: Transparency Chapter.

European Commission (2025c). Code of Practice for General Purpose AI Models: Safety and Security Chapter.

European Parliament (2024). Artificial Intelligence Act: European Parliament Legislative Resolution of 13 March 2024, P9_TA(2024)0138.
European Union (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L 2024/1689, 12 July 2024.

Italy (2025). Legge 23 settembre 2025, n. 132, Disposizioni e deleghe al Governo in materia di intelligenza artificiale, Gazzetta Ufficiale n. 223, 25 September 2025.

Japan (2025). Act on Promotion of Research and Development and Utilization of Artificial Intelligence-related Technologies, passed 28 May 2025.

U.S. Copyright Office (2025a). Copyright and Artificial Intelligence, Part 2: Copyrightability. Report of the Register of Copyrights, January 2025.

U.S. Copyright Office (2025b). Copyright and Artificial Intelligence, Part 3: Generative AI Training. Prepublication Version, May 2025.

When the Machine Keeps Its Hands Clean: AI, The Illusion of Purity and the Problem of Dirty Hands

By

Dr Demetris Tillyris

Director of Postgraduate Studies, CIM-Cyprus Business School

There is something reassuring about the prospect of entrusting difficult decisions to a machine. Machines process information swiftly, identify patterns and provide us with an answer; and they perform these functions, without hesitation, without harbouring resentments or seeking approval, and without losing sleep over the consequences of an unpopular decision. In an age in which organisations are increasingly expected to make decisions more expeditiously, and on the basis of ever larger quantities of information, the attraction of artificial intelligence (AI) is hardly surprising. AI promises not just to automate tasks that we find tedious, but to augment capacities traditionally regarded as distinctively human: judgement, practical reasoning, and the ability to make informed predictions (Raisch and Krakowski, 2021). Used well, AI may even improve ethical decision-making by exposing biases and inconsistencies that so often evade human attention (De Cremer and Narayanan, 2023). To wholeheartedly dismiss AI by appeal to the familiar anxieties about machines replacing human judgement would thus be asinine. The more puzzling question, however, is what happens when machines become quite capable at helping us decide.

The trouble is that more information does not necessarily produce better judgement. Indeed, there are circumstances in which the technologies through which we acquire information might well alter what we deem relevant, and what we might have otherwise regarded as important in the first place.

The 2008 financial crisis offers an instructive example. Donald MacKenzie's work on the sociology of finance shows how financial markets are constituted not simply by human actors exchanging assets, but via intricate interactions between people, models, technologies, institutions and representations (MacKenzie, 2008; 2011). Financial instruments increasingly came to be encountered through screens, prices, mathematical models and quantitative measures of risk. The significance of such a transformation lies not in the fact that traders somehow ceased to be aware that real people stood behind the mortgages and securities traded; rather, technical representations altered the relationship between the decision-maker and the object of the decision. For instance, a collection of households became a mere portfolio, and a complex social reality was reduced to a number on a screen. MacKenzie's analysis of the credit crisis illustrates how particular forms of financial knowledge and valuation rendered complex financial instruments intelligible on the one hand, whilst, at the same time, contributing to the conditions in which the risks associated with such instruments were systematically misunderstood (MacKenzie, 2008; 2011).

The point here is not that technology caused the crisis, or that human judgement would have sufficed in its absence. Technologies of decision-making do not merely supply us with more information; they shape how that information is seen, rendering some consequences easier to calculate, while making others less salient. The problem, however, is that, in matters of leadership, a decision can become technically clearer, but without necessarily becoming morally easier.

It is at this point that the problem of dirty hands (DH) becomes relevant – a longstanding problem in political theory (see Nagel, 1972; Walzer, 1973; de Wijze, 1994; Tillyris, 2014; 2016a; 2016b; 2017; Nick and de Wijze, 2023), but one that has also recently received attention in business ethics (Badaracco, 1997; Brenkert, 2009; Monge, 2015). In his ground-breaking essay in the field titled *Political Action: The Problem of Dirty Hands*, Michael Walzer puts forward a disquieting recognition that flies in the face of most ethical frameworks on offer: those in positions of power might, in certain circumstances, be "forced to choose between two courses of action both of which it would be wrong for [them] to undertake"; that, in certain cases, doing what is politically necessary might well entail doing something that is morally wrong (Walzer, 1973: 160). In such instances, the problem need not emanate from the lack of sufficient information; nor, by implication, need it be resolved cleanly, by merely pouring in more information in support of better decision-making. For, the core insight of the DH problem, is that there may be no course of action that allows those imbued with power to honour all of the obligations they have (Stocker, 1990; de Wijze, 1994; Tillyris, 2014). As Bernard Williams, in his discussion of politics and moral character, notes, the difficulty concerns not simply identifying the overall

right course of action, but what it means for an agent to retain a modicum of integrity, when all available options involve morally disagreeable consequences (Williams, 1981).

Taken seriously, the DH problem provides a useful lens through which to think about AI and organisational decision-making. The issue is not whether an algorithm can identify the most efficient course of action, or even whether it can correctly predict the consequences of different choices. Rather, the problem is that a perfectly informed decision can remain morally troubling because the trouble lies in deciding which competing obligations and values ought to prevail (and, what follows for those values and obligations which have been suspended). AI may thus solve an epistemic problem—i.e., it might help us to better appreciate what is likely to happen under certain scenarios—without resolving the normative problem of what should be done when every available course of action involves some form of moral loss.

Though early analyses of the DH problem have tended to focus on those who wield and exercise political power, more recent work on this phenomenon, including my own (see Stocker, 1990; de Wijze, 1994; Gowans, 2001; Tillyris, 2019b; 2023a), has sought to extend this issue beyond the narrow confines of political life. Consider, for instance, a company the most important division of which has become financially unsustainable. An AI system analyses its financial position, productivity, market conditions and future prospects and concludes that closing the division offers the best prospect of preserving the organisation as a whole. Doing so would save hundreds of other jobs and help protect a company faced with a serious risk of collapse. Yet, such a decision would not be without loss: for, it would also fail to protect some of the employees who work there, renege on commitments already made to them, and sacrifice an important part of the organisation's identity. Suppose the algorithm is right. Suppose, moreover, that there is no hidden error, no discriminatory variable and no obvious alternative. The moral difficulty does not evaporate. The leader must still decide whether the survival of the organisation justifies sacrificing those whose livelihoods depend upon the division; whether an obligation to the employees concerned can be overridden by obligations to everyone else; and, whether a commitment made under different circumstances remains binding, even when circumstances have changed. The model can estimate the consequences of each course of action with surgical precision. It cannot, however, wish the moral loss away. Hence, the decision may therefore be entirely defensible; but it will not be morally clean.

Taken together, MacKenzie's insights from the sociology of finance, and the DH problem issue a concern about AI proliferation which is different from the rather familiar by now concerns pertaining to algorithmic bias or technological control: AI may make a morally momentous decision appear to be a mere technicality, effacing its moral complexity. To be sure, AI is not itself culpable for creating the underlying problem. Organisations have long relied upon models, metrics, consultants, experts, public-relations professionals and other technologies of decision-making to render difficult choices more calculable and defensible. The significance of AI lies partly in its capacity to intensify and integrate these existing tendencies. What previously required several distinct stages - calculation, interpretation, justification and communication - and, quite often, several distinct actors, can now increasingly be brought together through a single system. The danger, therefore, is not that AI somehow introduces moral ambiguity into organisational life, but, rather, that it may amplify existing forms of moral simplicity.

This is not a trivial matter for questions of leadership. Leaders are increasingly encouraged to make decisions on the basis of evidence, metrics and measurable outcomes. There is nothing inherently objectionable about this. The difficulty arises when the availability of better calculations encourages us to forget that one of the most crucial aspects of any decision often resides outside the calculation itself. Someone must still decide what the organisation stands for, whose interests matter, what constitutes an acceptable loss and which obligations can legitimately be cast aside when they collide. The machine may perform the calculation. The leader still owns the judgement, and the onus of the responsibility which accompanies it (Gaczek et al., 2026).

There is, however, another reason why the seductive image and language of clean decision-making which a thoughtless utilisation of AI might perpetuate, is potentially misleading. The problem of moral decision making under uncertainty and ambiguity, does not cease when the decision has been made. Responsible leadership also requires leaders to explain, justify and communicate what they have done.

It is at this point, in which another uncomfortable feature of political and organisational life becomes vivid: the possibility or, perhaps even, the inevitability of hypocrisy and deception. We tend to regard the gap between what we believe, what we say and what we do as straightforward evidence of moral failure. Yet, leaders rarely speak to a single, homogeneous audience; nor, is leadership exercised under conditions in which one can say exactly what one believes, fully honour what one says, and satisfy every constituency to which one is accountable (Tillyris, 2016a; 2017). Leaders must negotiate competing expectations, preserve relationships, accommodate disagreement and, at times, communicate differently to differently audiences. Some degree of disjuncture between what one believes, what one says and what one does may be an inescapable feature of acting within complex political and organisational landscapes.

This is also where my most recent work, which focuses on political communication and dog-whistling, becomes relevant (Tillyris, 2024). Dog-whistling illustrates, in particularly stark form, how a single piece of communication, verbal or non-verbal, can acquire different meanings for different audiences - groups with different worldviews or aspirations. Otherwise put, dog-whistling, often labelled as a coded communication, entails the dissemination of messages which may appear innocuous or ambiguous to a group, while surreptitiously conveying a more specific meaning to others. Such a phenomenon is politically troubling, not just because it entails, by virtue of its very nature, aspects of deception and dissimulation, but, rather, because it also complicates the relationship between intention, communication and responsibility. Yet, neither dog-whistling, nor its underlying ethical ambiguity ought to be seen as confined to electoral politics. For, organisations, too, communicate simultaneously with employees, investors, customers, regulators and the wider public, whose interests and expectations may not coincide. For instance, multivocal appeals are regularly observed in the advertising literature. As Brumbaugh argues targeted advertisements that include "both similar subculture sources and other subculture relevant cues... induce members of subcultures to draw on this knowledge to allow them to create meaning that others outside this subculture cannot" (2002: 259). The rather perverse nature of such a phenomenon directs our attention to a more general problem of leadership: communication does not always have a single, transparent meaning, and the responsibility of the communicator cannot be understood simply by asking whether a statement corresponds neatly to an underlying intention. What matters is also how messages are received, what they enable different audiences to understand, and whether the management of these differences can be justified. The question, therefore, is not simply whether leaders communicate consistently, but how, and under what circumstances, the management of different audiences can be justified as a legitimate feature of leadership rather than condemned simply because it involves ambiguity, inconsistency or even deception.

AI renders this distinction even more important. A leader can now ask a system to produce a message that sounds empathetic and authentic, tailored carefully to the language of employees, investors, customers or the public, whilst pre-empting legitimate objections, and softening difficult news. An expedient utilisation of AI could produce what could appear as earnest explanation and justification of almost any decision. The problem is not simply that such communication might be deceptive. It is that the technology makes it increasingly easy to perform, what we might term, the virtues of leadership – even those which are morally questionable. This recognition captures a disquieting possibility. The easier it becomes to manufacture the appearance of ethical leadership, the harder it may become to distinguish ethical leadership from its mere performance. An organisation may become exceptionally good at sounding concerned, without becoming any more responsible for the consequences of its decisions.

None of this should lead us to reject AI. Nor should it encourage nostalgia for a supposedly purer form of human decision-making. There is no reason to romanticise human judgement. It is partial, fallible and frequently self-interested. There is, however, an important distinction between improving judgement and escaping it altogether. As noted earlier, one of the defining problems of leadership is not simply that leaders sometimes lack sufficient information; rather, it is that the world that leaders inhabit and act, is characterised by competing goods, values and priorities that cannot always be reconciled in a perfect and harmonious whole. The romantic aspiration to eliminate these tensions is understandable and, - though recently reinvigorated with the proliferation of AI -, is hardy new; it goes back to the very beginnings of social and political theory (see Berlin, 1990; Tillyris and Edyvane, 2019c). But such an aspiration rests upon a mistaken picture of the conditions under which practical judgement takes place. AI can help us calculate, predict and compare. It can help leaders see more clearly. What it cannot do is decide which sacrifices are acceptable, which obligations should prevail, or what the organisation ought ultimately to stand for. That remains a burden of responsible leadership. And perhaps that is precisely why such a burden matters. The danger is not that artificial intelligence will acquire dirty hands, but that we will use artificial intelligence to convince ourselves that ours are clean.

References

Badaracco, J. L. (1997) *Defining Moments: When Managers Must Choose Between Right and Right*. Boston, MA: Harvard Business School Press.

Berlin, I. (1990) 'The Pursuit of the Ideal', *New York Review of Books*, 37(19), pp. 11–18.

Brumbaugh, A. M. (2002). Source and nonsource cues in advertising and their effects on the activation of cultural and subcultural knowledge on the route to persuasion. *Journal of Consumer Research*, 29(2), pp. 258–269.

De Cremer, D. and Narayanan, J. (2023) 'How AI Tools Can—and Cannot—Help Organizations Become More Ethical', *Frontiers in Artificial Intelligence*, 6, 1093712.

de Wijze, S. (1994) 'Dirty Hands: Doing Wrong to Do Right', *South African Journal of Philosophy*, 13(1), pp. 27–33.

Gaczek, P., Leszczyński, G., Wei, Y. and Sun, H. (2026) 'The Bright Side of AI in Marketing Decisions: Collaboration with Algorithms Prevents Managers from Violating Ethical Norms', *Journal of Business Ethics*, 204, pp. 771–794.

MacKenzie, D. (2008) *Material Markets: How Economic Agents Are Constructed*. Oxford: Oxford University Press.

MacKenzie, D. (2011) 'The Credit Crisis as a Problem in the Sociology of Knowledge', *American Journal of Sociology*, 116(6), pp. 1778–1841.

MacKenzie, D. (2012) 'Knowledge Production in Financial Markets: Credit Default Swaps, the ABX and the Subprime Crisis', *Economy and Society*, 41(3), pp. 335–359.

Monge, R. (2015) 'Institutionally driven moral conflicts and managerial action: Dirty hands or permissible complicity?', *Journal of Business Ethics*, 129(1), pp. 161–175.

Nagel, T. (1972) 'War and Massacre', *Philosophy & Public Affairs*, 1(2), pp. 123–144.

Nick, C. and de Wijze, S. (2023) '50 Years of Dirty Hands: An Overview', *The Journal of Ethics*, 27(4), pp. 415–439.

Raisch, S. and Krakowski, S. (2021) 'Artificial Intelligence and Management: The Automation–Augmentation Paradox', *Academy of Management Review*, 46(1), pp. 192–210.

Stocker, M. (1990) *Plural and Conflicting Values*. Oxford: Clarendon Press.

Tillyris, D. (2014) '"Learning How Not to Be Good": Machiavelli and the Standard Dirty Hands Thesis', *Ethical Theory and Moral Practice*, 18(1), pp. 61–74.

Tillyris, D. (2016a) 'After the Standard Dirty Hands Thesis: Towards a Dynamic Account of Dirty Hands in Politics', *Ethical Theory and Moral Practice*, 19(1), pp. 161–175.

Tillyris, D. (2016b) 'The Virtue of Vice: A Defence of Hypocrisy in Democratic Politics', *Contemporary Politics*, 22(1), pp. 1–19.

Tillyris, D. (2017) 'Political Integrity and Dirty Hands: Compromise and the Ambiguities of Betrayal', *Res Publica*, 23(4), pp. 475–494.

Tillyris, D. (2019a) 'Dirty Hands and Suffering', *Theoria*, 66(3), pp. 95–121.

Tillyris, D. (2019b) 'Political Realism and Dirty Hands: Value Pluralism, Moral Conflict and Public Ethics', *Philosophia*, 47(5), pp. 1579–1602.

Tillyris, D. and Edyvane, D. (2019c) 'Value Pluralism and Public Ethics', *Theoria*, 66 (160), pp. 1–8.

Tillyris, D. (2023a) 'Dirty Hands as a 'Weapon of the Weak': 'Heroism', 'Aristocratism', and the Ambiguities of Everyday Resistance', *The Journal of Ethics*, 27, pp. 601–623.

Tillyris, D. (2024) 'Dog-whistling and democracy', *The British Journal of Politics and International Relations*, 26(4), pp. 1015–1032.

Vogler, G. and Tillyris, D. (2021) 'Arendt and Political Realism: Towards a Realist Account of Political Judgement', *Critical Review of International Social and Political Philosophy*, 24(6), pp. 821–844.

Walzer, M. (1973) 'Political Action: The Problem of Dirty Hands', *Philosophy & Public Affairs*, 2(2), pp. 160–180.

Williams, B. (1981) "Politics and Moral Character." In *Moral Luck: Philosophical Papers 1973–1980*, pp. 54–70. Cambridge: Cambridge University Press.

Leave AI, End Your Love-Hate Relationship!

By

Dr Maryam Kazemi Manesh

Senior Project Manager

Abstract

Artificial Intelligence (AI) continues to dominate business discussions, which oscillate between fascination and fear. For small and medium-sized enterprises (SMEs), the promise of AI-driven transformation often collides with limited resources, unclear strategies and unreliable data. This article argues that the real foundation for successful AI adoption is not technology itself, but data quality, organisational competence and readiness. Drawing on practical insights and examples from various industries, it outlines a step-by-step roadmap for SMEs to move from hype to genuine value creation. The central message is, before pursuing AI, businesses must strengthen their data infrastructure, cultivate analytical literacy, and adopt a measured, competence-based approach to innovation.

Keywords: Artificial Intelligence (AI), Analytics, Analytical Literacy, Small and Medium-sized Enterprises (SMEs), Innovation, Business Strategy, AI Readiness, Data Strategy and Governance, Data Quality.

Introduction

If you—the reader[1]—find yourself both fascinated and yet secretly frightened by AI, you're not alone. It's everywhere, from boardrooms to small-business meetups. The question echoing in every hallway seems to be: *"How can we apply AI to our business?"*

Every day brings new headlines about AI breakthroughs, in healthcare and medicine, manufacturing, transportation and autonomous vehicles, finance and risk management, and even art and literature. Yet right alongside the excitement runs a current of anxiety: job losses to automation, algorithms taking over decision-making, and data being processed by machines you can "barely understand and trust".

For SMEs, the fear of missing out (FOMO) is real. Competitors seem to jump on new tools every month, while vendors promise "plug-and-play AI" that will magically transform everything overnight. It's no wonder that many business owners feel torn; overwhelmed by the AI express train on one side, and uncertain where to start (or whom to trust) on the other. For SMEs especially, who have neither the budget for endless experimentation nor the luxury of failure in a competitive market, it feels as helpless as a hate-love relationship!

Step One: A Reality Check

A study by the U.S. Chamber of Commerce's Technology Engagement Center in 2024 reported that 98% of small businesses already use some form of AI [1]. That sounds impressive, but further research brings more clarity: the majority are simply using tools like ChatGPT as a "glorified internet search" or employing integrated AI features inside software they already had [2]. In reality, few are doing anything truly new or transformative.

In Cyprus, as *Cyprus Mail* reports, SMEs remain the backbone of the economy [3]. With tight budgets typical for SMEs, every euro counts. And when everyone around you seems to be doing something with AI, the temptation to "just try something" can be strong.

The heartfelt advice of the author of this article is as follows: don't rush into this relationship! Leave AI (for now)! The breakup statistics are seriously brutal! Would you marry someone just because all your friends are getting married? Probably not. According to MIT, 95% of enterprise AI initiatives fail [4]. Now that's a divorce rate worth avoiding!

[1] **Note:** The reader(s) is referred to as "you" in this article, i.e. the "reader" and "you" are used interchangeably.

Step Two: Get Honest About Your Data

If you're already asking the MIT studies, *"How did the 5% succeed?"*, you're thinking in the right direction. Those successful organisations had at least one thing in common: solid, reliable data.

Because here's the truth: AI is not magic. It doesn't invent insights; it amplifies what's already in your data. If that data is incomplete, fragmented or inaccurate, AI will only multiply the confusion. The old saying "garbage in, garbage out" applies here more than ever.

So, before you even whisper the word "AI," ask yourself a much more important question: *What is the state of our data?*

Everything, from client interactions and market trends to operational processes, generates valuable data. That data is your **asset.** Thus, the next questions follow naturally:

- Are your business decisions truly data-driven?
- Is your data centralised, secure and accessible to those who need it?
- Can your team read the story that your data is already telling you?

If the answer to any of these is "not yet," congratulations! You've just found your first real innovation project!

Step Three: Build Competence for AI

You are the expert in your profession. When your business is growing, you need further resources to deliver all your requests. Would you then employ someone whose expertise you cannot evaluate or measure in any way known to you?

The relationship with innovations, especially with AI, can only work out similarly. Your expert team, or a third-party on your behalf, must know what AI is producing. If AI is producing hallucinations, mistakes or injustices, the competence to reject and correct must be in place. Only then can you rely on results and boost your business with new technologies which can be successfully integrated into your processes. Implementing new machines doesn't take the responsibility away from humans [5] but facilitates the tasks, given the following five prerequisites [6].

The Five Prerequisites Before AI

1. **Identify** the data you already have (sales, customers, operations, compliance, etc.).
2. **Evaluate** its quality. Is it accurate, complete and timely?
3. **Centralise** and store it securely for easy access and analysis.
4. **Start small** with descriptive analytics: understand what has happened before predicting what *will* happen.
5. **Build competence** for interpreting the predictions.

Across years of leading implementation projects, the author of this article has seen it too often. Clients realising mid-way through that they don't actually have the data they need! Small businesses often rely on manual workarounds for practicality, which means key information never gets recorded in structured systems. Sometimes even large enterprises fall into this trap.

The bold question is: if your data has blind spots, what do you need AI for?

Reliable data alone already improves decision-making. You don't need AI to start making smarter moves. Just descriptive analysis and a clear view of your operations will take you a long way. Once that foundation is in place, you can gradually move toward predictive and prescriptive models powered by AI, which require also the competence for the justifying and interpreting of AI results.

Practical Use Cases: Grounded Before Glamorous

Next, we explore some examples across industries where better data and simple descriptive analytics [not AI] unlock real value. The following examples are perfectly applicable for SMEs. They are simple and require no sophisticated tools or systems, once you have the data. By applying a Microsoft Office tool or possibly low-priced software as a service (SaaS) for analytics or visualisation, or using embedded apps in websites you can already achieve fundamentally valuable results. The examples are chosen for the most relevant industries for Cyprus. Additionally, the suggested approach for an industry could be transferred across industries.

1. Retail and E-Commerce

Before investing in AI-driven recommendation engines, structure your transaction data:

- Which products sell most often?
- When do you see peaks and dips? Do you see seasonal effects or similar particular market dynamics?
- How do different customer segments behave?

Answering these with simple dashboards can already optimise inventory, marketing and pricing. Only *then* will AI personalisation make sense.

2. Construction and Engineering

In many construction firms, data still lives in spreadsheets or handwritten notes. Before implementing AI-based project forecasting, start by consolidating and digitising project data, e.g. labour hours, material use, equipment availability, subcontractor performance and/or delay reasons. You might be surprised that, for instance, missing clear communication and planning with subcontractors leads to delays, not necessarily externalities like the weather. Later, AI can help with predictive scheduling and resource optimisation.

3. Manufacturing and Supply Chain

Machine sensors and production logs are goldmines which are usually underused. Just by analysing downtime records, output rates and maintenance logs, inefficiencies become visible. With clean, structured data, AI can later predict failures and optimise production schedules.

4. Professional Services

Law firms, consultancies and accounting practices can review their project timelines, costs and outcomes to better allocate staff and improve profitability. Once the basics are in place, AI tools can help analyse contracts, simulate financial scenarios and evaluate financial reports.

5. Healthcare and Wellness Providers

Digitising and standardising patient and appointment records improves efficiency and compliance. Because of the high sensitivity of such information, data privacy and IT-security must be top priority. Later, AI can support triage or early diagnosis, but only if the data foundation is trustworthy.

6. Hospitality and Tourism

Hotels and restaurants can first analyse reservation patterns, customer feedback and seasonal cycles. With that, they can optimise staffing, prices and services. Only then should they consider AI assistants for personalised recommendations.

Focus on Efficiency and Competence, Not Hype

Don't misunderstand the author please. Innovation is inevitable and a great joy, despite all the risks, in the author's personal opinion! But business in general and especially SMEs thrive when they focus on what really matters, i.e. efficiency, customer satisfaction and the smart use of limited resources. The real competitive edge lies in strengthening data-driven decision-making and not trend-driven spending. Starting with data will:

- Increase operational efficiency
- Free employees from repetitive tasks
- Reallocate talent to creative, human-centered activities.

AI should be introduced gradually, as a helpful *colleague*, not a hasty *replacement*. Done right, it supports your staff, enhances customer trust and safeguards your budget.

A Step-by-Step Roadmap for SMEs

The author assumes you have already experienced many changes and integrated many innovations in business, your professional career and even your private life. AI is just another innovation with both advantages and disadvantages. From best practices and the author's experiences of integrating new innovations and change management, below is a summarised, step-by-step roadmap, also well-tailored for SMEs:

1. **Awareness:** Recognise that data is an asset, not a byproduct.
2. **Culture and Competence:** Steer your business culture and lead competence towards a data-driven approach and strategy setup.
3. **Collection:** Centralise scattered information digitally.
4. **Quality:** Clean, validate and standardise data.
5. **Visualisation:** Make insights visible with dashboards.
6. **Decision-Making:** Use data-driven insights to guide strategy.
7. **AI Readiness:** Once the data foundation is strong, explore AI use cases aligned with your business goals.

Conclusion: *Final Thoughts*

The love-hate relationship with AI is understandable. One day it's the hero promising to cut costs and boost profits; the next day it's the villain taking jobs and spreading chaos. But the truth is simpler: AI is neither hero nor villain.

AI is not a shortcut. AI is an amplifier. And it can only amplify the value of the data you already own.

So, before you chase the next shiny AI promise, take a step back. Organise, clean and value your data. Then, before you curse AI for hallucination and injustice, learn the real story deep within your data and business thoroughly, because for SMEs **the smartest way to prepare for the future is to understand the story your data is already telling you.**

References

[1] U.S. Chamber of Commerce (2024) "New Study Reveals Nearly All U.S. Small Businesses Leverage AI-Enabled Tools, Warns Proposed Regulations Could Hinder Growth". Available at: https://www.uschamber.com/technology/artificial-intelligence/new-study-reveals-nearly-all-u-s-small-businesses-leverage-ai-enabled-tools-warns-proposed-regulations-could-hinder-growth (Accessed: 29 October 2025).

[2] Kochkodin, B. (2025) "How Small Businesses Are Really Using AI", *Forbes*. Available at: https://www.forbes.com/sites/brandonkochkodin/2025/09/19/how-small-businesses-are-really-using-ai/ (Accessed: 29 October 2025).

[3] Psara, S. (2025) "Small businesses keep Cyprus economy humming", *Cyprus Mail*. Available at: https://cyprus-mail.com/2025/09/11/small-businesses-keep-cyprus-economy-humming (Accessed: 29 October 2025).

[4] Challapally, A., Pease, C., Raskar, R. and Chari, P. (2025) "The GenAI Divide: State of AI in Business 2025". Available at: https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf (Accessed: 29 October 2025).

[5] Madiega, T. (2019) "EU guidelines on ethics in artificial intelligence: Context and implementation", European Parliament. Available at: https://www.europarl.europa.eu/RegData/etudes/BRIE/2019/640163/EPRS_BRI(2019)640163_EN.pdf (Accessed: 29 October 2025).

[6] European Union (2024) "Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)". Available at: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng (Accessed: 29 October 2025).

Between dependence and autonomy: Commercialising AI Sovereignty in the European Union, United Kingdom, Canada and India, amid US–China technological competition

By

Ioannis Sidiropoulos

Lecturer in commercial law, CIM-Cyprus Business School

and PhD candidate in AI Governance, European University Cyprus

Abstract

As artificial intelligence (AI) moves from a specialist research domain into the commanding heights of economic competitiveness, public administration, critical infrastructure and defence, states are increasingly framing "AI sovereignty" as a strategic objective. Yet sovereignty in AI cannot be reduced to regulatory control, data localisation or public ownership of compute alone. A state may regulate AI extensively and subsidise domestic research, yet still remain structurally dependent if value capture, deployment and scale are mediated by foreign hyperscalers, foreign model providers or foreign chip supply chains.

This paper argues that the neglected commercialisation dimension is therefore central to AI sovereignty. A sovereign AI strategy must convert domestic research, data, talent, infrastructure and regulatory capacity into viable domestic markets and internationally competitive firms. The article first situates AI sovereignty within the wider literature on digital sovereignty and strategic autonomy and also clarifies the difference between regulatory, infrastructural, market and security sovereignty. It then shows how the European Union (EU), the United Kingdom (UK), Canada and India are increasingly linking compute, procurement, startup finance, industrial policy and sectoral deployment in an effort to build commercially sustainable sovereign capability. Brief reference to the United States, China and Singapore helps illuminate alternative institutional models.

The paper then examines the legal tensions this strategy generates between openness and resilience, competition and national champion policies, innovation and public values, domestic value capture and global legitimacy, and autonomy, power projection and security. It concludes by proposing a framework for commercially driven sovereign AI development in which the law does not merely constrain AI risks, but actively structures domestic value creation, market access, interoperability and strategic retention of control across the AI stack (Roberts, 2024).

Keywords: AI sovereignty; commercialisation; strategic autonomy; industrial policy; public procurement; managed interdependence.

Introduction

Recent scholarship makes clear that "AI sovereignty" is not a neutral or self-explanatory term. Roberts argues that digital sovereignty should be understood not as mere factual control, but as a normative question of legitimate authority; Mügge shows that AI sovereignty masks concrete trade-offs about whose interests are being served, whether sovereignty is directed against foreign powers or private gatekeepers, and whether the goal is competitiveness or emancipation from competitive logics. Further, Chesterman situates contemporary AI politics within a broader transfer of effective power from public institutions to private technology conglomerates (Roberts, 2024; Mügge, 2024; Chesterman, 2026). Taken together, this literature suggests that sovereignty in AI is not simply about who writes the rules, but also about who controls infrastructure, captures profits, determines market access and shapes the terms on which foreign technology can be adopted without translating into structural dependence.

For legal analysis, it is useful to disaggregate sovereignty in AI into at least four dimensions. First, there is regulatory sovereignty: the ability to set binding legal rules for AI systems and models in one's jurisdiction. Secondly, there is infrastructural sovereignty: access to compute, cloud, semiconductors, energy and specialised research resources. Thirdly, there is market sovereignty: the capacity of domestic firms and public institutions to develop, deploy and scale AI products in strategically significant sectors rather than merely consume tools built and monetised elsewhere. Fourthly, there is security sovereignty: the protection of sensitive data, model weights, critical infrastructure and strategic capabilities against coercive dependence, hostile access or extraterritorial constraint. The present article focuses on market sovereignty because it is the least theorised

in current debates and the most practically determinative of whether sovereign capability can survive beyond subsidy or regulatory symbolism (Roberts, 2024).

The law-and-policy debate, in contrast, has tended to overemphasise regulation and underemphasise commercial structure. The EU AI Act, for example, was enacted to improve the functioning of the internal market through a harmonised legal framework for the development, placing on the market and use of AI systems, while promoting the uptake of human-centric and trustworthy AI. That is a major constitutional and market-ordering achievement. But regulation alone cannot sustain sovereignty if domestic developers lack affordable compute, if promising firms must outsource deployment and distribution to foreign cloud providers, if domestic applications are trained on foreign-controlled foundation models under restrictive commercial terms, or if public procurement channels most demand to external suppliers. In such circumstances, compliance may be sovereign while the market remains dependent (Regulation (EU) 2024/1689).

The more precise question, therefore, is not whether states or markets should lead AI development, but how law structures the transition from strategic inputs to durable domestic value creation. Commercialisation should be understood as a legal-design problem. The relevant issue is whether firms and public institutions can obtain lawful access to data, protect or license intellectual property, secure compute at viable cost, interoperate with standards, comply without disproportionate burden, sell to anchor customers, attract finance and avoid exclusion by dominant incumbents. Where one or more of these links is missing, "sovereignty" collapses into dependence on grants, protected domestic demand or foreign platform intermediation.

A law-and-political-economy perspective also shows why AI sovereignty cannot be reduced to data localisation or state ownership. In some contexts, sovereignty may be exercised through trusted openness rather than closure. A polity may preserve strategic autonomy not by eliminating cross-border interdependence, but by ensuring that cross-border flows occur on terms compatible with domestic governance, bargaining power and public values. That insight is particularly important for medium-sized and trade-dependent jurisdictions, for which total stack autonomy is neither attainable nor desirable. Commercial sovereignty is therefore not synonymous with autarky. It is better understood as the legally and institutionally supported capacity of a polity to choose dependencies, bargain from strength and retain a meaningful share of value creation in strategically important AI layers.

The critical implication is that AI sovereignty is increasingly a problem of market design. The relevant legal toolkit extends well beyond horizontal AI regulation. It includes state aid and industrial policy, procurement law, data access regulation, copyright and licensing rules, competition law, foreign investment screening, export controls, public-private partnership design, interoperability rules, assurance standards and the governance conditions attached to sovereign compute infrastructure. The central commercialisation question is whether these instruments can convert public inputs into enduring domestic productive capacity rather than one-off subsidised experimentation (OECD, 2023).

From regulation to commercial sovereignty

A useful starting point is the Organisation for Economic Co-operation and Development's (OECD's) distinction between national AI compute capacity, effectiveness and resilience. The OECD's blueprint for national AI compute capacity argues that states need not only infrastructure availability, but also policies concerning people, innovation, access, security, sovereignty and sustainability. This matters because sovereignty fails whenever infrastructure exists without effective access, or access exists without domestic firms being able to monetise it. The same organisation's later work on domestic public-cloud availability underscores that governments are now measuring not merely abstract digital capacity, but the geographic distribution of commercially usable

AI compute. Commercialisation is therefore not an optional "later stage"; it is part of the very definition of effective sovereign capability (OECD, 2023; Lehdonvirta et al., 2025).

This point follows also from the economics of AI infrastructure. OECD work on competition in AI infrastructure finds high concentration, increasing vertical integration and extensive investments and partnerships across the supply chain. The Federal Trade Commission's 2025 report on partnerships between cloud service providers and AI developers adds that such agreements may give cloud providers access to privileged technical and business information, including model-development methods, chip co-design, customer usage and revenue data. Other reports have also documented multibillion-dollar investments by major cloud providers in AI model developers, often accompanied by contractual commitments to purchase substantial amounts of cloud-computing capacity from the investing provider. The Competition and Markets Authority (CMA) has likewise warned that a small number of incumbent technology firms could shape foundation-model markets in ways that undermine fair and effective competition and produce over-dependence across critical sectors. In short, the commercial pathways through which AI reaches the market are also pathways through which sovereignty can be lost (OECD, 2025; Federal Trade Commission, 2025; Competition and Markets Authority, 2024a; Competition and Markets Authority, 2024b).

Commercial sovereignty should therefore be defined more precisely as the legally and institutionally supported capacity of a polity to retain meaningful control over the AI stack while enabling domestic actors to build, deploy, scale and export AI products on commercially sustainable terms. That definition is deliberately narrower than autarky. None of the leading jurisdictions analysed here aims for complete self-sufficiency. What they seek instead is selective control over choke points: compute access, strategically significant datasets, public-sector demand, key financing pathways, high-value technical standards and the legal terms under which foreign capital and foreign platforms may participate. This is best understood as managed interdependence, not isolation (Roberts, 2024).

Viewed functionally, the commercialisation problem clusters around a recurring set of legal levers. The first concerns infrastructure and compute. Unless researchers, startups and sectoral deployers can obtain compute at predictable and affordable cost, domestic model development and application innovation cannot scale. The second concerns data governance. Commercial sovereignty requires lawful access not only to personal data subject to stringent protections, but also to industrial, public-sector and domain-specific data relevant to training, fine-tuning and deployment. The third concerns intellectual property and training-data law. Domestic firms need clarity on invention, licensing and permissible training pathways if they are to capture value rather than merely provide low-margin implementation services. The fourth concerns procurement and sandboxing. Public demand can create first markets that validate domestic technology and allow compliant firms to move from pilot to revenue. The fifth concerns competition law and market structure. Where chokepoints are controlled by a handful of integrated firms, sovereignty policy that ignores concentration simply transfers dependence from the foreign state to the foreign platform. The sixth concerns standards and assurance. Standards are not mere compliance afterthoughts; they are market infrastructure, reducing transaction costs, enabling interoperability and allowing firms to scale across regulated sectors (OECD, 2025; Regulation (EU) 2024/1689; Regulation (EU) 2023/2854; National Institute of Standards and Technology (NIST), 2023).

Data governance is especially important because sovereignty is often discussed in relation to data retention, but less often in relation to data usability. The EU's Data Act and Data Governance Act point towards a model in which lawful sharing, portability and fair access can support downstream innovation without abandoning rights protection. At the same time, data law can entrench dependence if access regimes are too fragmented, slow or burdensome for smaller firms to use in practice. The legal challenge is therefore to link data protection and fairness to commercially usable data ecosystems rather than to treat them as separate policy domains (Regulation (EU) 2023/2854; Regulation (EU) 2022/868).

Intellectual property and copyright law also deserve more attention in sovereignty debates than they have usually received. A state may fund frontier research and even provide sovereign compute, yet still fail to build sovereign capability if the rules governing training, fine-tuning, inventorship or output exploitation leave domestic firms uncertain about their ability to commercialise what they build. That does not mean AI sovereignty requires maximalist proprietary enclosure. On the contrary, open-source models, open technical standards and shared evaluation tools may strengthen sovereignty where they reduce lock-in and facilitate domestic entry. But legal clarity about who can train, adapt and commercialise AI systems remains central to whether value accrues domestically or migrates to firms with stronger bargaining positions upstream.

Public procurement is similarly underappreciated. Procurement does not simply allocate public spending; it can shape industrial structure. Initial public contracts create reference customers, reduce information asymmetry and force the development of auditable, secure, interoperable products. The EU's recent emphasis on public-sector AI adoption, the UK's procurement-linked hardware strategy and Canada's sovereign cloud procurement all reflect the same insight: a state that wants commercially sovereign AI cannot remain merely a regulator. It must also become a sophisticated early customer (European Commission, 2025c; Government of the United Kingdom, 2026c; Government of Canada, 2026b).

Finally, standards and assurance matter, because trust, portability and market access increasingly depend on them. The AI Act expressly relies on harmonised standards and, where necessary, common specifications as part of its internal-market architecture. Outside Europe, the National Institute of Standards and Technology's (NIST's) AI Risk Management Framework (US) and Singapore's AI Verify initiative show how governance and testing tools can function not just as ethics instruments but as commercial infrastructure. Jurisdictions able to shape assurance language, testing practice and interoperability expectations may exercise a form of sovereign influence without controlling the whole technological stack (Regulation (EU) 2024/1689; NIST, 2023; IMDA, 2026).

Comparative trajectories of commercial sovereign AI

The European Union offers the clearest example of a move from regulatory sovereignty towards commercial sovereignty. The AI Continent Action Plan presents AI as pivotal both for growth and for the preservation of European values, while the Apply AI Strategy expressly links sectoral AI adoption to competitiveness and technological sovereignty. Crucially, the Apply AI Strategy does not stop at compliance. It promotes an "AI first" policy, includes targeted sectoral flagships and, notably, speaks of a "buy European" approach for the public sector with attention to open-source AI solutions. In parallel, the EU's AI Factories and planned AI Gigafactories are designed to combine large-scale compute, reliable supply chains and public-private finance, with the InvestAI facility mobilising up to €20 billion to support up to five Gigafactories. The legal significance is that the Union is increasingly using industrial and procurement instruments to make its internal market a launch platform for European AI firms rather than merely a regulated destination market for foreign ones (European Commission, 2025a; European Commission, 2025b; European Commission, 2026).

The EU is also beginning to frame public demand as a sovereignty instrument. A Commission-backed report on AI in the public sector notes that EU public authorities spend nearly €2 trillion annually on goods and services, about 14 per cent of EU Gross Domestic Product (GDP), and argues that strategic procurement could accelerate AI adoption while giving startups and small and medium-sized enterprises (SMEs) opportunities to scale. The report also acknowledges that current frameworks often favour established suppliers through lengthy procedures, strict eligibility thresholds and large contract sizes. Read together with the Apply AI Strategy and the AI Act, this suggests that the next frontier of European AI law will not be only ex ante risk regulation, but

the redesign of procurement and adoption rules so that compliant European suppliers can actually win markets (European Commission, 2025c; European Commission, 2025b; Regulation (EU) 2024/1689).

Yet the European model also illustrates the limits of regulation-led sovereignty. The Union has strong rule-making capacity, but its own policy materials increasingly acknowledge problems of financing, commercial scale and dependency on foreign compute and platform infrastructure. That makes the EU especially important for the article's core thesis. It shows that trusted regulation may be a necessary condition of sovereign AI markets, but it is not a sufficient one. Without domestic scaling pathways, even a highly sophisticated internal-market regime risks becoming a governance framework for technology commercialised elsewhere.

The United Kingdom has moved even more explicitly towards a commercial sovereign-AI model. In its response to the AI Opportunities Action Plan, the Government committed to expand sovereign compute capacity by at least twentyfold by 2030. In a progress report published one year after the Plan's launch, the Government stated that it had committed £2 billion towards that expansion, designated five AI Growth Zones, and established up to £500 million in funding for the Sovereign AI Unit to invest in and support UK AI companies in critical parts of the value chain. The UK AI Hardware Plan goes further by connecting sovereign funding to procurement, export finance and deployment. It proposes a national cloud procurement model that creates opportunities for UK hardware companies to deploy at scale, uses the Sovereign AI Fund to help AI hardware firms scale, and links those firms to export-finance support and export-control guidance. This is a textbook example of commercial sovereignty: the state is not only de-risking research, but building a legal and financing architecture through which domestic firms can move from prototype to scale to export (Government of the United Kingdom, 2025; Government of the United Kingdom, 2026a; Government of the United Kingdom, 2026c).

The UK case is particularly revealing because it couples internal market-building with external commercial ambition. Sovereignty here is not conceived only as defensive autonomy but also as outward capability. Procurement, sovereign compute, financing and export support are tied together in a strategy that assumes domestic capability will be sustainable only if it can achieve scale through deployment and international reach. That connection between internal resilience and external projection is one of the most important developments in current sovereign-AI policy, and one that legal scholarship has only begun to appreciate.

Canada's recent policy is similarly instructive because it makes the sovereignty–commercialisation link legally explicit. The Canadian Sovereign AI Compute Strategy combines near-term augmentation of public compute infrastructure with an AI Compute Access Fund of up to $300 million to help innovators and businesses purchase compute resources. More strikingly, Canada's Sovereign Compute Infrastructure Program requires applicants to demonstrate sustained sovereign control in the national interest through Canadian ownership or legally enforceable control across the full lifecycle of a facility, alongside safeguards against foreign interference, extraterritorial restriction and supply-chain dependency risks. At the same time, Shared Services Canada has launched procurement for sovereign Canadian cloud capabilities that prioritise Canadian-owned and controlled providers, and the federal government has signalled plans for a made-in-Canada AI platform for government use. Canadian policy thus shows how sovereignty can be operationalised not only through ownership rules, but also through access funds and guaranteed public demand (Government of Canada, 2026b; Government of Canada, 2026a).

Canada is especially important because it foregrounds the legal meaning of control. Many discussions of sovereign AI speak casually of "domestic capacity", but the Canadian materials ask a more exacting question: who ultimately controls the facility, the lifecycle, the governance conditions and the exposure to extraterritorial restriction? That is a more analytically rigorous approach. It recognises that sovereign capability may be undermined not only by foreign ownership in a narrow corporate sense, but also by contractual dependence, export-control exposure, financing structure or service-layer lock-in.

India provides a different, but equally important, commercial pathway. IndiaAI has launched calls for indigenous foundational models, coupled them with access to 18,000-plus GPUs through empanelled AI cloud providers, and made commercial viability, scalability, sustainability and financial viability explicit evaluation criteria. It also offers direct funding, compute credits, equity-based investment and co-financing options from venture capital, angel investors and government grants. In addition, IndiaAI's compute-capacity programme is already allocating GPU resources to researchers, startups, MSMEs, students and public bodies. The legal significance is that India is not treating sovereign AI as a purely public-sector buildout: it is constructing a mixed ecosystem in which public compute, domestic data relevance, startup finance and commercial scale are meant to reinforce each other (Government of India, 2025; Government of India, 2026).

India's model also points to a wider lesson. In jurisdictions attempting strategic catch-up, sovereignty may depend less on immediate frontier-model parity than on creating an ecosystem in which compute access, public digital infrastructure, language diversity and mission-driven deployment support domestic firms' ability to move into commercially relevant niches. That is a more incremental but potentially durable pathway, especially where domestic demand is large and policy is oriented towards practical deployment rather than symbolic frontier competition.

Brief external comparators sharpen the broader argument. The United States exemplifies a scale-first model of commercial sovereignty in which private-sector leadership, research-resource sharing, national-security integration, procurement reform and export promotion are used to preserve domestic command of the AI stack. The White House's AI Action Plan and related security measures illustrate a state that treats market leadership and strategic power projection as mutually reinforcing rather than conceptually separate. The strength of this model lies in speed, capital depth and commercial scale. Its weakness is that sovereignty is mediated through a very small number of exceptionally powerful firms, creating internal concentration risks even as the state seeks external strategic advantage (White House, July 2025; White House, 2026a; White House, 2026b).

China, by contrast, presents the deepest integration of commercialisation and national-security logic. The 2017 New Generation Artificial Intelligence Development Plan tied AI directly to economic transformation and set out an explicitly long-term national trajectory. Subsequent algorithmic, data-security and generative-AI governance measures reflect a system in which commercial AI development is supported, but under conditions of sustained central steering and high regulatory salience. This model can build domestic capability quickly where a large internal market and policy coordination already exist, but it does so at the cost of greater state direction and a narrower tolerance for open-ended cross-border dependence (State Council of the People's Republic of China, 2017).

Singapore illustrates a third, small-state pathway. Its National AI Strategy 2.0 and AI Verify initiative suggest that sovereignty can be pursued through trusted governance, assurance, interoperability and digitally open market positioning rather than full-stack autarky. For smaller jurisdictions, that may be the more realistic route: not control of every layer, but influence over the governance, testing and trusted deployment conditions under which AI markets operate (Smart Nation Singapore, 2023; IMDA, 2025).

Legal and institutional tensions

The shift towards commercial sovereign AI creates at least five legal tensions.

The first is the tension between openness and resilience. AI development depends on globally distributed chips, cloud regions, finance, talent and datasets; yet the more concentrated those inputs become, the more vulnerable domestic systems are to external leverage. The United States' export controls and AI-related security measures illustrate the point starkly: access to key hardware, software and strategic partnerships can be

reshaped by the policy decisions of supplier states. Canada's insistence on protection against "extraterritorial restriction" in sovereign compute governance is a direct legal response to this structural vulnerability. Sovereignty, in other words, is being redefined against a background of very real extraterritorial power (Government of Canada, 2026a; White House, 2026a; White House, 2026b).

The second is the tension between competition and national-capability policy. Concentrated AI markets make domestic scaling difficult, yet industrial policy that prefers local firms can itself reduce contestability or entrench politically favoured incumbents. The EU's use of Important Projects of Common European Interest in microelectronics, the Foreign Subsidies Regulation and related economic-security tools all show a legal order attempting to remain open while defending its market against dependency and distortion. Similarly, the CMA, Federal Trade Commission (FTC), and OECD all warn that vertical integration and strategic partnerships in AI infrastructure require close scrutiny. Commercial sovereign AI therefore cannot mean indiscriminate national championism. A credible legal model must maintain rivalry, interoperability and switching possibilities inside the domestic market (European Commission, 2023; Regulation (EU) 2022/2560; Competition and Markets Authority, 2024a; Federal Trade Commission, 2025; OECD, 2025).

The third is the tension between rule-making legitimacy and economic urgency. The AI Act and related governance frameworks are intended to preserve rights and trust, but commercialisation strategies press governments to accelerate adoption, simplify compliance and create rapid public-sector demand. The EU's Apply AI Strategy explicitly aims to strengthen competitiveness and technological sovereignty; the UK's sovereign AI agenda is designed to move at the pace of venture capital and infrastructure deployment; India's foundational-model call combines ethical compliance with commercial viability. The law is therefore asked to do two things at once: slow harmful deployment where necessary and speed beneficial scaling where strategically desirable. That dual function is possible, but only if regulation becomes more tightly integrated with accreditation, testing, procurement and sector-specific support (European Commission, 2025c; Government of the United Kingdom, 2026a; Government of India, 2025).

The fourth is the tension between domestic value capture and global legitimacy. Mügge warns that AI sovereignty discourse can become jurisdictional, competitive and Eurocentric; Roberts warns that descriptive control is not itself normatively sufficient. A commercially successful sovereign AI ecosystem could still fail on legitimacy grounds if it entrenches opaque public-private power, excludes smaller domestic challengers, externalises harms abroad or uses sovereignty rhetoric to cloak protectionism without genuine public benefit. For legal design, this means that commercial sovereign AI must be evaluated not only by domestic market share or valuation growth, but by plural criteria: democratic accountability, due process, competition, security, diffusion of benefits and compatibility with broader public values (Mügge, 2024; Roberts, 2024).

The fifth is the tension between autonomy, power projection and security. States do not seek AI sovereignty only to reduce exposure; they also seek the capacity to project influence through exported technology stacks, governance norms, standards and supply-chain dependencies. The United States increasingly links domestic capability to the outward promotion of the American AI stack. The European Union seeks to project influence through rule-setting and regulatory benchmarking. China combines domestic ecosystem steering with outward standardisation and platform reach. Singapore's influence lies more in governance and assurance. Yet these forms of projection can conflict with domestic competition, with trusted openness, and with the plural values sovereignty is often said to protect. Once sovereignty becomes a vehicle for exporting dependency rather than merely reducing it, the line between legitimate strategic autonomy and geopolitical market power becomes harder to draw (White House, July 2025; European Commission, 2025a; State Council of the People's Republic of China, 2017; Smart Nation Singapore, 2023).

None of these tensions admits of a simple doctrinal solution. They instead suggest that "AI sovereignty" should be approached less as a slogan and more as a structured balancing exercise among openness, innovation, autonomy, power projection and security. Excess openness can deepen dependence on foreign cloud, chips, model providers and standards setters. Excess closure can fragment data, raise compliance costs, reduce downstream adoption and make domestic champions reliant on permanent subsidy. Excessive security controls can slow knowledge diffusion and market scaling; insufficient controls can offshore strategic value capture and expose critical infrastructure or sensitive data. The jurisprudential challenge is to articulate public-law principles capable of disciplining these trade-offs rather than obscuring them under the language of sovereignty.

A framework for commercially driven Sovereign AI

A workable framework for commercial sovereign AI should rest on five linked legal pillars. First, states need assured access to critical inputs: compute, cloud, energy, semiconductors and strategically necessary talent. The relevant instruments include sovereign compute programmes, infrastructure permitting, export-control diplomacy, research-resource sharing and lifecycle governance rules that reduce exposure to foreign interference or sudden denial of access. The point is not that every layer must be domestic, but that critical bottlenecks must not be governed by dependencies from which domestic actors cannot realistically exit (OECD, 2023; Government of Canada, 2026a; Government of the United Kingdom, 2026a).

Secondly, data governance must do more than protect rights; it must make high-quality lawful training and fine-tuning data available on terms that support domestic deployment. The Data Act's concern with fair access and use of data points in this direction, but the commercial challenge is to couple access rights to usable industrial and public-sector data ecosystems. This includes public-data access design, licensing pathways, sectoral data infrastructure and clear rules for secure sharing and portability. Without these, sovereignty policy may achieve territorial control over data while failing to translate that control into domestic productive capability (Regulation (EU) 2023/2854; Regulation (EU) 2022/868).

Thirdly, states need market-creation mechanisms. Public procurement is the most underused of these. If governments remain reluctant buyers, domestic firms will continue to depend on foreign distribution channels for scale. By contrast, procurement reform, pre-commercial procurement, targeted sandboxes, public compute credits, challenge funds and sectoral deployment missions can create early demand while preserving contestability. The EU public-sector report, the Apply AI Strategy, the UK's procurement-linked hardware plan, and Canada's sovereign cloud procurement all point towards this model (European Commission, 2025c; European Commission, 2025b; Government of the United Kingdom, 2026c; Government of Canada, 2026b).

Fourthly, commercial sovereignty requires capital-retention and anti-capture rules. Access funds, co-financing, export finance, public investment vehicles and investment screening should ensure that domestic breakthroughs do not immediately migrate abroad through dependency-inducing financing structures. At the same time, because concentration is endemic to AI infrastructure, states should combine industrial policy with robust competition enforcement, interoperability requirements where appropriate, open technical standards and careful scrutiny of exclusivity, revenue-sharing and information-access provisions in cloud-model partnerships. The objective is not to exclude foreign firms simply because they are foreign; it is to prevent the domestic ecosystem from becoming a dependent reseller layer on top of foreign-controlled chokepoints (Federal Trade Commission, 2025; OECD, 2025).

Fifthly, commercial sovereign AI must be standards-enabled and trade-compatible. A domestically validated AI product that cannot interoperate, certify compliance or travel across borders is not sovereign at scale; it is merely local. Hence the importance of harmonised standards under the AI Act, assurance frameworks such as

NIST's AI Risk Management Framework, and trust infrastructure such as Singapore's AI Verify. The right objective is not to eliminate external links but to prevent coercive dependence while preserving outward commercial reach (Regulation (EU) 2024/1689; NIST, 2023; IMDA, 2025).

These five pillars can be operationalised in a six-stage sequence. The first step is domain selection. No state can be sovereign across every AI layer. Governments should identify sectors where dependence would generate either systemic vulnerability or large national spillovers: defence, critical infrastructure, public administration, health, language technologies, industrial automation, scientific discovery or safety tooling. The second step is dependency mapping: identifying whether the real bottleneck lies in compute, foundation models, high-quality sectoral data, cloud lock-in, procurement access, skills or financing. The third step is legal-tool matching: targeting the relevant bottleneck rather than expanding regulation in the abstract. If compute is the problem, the tools are infrastructure finance and access regimes; if data is the problem, access rights and trusted sharing matter; if commercial demand is the problem, procurement and sandboxes become central; if concentration is the problem, competition and interoperability cannot be avoided.

The fourth step is anchor-market creation. Sovereign AI capabilities become durable when firms can sell into stable, demanding domestic markets that reward compliance, performance, portability and security. Defence and intelligence may be one anchor, but civilian procurement is at least as important. The fifth step is scaling through standards, assurance and outward-facing openness. Domestic validation should be convertible into broader market access rather than trapped within a narrow sovereign niche. The sixth step is metrics-based review. Sovereignty policy should be judged not by statutory volume or headline funding announcements, but by whether domestic firms can build, sell, scale and endure.

The most useful metrics are therefore commercial as well as constitutional. They include the domestic share of AI procurement spend; the number and growth of domestic firms above agreed revenue thresholds; domestic and allied compute availability; the proportion of strategic public datasets accessible through lawful frameworks; time-to-compliance for SMEs under AI regulation; concentration indicators in cloud and model markets; the proportion of public AI deployments using portable or interoperable standards; dependency ratios on foreign frontier models in critical sectors; and export revenues from domestically developed AI products and assurance services. These measures better capture commercially durable sovereignty than the more common rhetoric of "national leadership" or "strategic autonomy".

On this view, the best measure of sovereign AI success is not whether every layer of the stack is domestic. It is whether domestic firms, public bodies, and researchers can bargain, build, substitute and scale without coercive dependence. That is a more legally realistic and economically sustainable objective than either laissez-faire dependence or rigid techno-national isolation.

Conclusion

AI sovereignty will remain rhetorically attractive and analytically unstable unless public-law scholars and policy makers take commercial structure seriously. The present policy turn in Europe, the United Kingdom, Canada and India shows that states increasingly understand this. They are moving from a first phase centred on ethics, safety and general framework regulation towards a second phase focused on compute access, industrial capacity, procurement, domestic financing, sectoral deployment and strategic control over value capture. That shift is not secondary to sovereignty; it is the condition of its durability.

The core legal proposition of this article is therefore simple. Sovereign AI is sustainable only where law helps domestic actors commercialise strategic capability without surrendering control over the terms of access, deployment and value capture. Regulation without markets produces compliant dependency; subsidies without

scale produce transient capability; and protection without competition produces brittle national champions. A credible model of sovereign AI must instead combine public authority, commercial viability, interoperability and constitutional restraint.

In that narrower but more operational sense, the future of AI sovereignty will be decided not only in AI statutes, but in procurement rules, compute-access regimes, competition cases, cloud contracts, export-control frameworks, standards bodies and the legal architecture of public-private industrial policy. The commercialisation question is therefore not peripheral to sovereign AI. It is where sovereignty either becomes institutionally real or remains a language of aspiration.

References

Chesterman, Simon, "Silicon Sovereigns: Artificial Intelligence, International Law, and the Tech-Industrial Complex" (2026) 120 *American Journal of International Law*, https://www.cambridge.org/core/journals/american-journal-of-international-law/article/silicon-sovereigns-artificial-intelligence-international-law-and-the-techindustrial-complex/91E9751E1A9F586F4E4D7FD2FA36556B (Accessed: 20 July 2026).

Competition and Markets Authority, *AI Foundation Models: Update Paper* (2024a), https://www.gov.uk/government/publications/ai-foundation-models-update-paper (Accessed: 20 July 2026).

Competition and Markets Authority, *AI Foundation Models: Technical Update Report* (2024b), https://www.gov.uk/government/publications/ai-foundation-models-update-paper#full-publication-update-history (Accessed: 20 July 2026).

European Commission, Approved IPCEIs in the Microelectronics Value Chain (2023), https://competition-policy.ec.europa.eu/state-aid/ipcei/approved-ipceis/microelectronics-value-chain_en. (Accessed: 20 July 2026).

European Commission, *AI Continent Action Plan* (2025a), https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52025DC0165 (Accessed: 20 July 2026).

European Commission, *Apply AI Strategy* (2025b), https://digital-strategy.ec.europa.eu/en/policies/apply-ai (Accessed: 20 July 2026).

European Commission, *AI Adoption in the EU's Public Sector: An Opportunity to Better Serve Citizens and Support Startups* (2025c), https://digital-strategy.ec.europa.eu/en/library/ai-adoption-eus-public-sector-opportunity-better-serve-citizens-and-support-startups (Accessed: 20 July 2026).

European Commission, AI Factories policy page, including InvestAI and AI Gigafactories updates (last updated 23 April 2026), https://digital-strategy.ec.europa.eu/en/policies/ai-factories (Accessed: 20 July 2026).

European Union, Regulation (EU) 2022/868, *Data Governance Act* (2022) OJ L 2022/868, https://eur-lex.europa.eu/eli/reg/2022/868/oj/eng (Accessed: 20 July 2026).

European Union, Regulation (EU) 2022/2560, *Foreign Subsidies Regulation* (2022) OJ L 2022/2560, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02022R2560-20221223 (Accessed: 20 July 2026).

European Union, Regulation (EU) 2023/2854, Data Act (2023) OJ L 2023/2854, https://eur-lex.europa.eu/eli/reg/2023/2854/oj/eng (Accessed: 20 July 2026).

European Union, Regulation (EU) 2024/1689, Artificial Intelligence Act (2024) OJ L 2024/1689, https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng (Accessed: 20 July 2026).

Federal Trade Commission, *Partnerships Between Cloud Service Providers and AI Developers* (2025), https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-issues-staff-report-ai-partnerships-investments-study (Accessed: 20 July 2026).

Government of Canada, *Program Guide: Artificial Intelligence Sovereign Compute Infrastructure Program* (2026a), https://ised-isde.canada.ca/site/ised/en/program-guide-artificial-intelligence-sovereign-compute-infrastructure-program-scip (Accessed: 20 July 2026).

Government of Canada, Canadian Sovereign AI Compute Strategy (2026b), https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy (Accessed: 20 July 2026).

Government of India, IndiaAI, *Call for Proposals to Build Foundational AI Models* (2025), https://indiaai.gov.in/article/indiaai-mission-call-for-proposals-to-build-foundational-ai-models (Accessed: 20 July 2026).

Government of India, IndiaAI, *IndiaAI Compute Capacity* (updated July 2026), https://indiaai.gov.in/hub/indiaai-compute-capacity (Accessed: 20 July 2026).

Government of the United Kingdom, *AI Opportunities Action Plan: Government Response* (2025), https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan (Accessed: 20 July 2026).

Government of the United Kingdom, *AI Opportunities Action Plan: One Year On* (2026a), https://www.gov.uk/government/publications/ai-opportunities-action-plan-one-year-on/ai-opportunities-action-plan-one-year-on (Accessed: 20 July 2026).

Government of the United Kingdom, AI Firms Pioneering Drug Discovery, Cheaper Supercomputing and More Get First Backing Through UK's Sovereign AI (2026b), https://www.gov.uk/government/news/ai-firms-pioneering-drug-discovery-cheaper-supercomputing-and-more-get-first-backing-through-uks-sovereign-ai (Accessed: 20 July 2026).

Government of the United Kingdom, *UK AI Hardware Plan* (2026c), https://www.gov.uk/government/publications/uk-ai-hardware-plan (Accessed: 20 July 2026).

IMDA, *AI Verify* (2023), https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/press-releases/2023/singapore-launches-ai-verify-foundation (Accessed: 20 July 2026).

Lehdonvirta, Vili et al., "Measuring Domestic Public Cloud Compute Availability for Artificial Intelligence" (2025) OECD Artificial Intelligence Papers No 49, https://www.oecd.org/en/publications/measuring-domestic-public-cloud-compute-availability-for-artificial-intelligence_8602a322-en.html (Accessed: 20 July 2026).

Mügge, Daniel, "EU AI Sovereignty: For Whom, to What End, and to Whose Benefit?" (2024) 31(8) *Journal of European Public Policy* 2200–2225, https://www.tandfonline.com/doi/full/10.1080/13501763.2024.2318475 (Accessed: 20 July 2026).

National Institute of Standards and Technology, *AI Risk Management Framework* (2023) https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10 (Accessed: 20 July 2026).

OECD, *A Blueprint for Building National Compute Capacity for Artificial Intelligence* (2023), https://www.oecd.org/en/publications/a-blueprint-for-building-national-compute-capacity-for-artificial-intelligence_876367e3-en.html (Accessed: 20 July 2026).

OECD, *Competition in Artificial Intelligence Infrastructure* (2025), https://www.oecd.org/en/publications/competition-in-artificial-intelligence-infrastructure_623d1874-en.html (Accessed: 20 July 2026).

Roberts, Huw, "Digital Sovereignty and Artificial Intelligence: A Normative Approach" (2024) 26 *Ethics and Information Technology* art 70, https://link.springer.com/article/10.1007/s10676-024-09810-5 (Accessed: 20 July 2026).

Smart Nation Singapore, *Singapore National AI Strategy 2.0* (2023), https://www.smartnation.gov.sg/initiatives/national-ai-strategy/ (Accessed: 20 July 2026).

State Council of the People's Republic of China, *New Generation Artificial Intelligence Development Plan* (2017), https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/ (Accessed: 20 July 2026).

White House, *America's AI Action Plan* (July 2025), https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf (Accessed: 20 July 2026).

White House, *Promoting Advanced Artificial Intelligence Innovation and Security* (2026a), https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/ (Accessed: 20 July 2026).

White House, *Fact Sheet: President Donald J. Trump Signs Historic Directive on AI in the National Security Enterprise* (2026b), https://www.whitehouse.gov/fact-sheets/2026/06/fact-sheet-president-donald-j-trump-signs-historic-directive-on-ai-in-the-national-security-enterprise/ (Accessed: 20 July 2026).

Reverse Mentoring as a Catalyst for Authentic Reciprocal Learning in Multigenerational Workforces: An Empirical Research Study in Cyprus

By

Elena Constantinou

Special Teaching Scientist at the University of Cyprus.

Certified Team Coach.

Senior Partner and Lead Corporate Trainer at Evolving Minds.

Abstract

The increasing coexistence of multiple generations in contemporary workplaces has created new opportunities and challenges for talent management, organisational learning and knowledge transfer. Organisations today often employ talent from four or five different generations, resulting in diverse experiences and expectations regarding leadership, communication, learning and technology. Within this context, reverse mentoring, in which younger employees mentor more senior colleagues, has emerged as an innovative developmental strategy. This research note reviews the concept of reverse mentoring, its theoretical foundations, organisational benefits, as well as challenges and best practices for implementation. The paper argues that reverse mentoring extends beyond the need for technological upskilling of an ageing workforce, but functions as a strategic mechanism for reducing generational gaps, strengthening knowledge exchange and reciprocal learning, fostering innovation and promoting cultures of continuous feedback and feed-forward. Practical recommendations for designing effective reverse mentoring programmes are discussed, with particular relevance for relationship-oriented cultural contexts such as Cyprus. We conclude that reverse mentoring may be successfully adapted to organisational contexts where authority remains prominent, provided that organisations cultivate the relational conditions necessary for authentic reciprocal learning.

Keywords: reverse mentoring, reciprocal learning, intergenerational learning, multigenerational workforce, talent management, leadership development.

Introduction

Baby Boomers. Gen X. Millennials. Gen Z. Gen AI. All working in mixed teams on common projects representing different departments. How efficient and effective is communication in such teams? How easily can team members align their expectations? In an era of multiple generations in the workplace, what can organisations do to tackle the various challenges of multigenerational work environments?

Modern organisations are experiencing unprecedented demographic complexity. For the first time in history, many workplaces simultaneously include four or more generations of employees. This multigenerational composition certainly enriches organisations with diverse experiences, perspectives and competencies (Lyons and Kuron, 2014), but it can also create tensions related to communication styles, level of technological adoption, leadership expectations and work values (Stiglbauer, Penz and Batinic, 2022).

The rapid acceleration of digital transformation, artificial intelligence (AI), automation and platform-based collaboration has intensified these challenges. Older employees often possess substantial institutional knowledge, strategic insight and relationship capital, while younger employees frequently contribute digital fluency, technological agility and awareness of emerging social and market trends. Traditional talent development models, where expertise flows primarily from senior to junior employees, may therefore provide inadequate opportunities for organisational growth.

Against this backdrop, reverse mentoring has emerged as an important and strategic organisational practice. Originally popularised by Jack Welch at General Electric in the late 1990s, reverse mentoring challenged traditional assumptions about who holds valuable expertise. Younger employees began mentoring senior executives on internet technologies and digital trends. This initiative accelerated digital transformation whilst challenging assumptions with regards to expertise and hierarchy. GE demonstrated a critical insight: expertise in emerging domains may reside lower in the hierarchy. Since then, the concept has expanded significantly, for example, at Procter & Gamble, which uses reverse mentoring to help senior leaders better understand younger

consumer behaviour, digital marketing and social media ecosystems, at IBM which uses cross-generational learning initiatives to improve digital capability and innovation across teams and at PwC where reverse mentoring improves inclusion and diversity leadership.

Today, reverse mentoring is increasingly recognised not merely as a tool for technological training, but as a strategic intervention supporting intergenerational learning, reciprocal learning, diversity, inclusion, leadership development, innovation and talent retention.

What is Reverse Mentoring?

According to the European Mentoring and Coaching Council (EMCC), mentoring is defined as "a learning relationship involving the sharing of skills, knowledge and expertise between a mentor and mentee through developmental conversations, experience sharing and role modelling. The relationship may cover a wide variety of contexts and is an inclusive two-way partnership for mutual learning that values differences."

Traditional mentoring is typically defined as a developmental relationship, in which an experienced senior professional supports the growth of a less experienced junior colleague (Kram, 1985). It serves career development, psychosocial support and professional identity formation. Reverse mentoring inverts this model.

Murphy (2012) describes reverse mentoring as a relationship in which junior employees mentor senior employees by sharing expertise, perspectives and knowledge which is unavailable through traditional hierarchical structures. Gordon (2022) expands the definition to include a senior leader receiving mentorship from an individual belonging to an underrepresented demographic group, such as those differentiated by gender, age, ethnicity, social mobility or disability. This process reverses traditional hierarchical learning dynamics, positioning the senior leader as the novice and encouraging them to adopt a growth mindset. Through this engagement, leaders are challenged to critically reflect on their biases, broaden their perspectives and take meaningful action towards fostering equity and inclusion within the organisation (Gordon, 2025).

Modern scholarship increasingly emphasises **reciprocity** rather than simple role reversal. Gadomska-Lila (2020) argues that reverse mentoring is most effective when both participants contribute complementary forms of expertise: the mentor offers digital and cultural insights, while the mentee contributes strategic, relational and contextual knowledge. Thus, reverse mentoring should not be understood as *young teaches old*, but rather as *intergenerational co-learning.*

International Trends Supporting Reverse Mentoring

Interest in reverse mentoring has recently increased significantly due to three important macro trends. Firstly, the continuous need for digital transformation is encouraging organisations to constantly strive for better AI literacy, automation capability and digital adaptability. Secondly, many economies have ageing workforces, making older employee upskilling essential. Thirdly, what started with the Millennials is augmented with Generation Z, as younger employees increasingly seek voice, purpose and influence in organisations.

Rastogi, Vapiwala and Pandita (2026) note that AI adoption disproportionately challenges older workers, particularly Baby Boomers and Generation X, who may experience technostress and fear of skill obsolescence. Reverse mentoring offers an important mechanism for bridging this technological divide. Research further suggests that younger generations increasingly expect flatter hierarchies and collaborative learning environments – this reality creates an interesting alignment of reverse mentoring with these evolving workplace cultures.

The Benefits of Reverse Mentoring

Reverse mentoring creates benefits at three levels: the **organisation**, the **mentor** and the **mentee**

One of the most significant organisational benefits is bidirectional knowledge transfer and reciprocal learning. Organisations preserve institutional memory while simultaneously acquiring contemporary knowledge. Whilst senior employees contribute tacit knowledge, strategic judgement and stakeholder understanding, junior employees contribute digital fluency, AI literacy, social media and market trend awareness. Innovation can often emerge when diverse cognitive frames interact. Reverse mentoring introduces constructive friction between established and emerging viewpoints which fuels creativity. Moreover, organisations can benefit from lower employee turnover, as younger employees are more likely to remain engaged when they perceive that their expertise is recognised. Talent is better attracted and retained in an environment where early leadership experiences exist for high-potential younger talent.

Finally, reverse mentoring humanises the differences which can arise from the tension produced by stereotypes between generations, thus promoting empathy. Gadomska-Lila (2020) identifies organisational benefits including improved communication, enhanced cooperation, increased trust, diversity support and a stronger lifelong learning culture.

The **younger mentors** often experience substantial professional development, such as increased work satisfaction and confidence, stronger leadership capability, visibility to senior leadership, expanded professional networks and improved communication skills. Mentors also learn to influence without formal authority, an important leadership competency.

The **mentees**, or senior leaders, benefit through digital upskilling, greater learning agility and teaching capability, improved understanding of younger employees, reduced generational bias and enhanced adaptability. Perhaps most importantly, mentees cultivate humility and openness, qualities which are increasingly essential in complex business environments.

The Challenges of Reverse Mentoring

Despite the advantages, there are several challenges standing in the way of implementing reverse mentoring programs, from the viewpoints both of organisational culture and of employee resistance. The greatest barriers are often psychological rather than procedural and can include hierarchical resistance, i.e., when senior leaders perceive reverse mentoring as threatening because it challenges traditional status hierarchies. Hierarchical resistance is especially relevant in cultures characterised by higher power distance, for example, societies where people accept a high inequality in power and status as normal or legitimate.

Successful reverse mentoring requires psychological safety. If there is not an adequate level of trust between colleagues, junior mentors may censor themselves while senior mentees may become defensive. Furthermore, common generational stereotypes, such as, “young people lack commitment” or “older people resist change”, may undermine openness and inhibit authentic learning. Since the success of reverse mentoring programmes depends heavily on pairing, mismatch in a pair’s personality, communication style or expectations can reduce effectiveness. That is why many studies find that the effectiveness of a programme greatly depends on training and preparing participants adequately, building trust and creating participant engagement, as well as securing sustained organisational and leadership commitment (Gadomska-Lila, 2020).

Empirical Research: Cypriot Managers' Perceptions

While international literature on reverse mentoring is growing, empirical research in Cyprus remains limited. This presents a valuable opportunity for exploratory research examining how Cypriot managers perceive reverse mentoring. Consequently, we embarked on the current exploratory study by distributing an online questionnaire to over 80 managers and laying the foundations for future research using semi-structured interviews and implementing experiments with mentor-mentee pairs in various companies.

The purpose of the current study was to investigate how respondents perceive the purpose and value of reverse mentoring programs in multigenerational workplaces, examine their preferences for intended focus areas, identify organisational enablers and barriers, examine cultural attitudes toward hierarchy, establish important programme design principles and gauge the sample's readiness for future implementation. Given the Southern European / Mediterranean's relatively relationship-oriented and hierarchical business cultures (especially in countries with a large percentage of the workforce employed by family businesses, SMEs, traditional service companies and the public sector, such as Cyprus), we were curious to see if this research would reveal culturally specific facilitators and obstacles for reverse mentoring programmes.

The questionnaire informing this study was completed in June 2026 by 61 managers from various industries (including IT, manufacturing, banking and finance, education, hospitality, legal, retail, telecoms, coaching and government), with an average age of 48 years, 48% male and 52% female and an average management experience of 15 years (managing teams of 2-700 employees with over 90% of these teams being multigenerational). Reponses were anonymous and informed consent was obtained. A brief definition and explanation of reverse mentoring was offered at the beginning of the questionnaire. Questions included both open-ended prompts and questions requiring participants to select all which applied or to rank the importance of certain characteristics based on a five-point Likert scale. 49% of participants responded that they had participated in a traditional mentoring program at work before, whereas only three respondents had ever participated in a reverse mentoring program at work.

Open questions included:

- What should be the primary purpose of reverse mentoring?
- Which organisational structures best support it?
- What barriers are most significant in your organisations?
- Which principles should underpin programme design?

Overall, the findings indicate that Cypriot managers hold predominantly positive attitudes toward reverse mentoring. Figure 1 suggests that Cypriot managers are not resistant to the idea of reverse mentoring in principle, as only approximately 15% of them believe that reverse mentoring would be slightly or not important to their organisation. Consequently, 85% of managers believe reverse mentoring would be moderately important, important or extremely important. This finding is noteworthy, because reverse mentoring challenges conventional assumptions about where expertise resides within organisations.

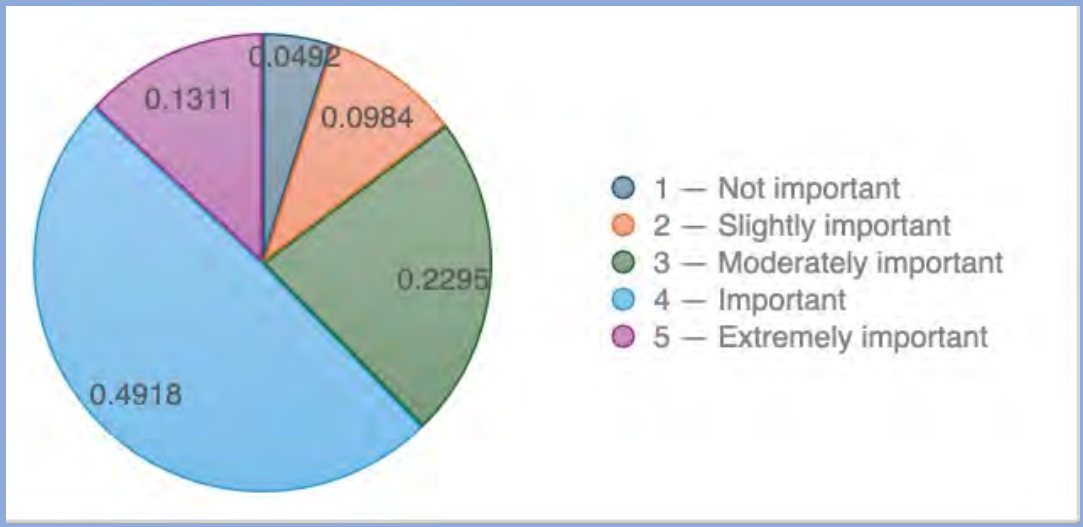


Figure 1: How Important do you believe reverse mentoring could be for your organisation?

Cypriot managers appear receptive to the concept of reverse mentoring, particularly when it is framed as a developmental and collaborative process rather than as a challenge to managerial authority. This distinction emerged repeatedly across both quantitative and qualitative responses and represents one of the most significant findings of this study. The data suggests that acceptance of reverse mentoring in Cyprus is best understood not as unconditional endorsement, but as openness to authentic reciprocal learning. Managers generally support reverse mentoring when it enhances communication, facilitates learning and improves collaboration, yet exhibit greater caution when it appears to imply deeper structural redistribution of influence or authority.

One of the clearest findings concerns the perceived primary objective of reverse mentoring. Respondents overwhelmingly prioritised **improving communication between generations**, ranking this more consistently than organisational culture transformation or technological upskilling, as can be seen in Table 1.

Ranking	Objectives of reverse mentoring program
1	Improving communication between generations
2	Knowledge sharing across generations
3	Digital skills and technology adoption of senior employees
4	Generating innovation and creative thinking
5	Retention of younger talent
6	Understanding younger generations (e.g., Generation Z values and expectations)
7	Organisational culture transformation
8	Diversity, equity and inclusion
9	Leadership development

Table 1: Resulting ranking of reverse mentoring program objectives

This finding is particularly interesting when contrasted with the international reverse mentoring literature. The original corporate examples of reverse mentoring focused primarily on technological knowledge transfer, with younger employees mentoring senior executives on internet usage and emerging digital technologies. Early scholarship similarly emphasised reverse mentoring as a mechanism for digital literacy and technological adaptation (Murphy, 2012).

The Cyprus data suggests that reverse mentoring is increasingly perceived through a broader and more relational lens. Rather than viewing it primarily as a vehicle for technological upskilling, managers appear to see reverse mentoring as a mechanism for strengthening intergenerational understanding and reducing communication barriers and experience gaps between age groups. This finding aligns closely with recent scholarship, which conceptualises reverse mentoring as a reciprocal learning relationship that supports knowledge transfer and relational cohesion (Gadomska-Lila, 2020).

This shift from technological emphasis toward relational emphasis may suggest that the most prominent generational divide perceived by respondents, may not be a digital skills gap alone, but also a gap in expectations, values and communication styles. Such an interpretation is highly consistent with contemporary talent management literature. Younger employees, particularly Millennials and Generation Z, frequently demonstrate different expectations regarding autonomy, flexibility, purpose, inclusion and feedback. Senior managers who developed professionally in more traditional work environments may struggle to interpret these expectations accurately. Reverse mentoring may therefore serve as an important interpretive bridge, helping senior leaders better understand the motivations, concerns and behavioural norms of younger employees (Gordon, 2025; Murphy, 2012).

The survey findings also indicate that **retention of younger talent** is viewed as an important strategic outcome of reverse mentoring. This reflects growing managerial awareness of the challenges associated with attracting and retaining younger employees in increasingly competitive labour markets. Reverse mentoring may be particularly valuable in this regard, because it signals recognition, voice and influence for younger employees. When junior staff perceive that their knowledge and perspectives are valued at senior organisational levels, they may experience greater engagement and stronger psychological attachment to the organisation, reducing employee turnover. This finding supports contemporary talent management research, which emphasises that younger employees increasingly value participation, visibility and meaningful contribution (Vieira, Costa and Santos, 2024).

By contrast, respondents showed greater variation in their views regarding reverse mentoring as a tool for **organisational culture transformation**. While many managers considered culture transformation important, responses were notably less uniform than for communication-related objectives. This divergence provides an important theoretical insight. Respondents appear comfortable supporting reverse mentoring as a means of improving dialogue and collaboration, yet somewhat less unanimous when reverse mentoring is framed as a mechanism for broader cultural or structural organisational change.

Survey participants recognise that barriers to implementing reverse mentoring programs primarily include senior manager resistance to role reversal, lack of trust, fear of judgement or loss of authority and a hierarchical organisational culture (Table 2). These concerns suggest that authority and role expectations remain meaningful in managerial thinking.

Table 2 can be found on the next page.

Barriers to reverse mentoring	Frequency
Senior managers resisting role reversal	40
Lack of trust between employees	32
Fear of judgment or loss of authority	30
Hierarchical organisational culture	28
Lack of time	25
Poor communication skills of employees	21
Generational stereotypes	20
Unclear benefits of program	19
Lack of HR support	11

Table 2: Which barriers are most likely to prevent reverse mentoring in your organisation?

This interpretation leads to what may be the central conceptual contribution of this study. The findings suggest that reverse mentoring in Cyprus does not necessarily function as a mechanism for dismantling hierarchy. Rather, it appears to support a model of hierarchical leadership with distributed expertise. Under this model, formal authority may remain largely hierarchical, but expertise is no longer assumed to correlate exclusively with seniority or organisational rank. Knowledge may therefore flow upward even while decision-making authority continues to flow downward. This distinction is theoretically important because it challenges assumptions that reverse mentoring requires flat organisational structures in order to succeed.

Responses to the question "Which principles should form the foundation of a reverse mentoring program" and answers to open questions further reinforce this interpretation by highlighting the importance of mutual respect, reciprocity, clear expectations, trust and psychological safety (Table 3). These relational conditions emerged repeatedly as essential prerequisites for successful reverse mentoring programmes. Respondents emphasised mutual respect, openness to learning and safe communication as core principles that should underpin any reverse mentoring initiative. This finding aligns strongly with both Murphy (2012) and Gadomska-Lila (2020), who argue that reverse mentoring depends heavily on relational quality and organisational support.

Table 3 can be found on the next page.

Principles	Frequency
Mutual respect between mentor and mentee	48
Reciprocity / two-way learning	42
Clear expectations of both parties	37
Trust	31
Psychological safety	28
Confidentiality during the process	24
Voluntary participation	23
Structured evaluation	19
Regular meetings	16
Equality between mentor and mentee	12

Table 3: Which principles should form the foundation of a reverse mentoring programme?

Clear expectations, trust and psychological safety appear particularly important because reverse mentoring creates vulnerability for both participants. Junior mentors may fear appearing disrespectful, arrogant or overly critical when providing feedback to senior leaders. Senior mentees, meanwhile, may experience discomfort associated with acknowledging knowledge gaps or limitations in front of younger employees. These dynamics suggest that reverse mentoring requires more than simple programme design; it requires a relational environment in which vulnerability is acceptable and learning is decoupled from status.

A particularly noteworthy theme which arose from the open-ended questions, was the importance of **humility** among senior leaders. Although not explicitly measured as a survey variable, many responses indirectly suggested that successful reverse mentoring requires leaders to demonstrate curiosity, openness and willingness to learn from younger colleagues. This observation introduces an important leadership dimension often underemphasised in reverse mentoring literature: **intellectual humility**. Reverse mentoring challenges senior leaders to acknowledge that strategically valuable knowledge may reside outside traditional centres of authority. In this sense, reverse mentoring does not merely redistribute knowledge; it redistributes epistemic authority, or authority over what constitutes valuable organisational knowledge.

Respondents also placed strong emphasis on programme design, including support for open communication, careful mentor–mentee matching, confidentiality, clear objectives, leadership sponsorship and structured facilitation (Table 4). This aligns with prior research demonstrating that reverse mentoring programmes rarely succeed when implemented informally or symbolically. Gadomska-Lila (2020) found that mentoring effectiveness depends heavily on organisational support structures, participant preparation and clear developmental goals. The present findings strongly support this conclusion.

Table 4 can be found on the next page.

Organisational Characteristics	Frequency
Open communication culture	54
Clear goals of the reverse mentoring program	41
Leadership support / sponsorship	37
Psychological safety	29
Diversity of age groups	28
Formal training for participants	18
Flat organisational structure	8

Table 4: Which organisational characteristics most support successful reverse mentoring in your opinion?

Taken together, the findings suggest that reverse mentoring in Cyprus may function less as a disruptive anti-hierarchical intervention and more as a mechanism for enabling reciprocal learning within existing organisational structures. This represents an interesting refinement of traditional reverse mentoring theory. Much of the international literature implicitly assumes that flatter organisational cultures are more conducive to reverse mentoring. The Cyprus findings suggest instead that organisations characterised by meaningful role boundaries may also successfully adopt reverse mentoring, provided that learning relationships are grounded in trust, mutual respect and psychological safety.

In conclusion, the empirical findings reveal that Cypriot managers perceive reverse mentoring as a promising talent management strategy, particularly for improving intergenerational communication and supporting retention of younger talent. However, managerial support appears strongest when reverse mentoring enhances collaboration without overtly destabilising traditional authority structures. Reverse mentoring in the Cypriot context therefore emerges not as a rejection of hierarchy, but as a mechanism for decoupling expertise from formal rank. This insight contributes to the broader literature by suggesting that reverse mentoring may be successfully adapted to organisational contexts where authority remains prominent, if organisations cultivate the relational conditions necessary for authentic reciprocal learning.

Best Practices for Implementing Reverse Mentoring Programmes

Drawing on both the literature and the empirical findings of this study, Figure 2 summarises the organisational, relational and programme-design conditions most likely to support effective reverse mentoring:



Figure 2: Best practices for implementing reverse mentoring programmes

Successful reverse mentoring can be approached not just as an isolated HR initiative, but as a deliberately designed organisational learning intervention. The effectiveness of such programmes depends not only on identifying suitable mentors and mentees, but also on the quality of the relational, structural and cultural conditions surrounding the mentoring relationship.

A first requirement is **visible senior leadership sponsorship**. Leadership sponsorship performs both a practical and symbolic function: it provides resources and legitimacy while also modelling the intellectual humility required for reverse mentoring to succeed.

A second critical factor **is clarity of purpose**. Organisations should establish why reverse mentoring is being introduced and what outcomes it is expected to achieve. That being said, programmes require what may be described as **structured flexibility:** sufficient structure to establish expectations, accountability and psychological safety, while leaving enough conversational freedom for unexpected learning to emerge.

Thirdly, **mentor-mentee matching** should be developmental rather than merely demographic. Successful matching usually considers personalities, learning needs and compatibility.

A fourth best practice relates to the **training** which mentors and mentees can receive before, during and after the mentoring programme. Such training can include active listening skills and techniques for giving and receiving constructive feedback.

A very important requirement for a reverse mentoring programme is that it is viewed and treated as an **iterative learning system rather than a one-off intervention**. Programmes should include agreed SMART objectives (especially specific, agreed and measurable), regular but not overly-engineered meetings, opportunities to review the relationship and a mechanism for rematching or exiting when a pairing is proving ineffective.

Last but not least, organisations should cultivate a communication climate based on **confidentiality, trust and psychological safety** before expecting candid upward learning to occur.

Ultimately, successful reverse mentoring requires organisations to create the conditions for **authentic, reciprocal learning**. When structure is combined with trust, openness and mutual respect, expertise can flow across generations and organisational levels, regardless of age or seniority. In this way, reverse mentoring can evolve from a discrete HR programme into an embedded mechanism for continuous organisational learning.

Conclusions

Reverse mentoring represents a significant evolution in organisational learning and talent management. In multigenerational workplaces, expertise is increasingly distributed across age groups rather than concentrated solely among senior employees. Reverse mentoring acknowledges this shift by creating structured opportunities for reciprocal learning.

Its strategic value extends beyond technological upskilling. Effective reverse mentoring improves intergenerational understanding, accelerates innovation, strengthens talent retention and fosters more inclusive organisational cultures. However, programme success depends heavily on organisational readiness, leadership support, trust and psychological safety.

For organisations operating in increasingly dynamic environments, reverse mentoring offers more than a development tool - it provides a framework for cultural transformation from hierarchy-based learning toward continuous, reciprocal learning.

In contexts such as Cyprus, where traditional leadership norms may remain strong, reverse mentoring may prove particularly valuable as both a talent management strategy and a catalyst for modernising organisational culture.

Future Work

This research note has laid the foundation for further research in Cyprus and, potentially, other Southern European and Mediterranean contexts. As a first step, we aim to verify and deepen the current findings through semi-structured interviews with managers and younger employees, particularly exploring the roles of hierarchy, trust and intellectual humility in reverse mentoring relationships. Exploring younger employees' willingness to

mentor senior colleagues, their perceptions of psychological safety and their experiences of influencing individuals with greater formal authority would provide a more complete understanding of the dynamics involved in reverse mentoring programmes.

Subsequent research could extend the study through collaborations with researchers in other countries, enabling cross-cultural comparison of the organisational and relational conditions that facilitate reverse mentoring. Of particular interest would be whether the model of hierarchical leadership with distributed expertise emerging from the present findings extends beyond the Cypriot context. In the longer term, we aspire to implement and evaluate reverse mentoring interventions involving mentor-mentee pairs across organisations in Cyprus and to use these findings to develop practical, evidence-based principles for the design and implementation of effective reverse mentoring programmes.

References

EMCC Global. Available at: https://emccglobal.org/home/leadership-development/leadership-development-mentoring/ (Accessed: 22 June 2026).

Gadomska-Lila, K. (2020) 'Effectiveness of reverse mentoring in creating intergenerational relationships', *Journal of Organizational Change Management*, 33(7), pp. 1313–1328.

Goldsmith, M. (2002) *Try feedforward instead of feedback*. New York: Leader to Leader Institute.

Gordon, P. (2025) 'Reverse Mentoring: Transforming Learning at Individual, Team, and Organizational Levels to Build Belonging in the Workplace', *International Journal of Human Resource Development: Practice, Policy and Research*, 9(1), pp. 53–60. https://doi.org/10.2478/ijhrd-2025-0006

Kram, K.E. (1985) *Mentoring at work: Developmental relationships in organizational life*. Glenview: Scott Foresman.

Lyons, S.T. and Kuron, L.K.J. (2014) 'Generational differences in the workplace: A review of the evidence and directions for future research', *Journal of Organizational Behavior*, 35(S1), pp. S139–S157.

Murphy, W.M. (2012) 'Reverse mentoring at work: Fostering cross-generational learning and developing Millennial leaders', *Human Resource Management*, 51(4), pp. 549–573. https://doi.org/10.1002/hrm.21489

Rastogi, S., Vapiwala, F. and Pandita, D. (2026) 'Upskilling ageing employees: reverse mentoring as a strategy in the era of artificial intelligence', *Working with Older People*, 30(3), pp. 277–285. https://doi.org/10.1108/WWOP-01-2025-0001

Stiglbauer, B., Penz, M. and Batinic, B. (2022) 'Work values across generations: Development of the New Work Values Scale (NWVS) and examination of generational differences', *Frontiers in Psychology*, 13, 1028072. https://doi.org/10.3389/fpsyg.2022.1028072

Stone, D., Cox, A. and Gavin, M. (2020) 'Technology adaptation and organisational capability', *Journal of Business Strategy*, 41(2), pp. 22–34.

Vieira, J., Gomes da Costa, C. and Santos, V. (2024) 'Talent management and Generation Z: A systematic literature review through the lens of employer branding', *Administrative Sciences*, 14(3), 49. https://doi.org/10.3390/admsci14030049

*** Acknowledgement on the use of AI—***Specifically, a statement to acknowledge the use of Generative AI in the writing process:* The author used ChatGPT (OpenAI, GPT-5.5) to assist with stylistic improvements. All content was reviewed and edited as needed and the author takes full responsibility for the content. This action was completed on 5th August, 2026.

Foreign Market Entry:
Traditional and Contemporary Developments.

By

Prof. Marios Georgiou

MBA Director at CIM-Cyprus Business School, Nicosia – Limassol

Abstract

Foreign market entry strategies are critical for organisations looking to expand their operations to other countries abroad, as these require significant commitments, financial and other, which are, more often than not, irreversible in the short and medium term. The area of foreign market entry has received considerable and growing academic attention over the last twenty years or so. International market entry is now one of the most researched areas in international management, together with the related fields of foreign market investments and internationalisation.This paper provides a review of traditional foreign market entry research and identifies the current state of the field, highlighting the main theories and frameworks. A systematic literature review showed that there is a plethora of empirical evidence on foreign market research which provides useful insights into this strategic process. However, the findings are often conflicting, which makes it difficult to draw general conclusions. The literature review has categorised the main theories and constructs of foreign market entry.

Keywords: Foreign Market Entry, Modes of Market Entry, International Expansion.

1. Introduction

The term foreign market entry refers to the process by which organisations enter international markets for the purpose of undertaking value-creating activities (Jones & Young, 2009). Sharma and Erramilli (2004, p. 2) defined entry mode as:

"a structural agreement that allows a firm to implement its product market strategy in a host country, either by carrying out only the marketing operations (i.e. via export modes), or both the production and marketing operations there by itself or in a partnership with others (contractual modes, joint ventures, wholly owned operations)."

Root (1987, p. 5) provided a different definition of foreign market entry mode as:

> *"an institutional arrangement that makes possible the entry of a company's products, technology, human skills, management or other resources into a foreign country."*

2. Modes of Entry – Classical approaches

An organisation has several options available through which it can enter foreign markets and these include in accordance with Datta, Herrmann and Rasheed (2002):

- Foreign direct investments (greenfield or acquisitions)
- Collaborative strategies (joint ventures)
- Contractual entries (licensing, franchising, agencies, distributorship)
- Exporting.

Pan and Tse (2000) argued that entry modes can be divided into equity (foreign direct investments and joint ventures) and non- equity options (exporting, licensing, agencies). These two broad categories are very different with regards to the level of investment and consequently, to the required level of control. The equity modes require higher involvement and control as the level of investment is higher. In contrast, non-equity modes require minimum owner involvement as the investment is limited.

In general, exporting is considered as the foreign market entry mode with the lowest capital investment and one which requires lower management control. According to Young et al. (1989), exporting can be affected either directly or indirectly. In the indirect mode, the firm exports its products through an exporting agency, or a trading company often based in the home country. More often, however, indirect exporting is carried out by a foreign company buying the products at the home country of the producer (Tihanyi, Griffith & Russell, 2005). All the operations associated with the export process are likely to be carried out by outsiders, without the involvement in, or knowledge of the company.

In contrast, in the direct exporting mode, the firm undertakes all the exporting activities, and the level of its involvement is relatively high. The products or services are exported to the host country through agents or specific outlets, and the home company may have limited or no involvement in the sales and marketing activities in the host market (Young et al., 1989).

Licensing is another form of foreign market entry in which the home company provides the host firm with a local license or access to a specific technology, knowledge, product or service for a specified period and under certain conditions and obligations from both parties (Young et al., 1989). The involvement of the home company (licensor) is higher and usually it has some knowledge of the sales and marketing activities in the host country, as the local firm (licensee) is obliged to provide such information.

Contractual agreements are very similar to licensing agreements, but extend to cover a wider spectrum of activities in the host country. For example, this mode of market entry includes franchising, distributorship agreements and/or sales and marketing alliances. The obligations of each of the two parties are written into the agreement, and often the home company has increased involvement in the sales and marketing activities within the host country, providing guidance and training to the local firm and sometimes sharing some of the local advertising and promotional expenses. In this mode, the two parties share skills and resources and the term of agreement is usually of medium to longer duration (5–10 years) and renewed if both parties are satisfied.

Equity type modes of foreign market entry include joint ventures and wholly owned subsidiaries. In the mode of equity joint venture, the home and host partners share risk, profits and capital investments and the two parties participate actively in the management of the joint venture (JV). The shareholding is usually governed by the relative contribution of each party in the technology, know-how, market access, capital investment, management and other resources (Agarwal, 1994). The agreement is often based on a longer term (for example 10 years) which may be renewed. The home and host firms are involved in all the operations and exert joint control. Wholly owned subsidiaries are another form of equity market entry in which the home organisation acquires an existing foreign firm or sets up a new company in the foreign market (greenfield). In such a mode, the home firm undertakes the entire risk and capital investment but fully retains management and control (Young et al., 1989).

Sarkar and Cavusgil (1996) have provided an integrative classificatory framework, identified the main and sub-themes of the literature and investigated the links between them. They also showed the evolutionary changes in entry mode research and suggested future directions for entry mode research. The authors suggested that the main antecedents in foreign market entry research may be broadly categorised into two main themes; those derived from the economic and behavioural theories which were named as efficiency driven factors (for example, product market factors, firm factors, host market factors and cultural factors); and strategy derived factors (for example, global industry structure and corporate strategy). According to this review, some of the research focused on the outputs of the foreign market choices and particularly to the firm's performance. Finally, the authors suggested the need to research the complex relationship between the economic and non-economic variables in the foreign market entry strategy.

Valuable insights into the foreign market research were provided by the review of Datta, Hermrman, and Rasheed (2002), which classified, reviewed and synthesised past empirical research. The authors have highlighted the main theoretical perspectives that have been proposed on foreign market entry, such as monopolistic advantages, transaction cost, internalisation, internationalisation and eclectic theories. The review focused on empirical evidence from US firms entering foreign markets or entry into the USA from foreign firms and categorised the antecedent factors and outcomes of foreign market entry into six distinct streams (for example, entry firm characteristics, home and host country characteristics, industry characteristics and global conditions). It also investigated how these influence the choice of foreign market entry mode. The authors concluded that the reviewed empirical evidence showed a lack of consensus and inconsistencies which necessitate more work in the area of foreign market entry and suggested directions for further research.

The literature review findings of Georgiou (2014) and Brouthers and Hennart (2007) on international entry mode choice also confirmed the need for new research directions into this area. The study identified weaknesses and limitations and proposed new directions for research. The authors emphasised the lack of consensus in some of the fundamental modes of entry (for example, contracts, joint ventures and wholly owned subsidiaries) and criticised some of the ways that the theories have been employed to explain specific entry mode choices. Moreover, the authors proposed the use of more realistic models to supplement the currently used rational models, thereby taking into consideration the attitudes and beliefs of managers making the decisions as well as the external factors influencing these.

This section has provided the basic modes of foreign market entry found in the literature but these are by no means exhaustive. For example, Erramilli and Rao (1990) identified nine variations of entry modes whereas Brouthers and Hennart (2007) reported sixteen different entry modes. The (definitional) boundaries of these entry modes are not always clear and different researchers or organisations may use their own terminology as well as employ characteristics from different modes for their own market entry mode definition (for example, the terms 'agents' and 'distributors' are used differently by various researchers and organisations).

3. Themes and models

Most of the research studies in this field offer various theoretical propositions regarding the factors influencing the choice of foreign market entry, and some provide empirical evidence to support their theories. The main themes and models which have been researched and it is suggested influence the choice of foreign market entry are outlined below:

- Transaction cost – This is the most frequently researched determinant in foreign market entry (Brouthers & Hennart, 2007; Canabal & White, 2008) and proponents of this theory claim that organisations will select the entry mode that will minimise the total transaction and operational costs and therefore maximise the overall efficiency of the process. In essence, foreign market entry is considered as a specific transaction and the benefits and costs of each option are estimated and the mode selected is the one which provides the lowest overall cost (e.g., Brouthers, 2013a; Madhok, 1997). One factor that has been suggested influences transaction cost economics 'is asset specificity', which is the degree of specialisation and uniqueness of the firm's assets that are employed for the given market entry (Maekelburger, Schwens & Kabst, 2012). The concept of asset specificity, however, has not been measured consistently by researchers (physical assets, technology, research and development (R&D) expenditure and human assets are some of the measures employed). As a consequence, the reported evidence is often contradictory (Brouthers & Hennart, 2007). The prevailing view is that high asset specificity favours higher control of entry modes by the home country, such as wholly owned subsidiaries (K.D. Brouthers & L.E. Brouthers, 2003; K.D. Brouthers, L.E. Brouthers & Werner, 2003).

- Cultural distance – This is based on Hofstede's cultural dimensions theory which proposes five dimensions for analysing cross-cultural communication and behaviour (Hofstede, 1984). Cultural distance refers to the difference in culture between the home and host countries and its influence on the choice of foreign market entry mode (e.g., Samiee, 2013; Tihanyi, Griffith & Russell, 2005). Research evidence does suggest that although cultural distance is an important determinant in foreign market entry decisions, that influence is more likely to be indirect and manifest itself as investment risk (Brouthers & Hennart, 2007). Furthermore, the cultures of the home and host countries separately have been suggested as influences on the entry market mode through the control of the new venture. For example, cultures characterised by uncertainty avoidance are more likely to pursue options with higher host control and ownership (Datta, Herrmann & Rasheed, 2002; Stevens & Dykes, 2013).

- Internationalisation – The firm's experience in international markets affects the foreign market decision (e.g., Dow & Larimo, 2009; Martinez & Lopez, 2009). It has been argued that cumulative international experience and familiarity with diverse international environments enables firms to undertake higher levels of resource commitment and risk and choose a foreign market mode with higher ownership and control (Castaner & Genc, 2011; Datta, Herrmann & Rasheed, 2002; Sarkar & Cavusgil, 1996). Related to this theme are studies of the personal experiences of the executive directors and management teams and how these influence decisions of foreign entry mode (Dikova, 2012; Herrmann & Datta, 2006).

- Risk and uncertainty – Political and other risks and uncertainties in the host countries influence foreign market decisions (e.g., K. D. Brouthers, L.E. Brouthers & Werner, 2000; Lopez-Duarte & Vidal-Suarez, 2010). A number of different ways have been adopted to measure risk (for example, perceived political and economic stability, perceived market potential and market volatility) and these inconsistencies have produced conflicting conclusions (Brouthers, 2002; Diego, Enrique & Laura, 2007; K.D. Brouthers & L.E. Brouthers, 2003). Internal uncertainty is another type of uncertainty related to the firm's experience in international business (Zhao, Luo & Suh, 2004) and similar to external uncertainty. It has been

measured by a variety of methods (for example, the number of years in international businesses, the number of total investments or the ratio to foreign investments to the total investments). Contradictory evidence is reported, but according to Zhao, Luo and Suh (2004), organisations with high experience and thus less uncertainty will select entry modes with higher resource commitment and higher ownership.

- Resource-based view – The proponents of this theory argue that an organisation will choose a foreign entry mode that will best exploit the unique resources and capabilities of the firm (e.g., Ekeledo & Sivakumar, 2001; Madhok, 1998; Sharma & Erramilli, 2004). The organisational resources and capabilities can only provide a competitive advantage if they are valuable, unique and difficult to replicate. Accordingly, such organisations may select an entry mode that will retain these unique resources and competencies within the firm, and will avoid entry modes that necessitate the sharing of these resources (Brouthers & Hennart, 2007; Datta, Herrmann & Rasheed, 2002). Amongst the organisational resources that have been studied are international experience, proprietary technology, skilled work force, know-how and R&D intensity.

- Internalisation – Firms will select a foreign entry mode in an effort to internalise all activities that can be conducted at a lower cost (e.g., Buckley & Casson, 1998; Madhok, 1997). This theory is based on the efficiency rational and can be considered as a variant of the transaction cost approach. If the external firm can perform relevant activities which are less costly than the home organisation, then the entry modes chosen are likely to be joint ventures, or some form of contractual agreement (Datta, Herrmann & Rasheed, 2002).

- Institutional approach – The institutional structure in the host country affects the decision on the mode of market entry (e.g., Brouthers, 2002; Brouthers, 2013a). One aspect of the institutional theory is that it suggests that a new firm entering a host country will attempt to copy some of the actions of local companies or competitors in an effort to establish itself in the market, otherwise it may be forced out. This effect is known by the term 'isomorphic pressure tendency' (Brouthers & Hennart, 2007; Canabal & White, 2008). Other dimensions of the host country's institutional environment that have been studied are the regulatory, cultural, marketing characteristics and level of corruption. The main findings are that these factors do influence entry mode choice (Brouthers, 2002; Meyer, 2001).

- Ownership, location and internalisation – This is known as the eclectic framework, on the premise that these three factors influence the foreign market entry decisions in varying degrees (e.g., K.D. Brouther, L.E. Brouther & Werner, 1999). It may be considered as an approach that brings together three of the main theories that were discussed earlier: the resource-based view (ownership), institutional (location) and internalisation (Brouthers & Hennart, 2007).

- Industry structure and bargaining power – The way a specific industry is structured influences the mode of entry (e.g., Markusen & Stahler, 2011). Relative bargaining power suggests that the sizes of the home and host markets may influence the outcome of negotiations and entry mode choice. For example, a multinational organisation with a large home market has greater bargaining power than a smaller host market firm and may dictate the outcome (Datta, Herrmann & Rasheed, 2002).

4. Contemporary approaches

4.1 Digital platforms

Digital platforms and born-digital firms have changed the entry landscape: firms can now reach large foreign markets with comparatively low upfront physical investments and little physical presence in the host country. The platform governance (rules, algorithms, data flows) becomes a core part of market access. Recent work proposes frameworks for platform firms' market selection and entry tactics, showing platform characteristics, data assets and regulatory constraints shape whether firms choose direct digital servicing, local partners or hybrid modes. Studies focusing on B2B digital platforms emphasise analytics-driven market targeting and modularisation of offerings as enabling factors for rapid and scalable entry (Eskiyerli, 2024).

4.2 Entry as an activity bundle rather than a choice among discrete modes

Contemporary academic literature repositions "entry mode" from selecting discrete governance forms (export, license, JV, wholly owned affiliate) to coordinating a mixture of activities (marketing, distribution, regulatory engagement, digital deployment). This perspective recognises that firms increasingly combine activities (e.g., digital market testing → local partnership → contract manufacturing) and that the concept of entry may be more dynamic and hybrid than earlier practices. Brouthers (2022) highlights the need to study the broader set of activities firms undertake during entry and how they fit together strategically.

4.3 Born-globals and international new ventures (INVs)

Recent syntheses and empirical work have revisited born-global and INV phenomena, emphasising different mechanisms of foreign market entry in earlier internalisation processes and pointing to alternative mechanisms for rapid international entry (entrepreneurial networks, parent-firm relationships, digital market access). Scholars argue that the original typology (born-global vs. gradual internationaliser) is too coarse: many early internationalisers differ by home-country context, founder experience and embedded networks, which shape entry mode and performance outcomes. This reconceptualisation has important implications for when Uppsala-style learning explanations apply, as they do not capture the full set of mechanisms for rapid internationalisation (Rumyantseva, 2023).

5. Future Research Directions

Given the preceding discussion, the following promising research directions are evident:

- Mechanisms of digital market access: identify how specific platform governance features (data portability, algorithmic match quality) mediate entry success.
- Sequencing and timing: empirically model how firms sequence digital testing, partnerships and investment, and which sequences perform best in which contexts.
- Liability of foreignness under political volatility: test how sudden political changes (trade restrictions, sanctions) reshape liability of foreignness and whether pre-emptive nonmarket investments buffer firms (Li et al., 2022).

- Integration of marketing and international business (IB) theory: investigate how local market positioning and branding strategies interact with ownership and governance choices (Rašković, 2025).

- Comparative born-global pathways: leverage cross-country samples to understand how home-country institutional contexts shape rapid internationalisation and subsequent growth (Rumyantseva, 2023).

6. Conclusion

Foreign market entry research has shifted beyond the static and traditional mode choice models toward richer, more dynamic conceptions that incorporate digitalisation, political economy, and entrepreneurial heterogeneity. Recent studies have usefully highlighted entry as a bundle of activities, the differentiated pathways of born-globals, and the evolving nature of liability of foreignness. However, the literature would benefit from greater theoretical integration, more rigorous longitudinal and causal methods, standardised operationalisations of hybrid modes, and stronger incorporation of nonmarket dynamics. Addressing these gaps will yield findings that are both academically cumulative and practically actionable for managers navigating increasingly complex international entry environments.

References

Agarwal, S. (1994). Socio-cultural distance and the choice of joint ventures: A contingency prospective. *Journal of International Marketing*, 2(2), pp. 63–80.

Ahsan, M. and Musteen, M. (2011a). Multinational enterprises' entry mode strategies and uncertainty: A review and extension. *International Journal of Management Reviews*, 13, pp. 376–392.

Ahsan, M. and Musteen, M. (2011b). Part IV: ethics, learning, and change: flexibility, learning and control: entry mode choice under uncertainty. *Current Topic in Management*, 15, pp. 199–224.

Angeles, G.M. and Ramos, H. (2009). The relationship between timing of entry into foreign market, entry mode decision and market selection. *Time & Society*, 18(2-3), pp. 306–331.

Brouthers, K.D. (1995). The influence of international risk on entry mode strategy in the computer software industry. *Management International Review*, 35(1), pp. 7–28.

Brouthers, K.D. (2002). Institutional, cultural and transaction cost influence on entry mode choice and performance. *Journal of International Business Studies*, 33(2), pp. 203–221.

Brouthers, K.D. (2013a). Institutional, cultural and transaction cost influences on entry mode choice and performance. *Journal of International Business Studies*, 44(1), pp. 1–13.

Brouthers, K.D. (2013b). A retrospective on: institutional, cultural and transactional cost influences on entry mode choice and performance. *Journal of International Business Studies*, 44(1), pp. 14–22.

Brouthers, K.D. (2022). Charting new courses to enter foreign markets: On the need to study entry activities as bundles. *Journal of International Management*, 28(3), 100908. https://doi.org/10.1016/j.intman.2022.100908

Brouthers, K.D. and Brouthers, L.E. (2001). Explaining the national culture distance paradox. *Journal of International Business Studies*, 32(1), pp. 177–189.

Brouthers, K.D. and Brouthers, L.E. (2002). Does international entry mode choice influence firm performance? *28th Annual Conference of the Academy of International Business*, pp. 194–204.

Brouthers, K.D. and Brouthers, L.E. (2003). Why service and manufacturing entry mode choices differ; the influence of transaction cost factors, risk and trust. *Journal of Management Studies*, 40(5), pp. 1179–1204.

Brouthers, K.D. and Dikova, D. (2010). Acquisitions and real options: The Greenfield alternative. *Journal of Management Studies*, 47(6), pp. 1048–1071.

Brouthers, K.D. and Hennart, J.F. (2007). Boundaries of the Firm: Insights from International Entry Mode Research. *Journal of Management*, 33(3), pp. 395–425.

Brouthers, K.D., Brouthers, L.E. and Werner, S. (2000). Influences on strategic decision making in the Dutch financial services industry. *Journal of Management*, 26, pp. 863–883.

Brouthers, K.D., Brouthers, L.E. and Werner, S. (2002). Industrial sector, perceived environmental uncertainty and entry mode strategy. *Journal of Business Research*, 55(6), pp. 495–507.

Brouthers, K.D., Brouthers, L.E. and Werner, S. (2003). Transaction cost enhanced entry mode choices and firm performance. *Strategic Management Journal*, 24(12), pp. 1239–1248.

Brouthers, K.D., Brouthers, L.E. and Werner, S. (2008). Real Options, International Entry Mode Choice and Performance. *Journal of Management Studies*, 45(5), pp. 936–960.

Brouthers, L.E., Brouthers, K.D. and Nakos, G. (2003). How national culture influences international entry mode choice. *30th Conference of Academy of International Business*, pp. 107–120.

Brouthers, L.E., Brouthers, K.D. and Werner, S. (1999). Is Dunning's eclectic framework descriptive or normative? *Journal of International Business Studies*, 30, pp. 831–844.

Brouthers, L.E., Brouthers, K.D. and Werner, S. (2000). Perceived environmental uncertainty, entry mode choice and satisfaction with EC-MNC performance. *British Journal of Management*, 11(3), pp. 183–195.

Buckley, P. and Casson, M. (1998). Analysing foreign market entry strategies: extending the internalization approach. *Journal of International Business Studies*, 29(3), pp. 539–561.

Canabal, A. and White, G.O. (2008). Entry mode research: past and future. *International Business Review*, 17, pp. 267–284.

Cao, M. (2021). Overcoming the liability of foreignness – A new perspective. *Journal of International Business Studies,* 52(5), pp. 879–901. https://doi.org/10.1057/s41267-020-00356-y

Carpenter, M.A., Geletkanycz, M.A. and Sanders, W.G. (2004). Upper echelons research revisited: Antecedents, elements and consequences of top management team components. *Journal of Management*, 30(6), pp. 749–778.

Castaner, X. and Genc, M.E. (2011). Expanding to institutionally different countries; reasons, firm international experience and entry mode choice. Strategic Alliances in a Globalizing World. *Chapter in Research in Strategic Alliances*, pp. 1–25.

Cieslik, A. and Ryan, M. (2012). Productivity differences and foreign market entry in an oligopolistic industry. *Open Economies Review*, 23(3), pp. 531–557.

Combs, J.G. and Ketchen, D.J. (2021). Franchising research: Major milestones and future directions. *Journal of Management,* 47(6), pp. 1412–1436. https://doi.org/10.1177/01492063211004969

Couturier, J. and Sola, D. (2010). International market entry decisions: the role of local market factors. *Journal of General Management*, 35(4), pp. 45–63.

Cumberland, F. (2006). Theory Development within International Market Entry Mode: An Assessment. *Marketing Review*, 6(4), pp. 349–373.

Datta, D., Liang, X. and Musteen, M. (2009). Strategic Orientation and the choice of foreign market entry mode. *Management International Review*, 49(3), pp. 269–290.

Datta, D.K. and Beamish, P.W. (2009). Discretion and internationalization: impact of environmental determinants of managerial discretion and host country experience on entry mode choice. *International Journal of Trade and Global Markets*, 2(3), pp. 267–285.

Datta, D.K., Herrmann, P. and Rasheed, A.A. (2002). Choice of foreign market entry modes: critical review and future directions. *Advances in International Comparative Management*, 14, pp. 85–153.

Davies, R.B., Egger, H. and Egger, P. (2010). Profit taxation and the mode of foreign market entry. *Canadian Journal of Economics*, 43(2), pp. 704–727.

Diego, Q., Enrique, C. and Laura, R. (2007). The impact of country risk and cultural distance on entry mode choice: an integrated approach. *Cross Cultural Management*, 14(1), pp. 74–87.

Dikova, D. (2012). Entry mode choices in transition economies: the moderating effect of institutional distance on managers' personal experiences. *Journal of East-West Business*, 18(1), pp. 1–27.

Dixon-Woods, M., Agarwal, S., Jones, D., Young, B. and Sutton, A. (2005). Synthesising qualitative and quantitative evidence: a review of possible methods. *Journal of Health Services Research Policy*, 10(1), pp. 45–53.

Dow, D. and Larimo, J. (2009). Challenging the conceptualization and measurement of distance and international experience in entry mode choice research. *Journal of International Marketing*, 17(2), pp. 74–98.

Dunning, J.H. (1980). Toward an eclectic theory of international production: Some empirical tests. *Journal of International Business Studies*, 11(1), pp. 9–31.

Ekeledo, I. and Sivakumar, K. (2001). Foreign market entry mode choice of manufacturing firms and service firms: a resource based perspective. *AMA Winter Educators' Conference Proceedings*, 12, p. 282.

Elbanna, S. and Child, J. (2007). The influence of decision, environmental and firm characteristics on the rationality of strategic decision making. *Journal of Management Studies*, 44(4), pp. 561–591.

Erramilli, M.K. and Rao, C.P. (1990). Choice of foreign market entry modes by service firms: Role of market knowledge. *Management International Review*, 3(2), pp. 135–150.

Eskiyerli, D.Y. (2024). A quantitative approach to foreign market selection for B2B digital platforms. *Journal of Business Research*, 167, 114021. https://doi.org/10.1016/j.jbusres.2023.114021

Fincham, R. (1992). Perspectives on power: Processual, institutional, and internal forms of organizational power. *Journal of Management Studies*, 29(6), pp. 741–759.

Fink, A. (2010). *Conducting research literature reviews: From the internet to paper* (3rd ed.). Los Angeles; London: SAGE.

Gannon, M. (1993). Towards a composite theory of foreign market entry mode choice: the role of marketing strategy variables. *Journal of Strategic Marketing*, 1(1), pp. 41–54.

Georgiou, M. (2014). Foreign market entry and exit process: The case of BP Lubricants in Europe and Africa. *DBA Thesis, University of Gloucestershire*.

Hennart, J. (2009). Down with MNE-centric theories! Market entry and expansion as the bundling of MNE and local assets. *Journal of International Business Studies*, 40(9), pp. 1411–1431.

Hennart, J.F. and Reddy, S. (1997). The choice between mergers/acquisitions and joint ventures: The case of Japanese investors in the United States. *Strategic Management Journal*, 18, pp. 1–12.

Hermann, P. and Datta, D. (2002). CEO successor characteristics and the choice of foreign market entry mode: an empirical study. *Journal of International Business Studies*, 33(3), pp. 551–569.

Hermann, P. and Datta, D. (2006). CEO experiences: effects on the choice of FDI entry mode. *Journal of Management Studies*, 43(4), pp. 755–768.

Hofstede, G. (1984). The cultural relativity of the quality of life concept. *Academy of Management Review*, 9, pp. 389–398.

Johanson, J. and Vahlne, J.E. (1977). The internationalization process of the firm: A model of knowledge development and increasing foreign commitments. *Journal of International Business Studies,* 8(1), pp. 23–32.

Johanson, J. and Vahlne, J.E. (2009). The Uppsala internationalization process model revisited: From liability of foreignness to liability of outsidership. *Journal of International Business Studies,* 40(9), pp. 1411–1431.

Jones, M.V. and Young, S. (2009). Does entry mode matter? Reviewing current themes and perspectives. In Jones, M.V., Dimitratos, P., Fletcher, M. and Young, S. (eds). *Internationalization, entrepreneurship and the smaller firm – evidence around the world.* Cheltenham, Edward Elgar, pp. 6–9.

Li, C., Wang, T. and Chen, Z. (2022). Political ideology and the liability of foreignness: Evidence from multinational subsidiaries. *Journal of International Management*, 28(1), 100876. https://doi.org/10.1016/j.intman.2021.100876

Lopez, D.C. and Vidal, S.M. (2010). External uncertainty and entry mode choice: cultural distance, political risk and language diversity. *International Business*, 19(6), pp. 575–588.

Madhok, A. (1994a). Mode of foreign market entry: the transaction and the firm. *World Business Congress, 3rd Conference*, Malaysia, pp. 755–757.

Madhok, A. (1994b). Understanding the foreign market entry decision of multinational firms: linking theory and practice. *World Business Congress*, pp. 758–760.

Madhok, A. (1996). Know-how, experience and competition related considerations in foreign market entry: an exploratory investigation. *International Business Review*, 5(4), pp. 339–366.

Madhok, A. (1997). Cost, value and the foreign market entry mode: the transaction and the firm. *Strategic Management Journal*, 18(1), pp. 39–61.

Madhok, A. (1998). The nature of multinational firm boundaries: transaction costs, firm capabilities and foreign market entry mode. *International Business Review*, 7(3), pp. 259–291.

Maekelburger, B., Schwens, C. and Kabst, R. (2012). Asset specificity and foreign market entry mode choice of small and medium sized enterprises: the moderating influence of knowledge safeguards and institutional safeguards. *Journal of International Business Studies*, 43(5), pp. 458–476.

Markusen, J. and Stahler, F. (2011). Endogenous market structure and foreign market entry. *Review of World Economics*, 147(2), pp. 195–215.

Martin, X. (2013). Solving theoretical and empirical conundrums in international strategy research: linking foreign market entry mode choices and performance. *Journal of International Business Studies*, 44(1), pp. 28–41.

Martinez, P. and Lopez, J. (2009). Making foreign market entry decisions. *Global Business & Organisational Excellence*, 20(2), pp. 52–59.

Merriam, S.B. (2009). *Qualitative Research: A Guide to Design and Implementation*. San Francisco: Jossey-Bass.

Meyer, K.E. (2001). Institutions, transaction costs and entry mode choice in Eastern Europe. *Journal of International Business Studies*, 32(2), pp. 357–368.

Miles, M.B. and Huberman, A.M. (1994). *Qualitative Data Analysis: An Expanded Sourcebook (2nd ed.)*. Thousand Oaks, Calif.; London: Sage.

Morschett, D., Schramm-Klein, H. and Swoboda, B. (2010). Decades of research on market entry modes: what do we really know about external antecedents of entry mode choice? *Journal of International Management*, 16(1), pp. 60–77.

Musteen, M., Datta, D.K. and Herrmann, P. (2009). Ownership structure and CEO compensation: implication for the choice of foreign market entry modes. *Journal of International Business Studies*, 40(2), pp. 321–338.

Nakos, G. and Brouthers, K. (2002). Entry Mode Choice of SMEs in Central and Eastern Europe. *Entrepreneurship Theory and Practice*, 27, pp. 47–63.

Omar, M. and Porter, M. (2011). Reducing risk in foreign market entry strategies: standardization versus modification. *Competitiveness Review*, 21(4), pp. 382–396.

Pan, Y. and Tse, D.K. (2000). The hierarchical model of market entry modes. *Journal of International Business Studies*, 31, pp. 535–554.

Pope, C., Mays, N. and Popay, J. (2007). *Synthesizing Qualitative and Quantitative Health Evidence*. McGraw-Hill, Open University Press.

Rašković, M.M. (2025). Current state of the literature and new research directions: Closer integration of marketing and international business theory. *Journal of Business Research,* 172, 113456. https://doi.org/10.1016/j.jbusres.2023.113456

Rivera, J. and Oh, C.H. (2013). Environmental regulations and multinational corporations' foreign market investments. *Policy Studies Journal*, 41(2), pp. 243–272.

Root, F.R. (1987) *Entry strategies for international markets*. Lexington, MA: Lexington Books.

Rumyantseva, M. (2023). The born global and international new venture revisited: A literature review and research agenda. *Journal of International Business Studies*, 54(2), pp. 233–259. https://doi.org/10.1057/s41267-022-00554-8

Samiee, S. (2013). International market entry mode decisions: cultural distance's role in classifying partnerships versus sole ownership. *Journal of Business Research*, 66(5), pp. 659–661.

Sarkar, M. and Cavusgi, S. (1996). Trends in International business thought and literature: a review of international market entry mode research: Integration and synthesis. *International Executive*, 38(6), pp. 825–847.

Sharma, V.M. and Erramilli, M.K. (2004). Resource-based explanation of entry mode choice. *Journal of Marketing Theory and Practice*, 12(1), pp. 1–18.

Simon, D.G. and Lane, P.J. (2004). A model of cultural differences and international alliance performance. *Journal of International Business Studies*, 35(4), pp. 306–319.

Simons, H. (2009). *Case Study research in practice*. Sage.

Stevens, C.E. and Dykes, B.J. (2013). The home country cultural determinants of firms' foreign market entry timing strategies. *Long Range Planning*, 46, pp. 387–410.

Tihanyi, L., Griffith, D. and Russell, C. (2005). The effect of cultural distance on entry mode choice, international diversification and MNE performance: a meta-analysis. *Journal of International Business Studies*, 36(3), pp. 270–283.

Werner, S. (2002). Recent developments in international management research: a review of 20 top management journals. *Journal of Management*, 28, pp. 277–305.

Young, S., Hamill, J., Wheeler, C. and Davies, J.R. (1989). *International Market Entry and Development*. Englewood Cliffs: Prentice Hill.

Zhao, H., Luo, Y. and Suh, T. (2004). Transaction cost determinants and ownership-based entry mode choice: A meta analytical review. *Journal of International Business Studies*, 35(6), pp. 524–544.

Performance Management: The good, the bad and the ugly

By

Dr Zano Tyrannis

M.Tech. Chiro., M.Sc. Med. (Pharm.), MBA | Human Potential Promoter

Abstract

This article critically examines the evolution of performance management in today's workplace and provides an examination of where we are currently in terms of performance management practices. It will draw upon relevant literature and case studies of modern companies to explore the "**good**" aspects of performance management systems (i.e., continuous feedback, using data to inform decision-making and aligning with OKRs); the "**bad**" aspects of performance management systems (i.e., bias, over reliance on rating scales, and managers' lack of capabilities); and the "**ugly**" aspects of performance management (i.e., use as a tool to downsize, as a means of surveillance via algorithms, and a way of ignoring employees' well-being). In addition, the review will examine how companies such as Spotify, 3M and Klarna have successfully implemented and/or abused PMS to enhance and/or impair employee engagement and company culture. Finally, the review paper will focus on the Cyprus context, which presents both a challenge and an opportunity in the form of its organisations' reliance on annual evaluations and willingness to embrace new and innovative ways of implementing performance management systems.

In conclusion, the review provides a number of guiding principles for improving performance management systems including: (1) focusing on development rather than assessment; (2) emphasising coaching as opposed to compliance; (3) ensuring transparency when utilising artificial intelligence; (4) fostering skill-based growth; and (5) transforming performance management into a human-centred approach to enable trust, innovation and organisational longevity.

Keywords: Performance Management (PM), Continuous Feedback, Objectives and Key Results (OKRs), People Analytics, Skill-Based Performance, Workplace Bias, Culture.

Introduction

Performance management, as defined by the Chartered Institute of Personnel and Development (CIPD, 2023), is "*a continuous process which aims to improve organisational performance by developing the performance of individuals and teams*." This definition underscores the shift from traditional, annual evaluations toward a more holistic, growth-oriented approach — one that aligns purpose, progress and performance.

Performance management is frequently referred to as the HR function "that everyone despises but no one can avoid." (Cappelli and Tavis, 2016). Organisations have been using the same methods to assess employee performance and to determine whether an employee receives a raise or promotion for nearly four decades. This has included annual reviews, numerical ratings and forced rank order lists. With the emergence of hybrid work models, multi-generational teams, AI tools and skills-driven roles, this method of assessing performance has outlived its usefulness and, in fact, is becoming both outdated and harmful.

In fact, in 2024, Gartner's research revealed that current talent management efforts often inhibit optimal employee and organisational performance (Gartner Inc., 2024a). This finding highlights a growing gap between the systems we use and the culture of engagement and agility we seek to build.

Many organisations using performance management systems have absolutely no idea why they are using them. The simple answer that they would give on questioning is that everybody else is doing it and it's a regulatory requirement which, by the way, it is not.

On deeper probing, the HR department might say it is to develop staff skills, to weed out poorly performing employees or to use to consider promotions and/or bonuses.

In the author's opinion, the first step is establishing *why* a performance management system should be implemented and what purpose it should serve both in terms of staff development and remuneration, etc.

Secondly, it is important to ask if all members of the organisation should have the same performance management system. Is it worthwhile or even fair to measure C-suite members of the organisation in the same way as, say, production line workers?

Thirdly, does the tenure of the employee influence the decision? Would new employees be measured on the same metrics and in the same way as employees that have been with the organisation for a longer time?

Alongside these considerations, organisations need to think about the negative implications of a performance management system in terms of culture, staff morale and retention.

The good: What's going well in performance management?

When done well, these systems improve communication, align individual goals with business strategy, and support leadership development (Gartner Inc., 2023). Continuous feedback loops and clear objectives also enhance employee motivation and performance clarity (Gouldsberry, 2022).

Positive developments in performance management include:

Continuous feedback and coaching

Replacing traditional, annual performance evaluations with continuous, informal feedback and regular check-ins has been a useful development. Regular feedback and check-ins provide timely alignment, allow employees to adjust their efforts quickly and offer clearer understanding of expectations. Many organisations have abandoned lengthy performance appraisals in favour of more frequent and more balanced check-ins, either after a project or a regular short time interval. For example, in Cyprus, weekly or bi-weekly structured check-ins are being implemented to supplement formal evaluations (Emerald Zebra, 2025). However, these need to be relevant and fair. For instance, one of the big four consulting companies used a spot check appraisal that was based on only four questions:

- Given what I know of this person's performance, and if it were my money, I would award this person the highest possible compensation increases and bonus. (Scale of 1 to 5)
- Given what I know of this person's performance, I would always want him or her on my team. (Scale of 1 to 5)
- This person is at risk of low performance. (Yes/No)
- This person is ready for promotion today. (Yes/No)[1]

Naturally, the above approach would be highly biased against a new employee.

[1] *For more info*: https://www.coursehero.com/tutors-problems/Human-Resource-Management/36959784-PERFORMANCE-MANAGEMENTTransforming-Performance-Management-The/.

Goal frameworks that drive focus: Objectives and Key Results (OKRs) and SMART goals

Setting performance goals tied to company-wide outcomes and employee skills in terms of OKRs rather than generic "team player" type of goals is a useful method that tends to be less biased and focused on skills development rather than talent inadequacies. Performance review methodologies here may take the following forms:

- An increase in online/digital training tools
- Gamification
- Personalised training.

For instance, SPOTIFY and LIVING SOCIAL replaced annual reviews with a mobile, gamified solution and over 90% of employees participated voluntarily (Bersin, 2014).

Data and people analytics

Using data and analytics to inform performance decisions: Companies are leveraging people-analytics tools and dashboards to track goal achievement, volume of feedback received and development needs - all of which bring a more objective lens to performance assessments.

What gets measured gets improved, but is it the right measure for the right behaviour?

These new approaches represent a more human-centred view: enabling employees to be successful rather than simply evaluating them. As artificial intelligence (AI) develops and more tasks can be automated, more emphasis will be placed on softer skills and relatedness. For instance, during the late 2000s and early 2010s many high street bank branch staff were replaced by a variety of machines. It was possible to walk in and find a bank of ATM's and specialised "ATM-like" machines for deposits, requesting statements, etc. The public, however, did not embrace this technological shift. It seems that when it comes to money, people still prefer interacting with people. A similar situation played out recently when Klarna CEO Sebastian Siemiatkowski admitted that people did not like engaging with AI chats and still wanted to explain their problem and get a resolution from an actual person, according to Entrepreneur (2024).

The bad: Challenges that remain

Many of the previous challenges continue to exist:

"Every single human alive today is a horribly unreliable rater of other human beings." (Buckingham, 2015)

Bias and fairness questions:

According to the HRM Guide (n.d.) and Menzies (2025), ratings which are thought to be more objective as they are normally target or task based are still subject to influences such as recency, manager memory, visibility and unconscious proximity bias (which may be exacerbated in hybrid environments). This is due to the fact that people generally tend to remember more negative aspects and experiences as opposed to positive ones. An employee who has delivered well on nine projects will be penalised for the one they didn't deliver well on. There is no averaging out in performance assessments, you are normally only as good as your last piece of work!

These types of assessments are particularly problematic for new employees where they may under-perform in early appraisals not because of lack of skill, but because their onboarding neglected cultural integration and network building—a gap that often leads to failure in the role (Watkins, 2017). A lack of cultural integration causes new employees to not be aware of "How things are actually done in the organisation" because often the espoused cultural values are inconsistent with the actual values within the organisation. A common example is "we value creativity", and new employees realise that what the company means is creativity within the constraints of a very rigid framework. This is known as a Culture Gap (HR Brain, 2024).

An immature internal network can also be a barrier to a satisfactory performance appraisal because the employee may be unaware of who the "power players" are and the internal political agendas and motives. Takeuchi et al. (2022) demonstrated that in organisations with strong political agendas the less "politically savvy" employees tend to do less favourably as opposed to those who are politically savvy and earn credit and favourable evaluations.

In addition, some companies have built in policies which discourage awarding employees maximum scores, as they feel it prevents the employee being reviewed on growth targets that they could aspire to.

Rating politics rather than growth:

Using forced rankings or rating curves creates behaviours that please managers rather than behaviours that create value, generate innovation or solve problems. Traditional productivity ratings often fail to capture the full spectrum of employee contribution because they narrowly measure output against predefined objectives, leaving little room for creativity, experimentation or intrapreneurial thinking. As such, many traditional measurement systems may unintentionally inhibit innovative ideas through the emphasis placed on short term productivity (efficiency) over long term value creation (Amabile, 1998). When employees are evaluated and rewarded solely on quantifiable metrics, they are less likely to take risks or engage in exploratory projects that may not yield immediate results, but have the potential to drive major breakthroughs.

The importance of fostering a culture of innovation and growth:

Many companies treat performance reviews as administrative burdens, rather than opportunities for growth. Gartner (2024b) notes that rigid systems stifle innovation and discourage honest dialogue between managers and employees. As a result, employees feel undervalued, managers feel overburdened, and HR teams chase compliance rather than culture. A great example of a company that encourages innovation and growth is 3M and its "15% Culture," which encourages employees to spend 15% of their time on projects that they choose. This policy has led to a number of successful new products over the years, including Post-it notes (Wright, 2025). Many companies would regard this project activity as inefficient and wasteful.

Untrained or under-skilled managers:

Many managers are not equipped as coaches or feedback providers. Unless those skills are developed, the system will fail. Generally, in most occupations and organisations, competency in the individual results in their promotion to higher levels of seniority. However as one moves up in an organisation, the technical skills of the role are less important than the people management abilities. A great software engineer seldom makes a good manager unless they have developed the correct people skills.

The ugly: New challenges requiring immediate focus

Downsizing via performance systems:

In some global corporations, performance frameworks are being utilised to increase the number of "below expectation" ratings to justify headcount reductions or "trimming the fat", so to speak, resulting in a lack of trust and disengagement. Nothing kills a culture faster! Ironically the star performers are the first ones to leave in such instances, as they get the feeling that their time for removal will come too.

AI/automation missteps:

As technology becomes more advanced, it could be possible to use measures of keystroke, email response or tool utilisation to measure performance rather than actual impact, potentially contributing to a culture of surveillance, algorithmic bias and the erosion of trust. Employees may be productive but disengaged. This is likely to lead to high staff turnover and any efficiencies gained from productivity would be quickly erased by increased recruitment and training costs. According to Gardiner (2024) these decisions need to be evaluated in terms of three foundational building blocks, namely employee mindset, manager effectiveness and quality performance data.

Ignoring wellbeing and human context:

When output metrics dominate, hidden costs, such as burnout, turnover and disengagement arise. Sure, machines and AI can be more productive in certain areas, but a human is still required in most instances. These humans need to be developed and know that their wellbeing is at the heart of the organisation. A recent Cyprus Business News (2025) commentary observed that "businesses in Cyprus spend much less time developing employees than they do evaluating them".

Cyprus-specific local context:

Cyprus organisations are feeling the pressure more than ever. Many are still operating with outdated systems that simply don't match the realities of modern work. Performance judgments often default to old metrics—face-time, tenure, visibility—even though they no longer predict effectiveness. As hybrid work expands, talent becomes scarcer, and younger generations demand transparency and fairness, Cypriot businesses are being forced into a much-needed evolution of how performance is understood, measured and rewarded. A recent article noted: "In Cyprus, performance management appraisal systems rooted in annual reviews remain common... yet 72% of workers and 61% of managers globally express distrust." (Cyprus Business News, 2025).

In another review, the **Creating People Advantage 2023** report (Cyprus country results), performance management was rated as a **priority**, but only **38% of Cypriot HR professionals** considered it **"very important"** compared to other HR functions. This reflects a gap between strategic recognition and perceived practical importance (BCG and WFPMA, 2023).

However, Cyprus still relies heavily on legacy performance-management systems, particularly in the public sector, where traditional appraisal processes remain deeply entrenched. For example, the Water Development Department (WDD) of the Cypriot civil service continues to use a manager-driven annual rating system, with research showing tendencies toward excessive positive ratings and concerns about transparency and meritocracy (Christou-Tsangarides, 2015). Other legacy systems in use are: Supervisor-Only Evaluation / Top-Down Feedback, Balanced Scorecard / KPI-Driven Legacy Models and Public Sector Appraisal Systems.

This combination of change pressure and legacy systems gives Cyprus organisations a unique opportunity: to leapfrog outdated models, rather than simply upgrade them. As many organisations and even government departments are modernising many of their systems, there is an opportunity to not only meet current standards but exceed them by building in the latest best practices as well as future trends.

Future trends

The following are future trends that are likely to shape performance management.

Enabling performance through AI:

The introduction of tools that assist in recording employee performance as well as suggesting developmental paths, personalised learning and coaching nudges. (However, transparency is essential here.)

Skill-based performance:

Moving away from performance based solely upon job titles to performance based upon skill taxonomies and career ladders. Better alignment of employee growth with organisational capacity.

Hybrid inclusive performance:

Ensuring that remote employees are assessed on the basis of performance outcomes rather than on the basis of visibility; reducing proximity bias through calibrating ratings.

Maturity of people analytics:

From simple reporting dashboards to predictive insights - e.g., identifying performance declines, coaching needs, retention risks.

A better path forward and guiding principles

Organisations willing to transform their performance management systems into something new and innovative should consider the following six fundamental principles:

1. Developmental support before assessment: Frame the system as “I want to help you grow” not “I need to monitor you”.
2. Regular, structured dialogue: Replace the annual performance review with regular check-in and developmental dialogues.
3. Performance outcomes and skills, not presence: Assess what employees produce and learn, not when or where they perform.
4. Transparent and ethical use of data: If you utilise AI or analytics, explain how you collect and process metrics - employees should be aware of how the metrics are collected and processed.
5. Train **managers as coaches**: Provide managers with the training and time to offer meaningful feedback and support.

6. Assess performance in terms of business outcomes, but maintain long-term behaviour: Tie performance goals to business outcomes by all means, but encourage behaviours that promote creativity, innovation, learning and growth - Do not simply use performance goals to achieve short-term key performance indicators (KPIs).

Implementation checklist for Cyprus organisations

- Identify the primary intention of your performance system (development vs. calibration vs. compensation).
- Replace or enhance the annual review with a quarterly or monthly structured check-in.
- Develop or update a skills taxonomy aligned to your strategic vision.
- Test pilot AI/people-analytics tools for transparency, and engage employees in the design process.
- Train managers in coaching, providing feedback and recognising and avoiding biases.
- Monitor the equity and inclusiveness of your performance system (for example, compare remote to on-site ratings) and hold calibration forums.
- Track the overall health of your performance system: employee engagement, perceived fairness, retention and developmental activities.

Conclusion: Closing thoughts

Performance management systems require significant evolution due to rapid changes in the world of work. For organisations in Cyprus, the decision is clear: upgrade or fall behind. The most effective performance management systems do not simply evaluate - they enable, empower and develop employee capabilities. By treating performance as a human system and not a scorecard, superior performance will be the focus and drive. This will cultivate a greater level of trust, agility and a growth mindset that today's workforce requires.

Disclaimer

The author wishes to acknowledge that AI tools were used in researching this article in the following ways: ChatGPT and Microsoft Copilot were used to search for appropriate examples and case studies of current and future trends.

References

Amabile, T.M. (1998) 'How to Kill Creativity', *Harvard Business Review*, 76(5), pp. 76–87. Available at: https://hbr.org/1998/09/how-to-kill-creativity (Accessed: 2 November 2025).

BCG and WFPMA (2023) 'Creating People Advantage 2023: Country Results – Cyprus', Cyprus Human Resource Management Association. Available at: https://www.cyhrma.org/wp-content/uploads/2024/09/CPA_2023_Country-Results_Cyprus_vS.pdf (Accessed: 2 November 2025)

Bersin, J. (2014) 'The Myth of the Bell Curve: Look for the Hyper-Performers', *Forbes*, 19 February. Available at: https://www.forbes.com/sites/joshbersin/2014/02/19/the-myth-of-the-bell-curve-look-for-the-hyper-performers/ (Accessed: 15 November 2025).

Buckingham, M. (2015) 'Most HR Data Is Bad Data', *Harvard Business Review*, 9 February. Available at: https://marcusbuckingham.co/wp-content/uploads/2017/08/Most-HR-Data-Is-Bad-Data-HBR.pdf (Accessed: 2 November 2025).

Cappelli, P. and Tavis, A. (2016) 'The Performance Management Revolution', *Harvard Business Review*, October 2016. Available at: https://hbr.org/2016/10/the-performance-management-revolution (Accessed: 2 November 2025).

Christou-Tsangarides, K. (2015) *Performance appraisal at the Water Development Department of Cyprus*. Master's thesis. Cyprus: Neapolis University. Available at: https://hephaestus.nup.ac.cy/handle/11728/6955 (Accessed:16 November 2025).

CIPD (2023) 'Performance Management: An Overview', Chartered Institute of Personnel and Development. Available at: https://www.cipd.org/uk/knowledge/factsheets/performance-management-factsheet/ (Accessed: 2 November 2025).

Cyprus Business News (2025) 'Unlocking human potential: Modernising performance management for business success'. Available at: https://www.cbn.com.cy/article/116828/unlocking-human-potential-modernising-performance-management-for-business-success (Accessed: 2 November 2025).

Emerald Zebra (2025) 'Optimizing Performance Management in Remote and Hybrid Work in Cyprus', 24 April. Available at: https://emeraldzebra.cy/optimizing-performance-management-in-remote-and-hybrid-work-in-cyprus/ (Accessed: 10 November 2025).

Entrepreneur (2024) 'Klarna CEO reverses course by hiring more humans, not AI', 22 April. Available at: https://www.entrepreneur.com/business-news/klarna-ceo-reverses-course-by-hiring-more-humans-not-ai/491396 (Accessed: 2 November 2025).

Gardiner, L. (2024) 'HR Leaders Should Evaluate Three Foundational Building Blocks Before Implementing AI in Performance Management', Gartner Newsroom, 28 October. Available at: https://www.gartner.com/en/newsroom/press-releases/2024-10-28-hr-leaders-should-evaluate-three-foundational-building-blocks-before-implementing-ai-in-performance-management (Accessed: 3 November 2025).

Gartner Inc. (2023) 'HR Leaders Survey Reveals Top Two Priorities in 2024 Are Leader and Manager Development and Organisational Culture', Gartner Newsroom, 7 November. Available at: https://www.gartner.com/en/newsroom/press-releases/2023-11-07-gartner-hr-leaders-survey-reveals-top-two-priorities-in-2024 (Accessed: 3 November 2025).

Gartner Inc. (2024a) 'HR Research Finds Organisations' Current Talent Management Efforts Inhibit Optimal Employee and Organisational Performance', Gartner Newsroom, 18 September. Available at: https://www.gartner.com/en/newsroom/press-releases/2024-09-18-gartner-hr-research-finds-orgs-current-talent-management-efforts-inhibit-optimal-employee-and-org-performance (Accessed: 3 November 2025).

Gartner Inc. (2024b) 'Performance Management: The Complete Guide and Essential Tips for HR'. Available at: https://www.gartner.com/en/human-resources/topics/performance-management (Accessed: 3 November 2025).

Gouldsberry, M. (2022) 'A Practical Guide to Continuous Feedback at Work', Betterworks, 28 October. Available at: https://www.betterworks.com/magazine/continuous-feedback/ & https://d3bql97l1ytoxn.cloudfront.net/app_resources/371763/documentation/1285691_1668454291017_en-US.pdf (Accessed: 18 November 2025).

HR Brain (2024) 'Gaps between desired culture and actual culture: analysis', HR Brain blog, 12 February. Available at: https://hrbrain.ai/blog/gaps-between-desired-culture-and-actual-culture-analysis/ (Accessed: 15 November 2025).

HRM Guide (n.d.) 'Bias in Performance Evaluation: Identifying and Preventing It'. Available at: https://hrmguide.io/essentials/ethical-decisions/performance-bias-in-hr/ (Accessed: 3 November 2025).

Menzies, F. (2025) 'Managing Bias in Performance Appraisal', LinkedIn, 11 August. Available at: https://www.linkedin.com/pulse/managing-bias-performance-appraisal-felicity-menzies-tyl5c (Accessed: 3 November 2025).

Takeuchi, R., Yu, N.Y. and Lin, C-C. (2022) 'Organizational politics and employee performance in the workplace', *Journal of Vocational Behavior*, 133. Available at: https://www.sciencedirect.com/science/article/abs/pii/S0001879121001494 (Accessed: 15 November 2025).

Watkins, M.D. (2017) 'Does your company really do a good job of onboarding new hires?' IMD Insight article. Available at: https://www.imd.org/research-knowledge/talent-management/articles/does-your-company-really-do-a-good-job-of-onboarding-new-hires/ (Accessed: 3 November 2025).

Wright, L. (2025) 'The Power Of Stepping Aside: Empowering Employees To Lead', *Forbes*, 18 June. Available at: https://www.forbes.com/councils/forbesbusinesscouncil/2025/06/18/the-power-of-stepping-aside-empowering-employees-to-lead/ (Accessed: 3 November 2025).

Nation Branding, Geopolitics and the Defence Industry: The Case of the Republic of Cyprus How the Cyprus Defence Industry could contribute to the formation of a new reinforced Nation Brand

By

Dr Panayiotis Hadjipavlis

CERIDES – Excellence in Innovation and Technology,

European University Cyprus, Research Fellow

Dr Constantinos Constantinou

Director of Research at CIM-Cyprus Business School, Nicosia – Limassol

Abstract

This synoptical commentary presents the growing local defence industry ecosystem and provides specific cases from around the region and beyond. Specifically, Cyprus and other small states—Israel and Singapore—are briefly examined. The context wherein the discussion takes place consists of nation branding and the critical role of geopolitics. Finally, this work also provides some practical recommendations for policy makers and identifies possible areas for future scholarly inquiry.

Keywords: Cyprus; Defence industry; Europe; Nation Branding/Geopolitics; Small states; Defence-related small and medium-sized enterprises (SMEs).

Introduction

Although the European Union (EU) seeks the ability to protect itself—on its own—from a possible attack from 2030 onwards, SMEs have long faced challenges in obtaining funding (Tidey, 2025). Even though we cannot change the past, we can nevertheless influence the future (Theophanous, 2025, p. 14). It is of the utmost importance, then, to understand how to reposition Cyprus in terms of the defence industry in order to create and enhance the Cyprus nation brand, especially within the EU.[1]

Besides investigating how to form the nation brand, it is imperative to stress the importance of restructuring the whole state apparatus to support Cyprus' defence industry. In other words, bolstering the entire defence infrastructure. Other issues are also significant—such as, the political identity, based on Cyprus' adverse political status and its significant geopolitical position.

Analogous to how Spain, for instance, now needs a strategic approach that aids the nation in restructuring its defence-industrial base during a sustained deterioration in global stability, trust and cooperation between leading world powers (Arteaga, Fiott and Simón, 2025), Cyprus should also consider its next move. This is particularly at a time that requires Europe to act—and quickly (Biscop, 2025). In other words, in an ever more uncertain global arena, Europe must enhance its defence capabilities in order to be ready to face any potential future threats (Fiott, 2024a).

This policy paper, therefore, aims at positioning Cyprus on the European defence industry map. In addition, it presents the highly specialised capabilities of Cyprus' development state-of-the-art, distinct market by efficient SMEs. By highlighting the sound and effective defence industry ecosystem in Cyprus, it illuminates how it could contribute significantly to the advancement of the European defence industry and the strengthening of its autonomy, resilience and supply chain security (Hadjipavlis et al., 2026). That is to say, in a wider context, the present analysis delves into the contribution of small states and SMEs to the broader defence industry ecosystem.

In terms of this article's structure, it commences with an introduction, providing some brief theoretical background—including a literature review (more empirical/theory-backed)—before describing the current contribution of small states and SMEs in the defence industry ecosystem. It then closes by highlighting the need for a robust and sound defence industry in Europe and, specifically, the case of the Republic of Cyprus (RoC). It then presents some practical recommendations and directions for further studies. In this context, the article argues that the development of a strong, state-of-the-art defence industrial sector contributes to the nation building of the RoC. The article now presents a brief theoretical context.

[1] In a broader context, Cyprus overall ranks 38th of 174 countries in total at the time of writing in terms of "The Good Country Index". Specifically, in terms of Global Contribution: (1) Science and Technology (ranked 25th); (2) Culture (ranked 37th); (3) International Peace and Security (ranked 90th); (4) World Order (ranked 2nd); (5) Planet and Climate (ranked 19th); (6) Prosperity and Equality (ranked 151st); and (7) Health and Wellbeing (ranked 64th). See more at: https://index.good-country.org/.

Nation branding comprises a relatively new concept whose field is still evolving (Granitsas, 2024). In general terms, nation branding refers to a country's strategy to carry a specific image of itself outside its borders to attain some advantageous goals. Simon Anholt, an independent policy advisor,[2] is widely considered a pioneer in the field of research dealing with the concept of country "branding". Specifically, in 1996, Anholt coined the term Nation Brand to show that the country's reputation behaves in a similar way to the brand images of businesses and products.[3] As such, a nation's brand can have a substantial impact on its prosperity and welfare, as well as on effective administration (Murphy, 2022).

The Anholt IPSOS Nation Brands Index – which is one of the first indices to have been developed – was created in 2005 and, since 2008, developed in association with global market and social research specialists IPSOS. This Index offers a global appraisal of a country in terms of its performance in relation to six criteria: (1) Exports—i.e., reputation of the products/services of the country; (2) Tourism—i.e., the level of interest in visiting the nation and its natural/artificial attractions; (3) Culture and Heritage—i.e., the value of a country's heritage and interest in its contemporary culture (that is music, art, cinema, literature and sports); (4) Governance—i.e., public opinion about national government competency and fairness, in addition to its commitments to international issues; (5) People—i.e., worldwide reputation with respect to openness, friendliness or tolerance; and (6) Investments and Immigration—i.e., a nation's ability to attract people (or businesses) to live (or settle), work or study, and the quality of life as well as the business environment it provides (Murphy, 2022).

The fundamental theory of nation branding maintains that the better a nation's image is abroad, the better it will be at attracting foreign direct investment, tourists and buyers for its products. Similarly, a nation with a positive image is able to exercise more influence in its foreign relations—essentially, what an international relations expert has called, soft power (Granitsas, 2024).[4]

The results of the study of Rojas-Méndez and Khoshnevis (2023), moreover, reveal that nation branding is influenced by six main factors—business and marketing, political, social and cultural, economic and labour, international and environmental factors.

To move now from perceptions in general to caution about the image of militarisation, perceptions of a nation tend to be entrenched in stereotypes carried by the media and popular culture (Granitsas, 2024). Adding to this, it has repeatedly been said that "perception is reality." (Green and Keegan, 2020, p. 455). In this context, caution is needed lest the image of the militarisation of a nation becomes evident due to heavy defence-related investment. Nonetheless, the prevailing notion within the EU towards the strengthening of Europe's defence capabilities, due to the adverse global geopolitical situation and numerous emerging threats, has been conducive to the Member States' overall efforts to enhance their defence (Fiott, 2023). Cyprus, being a state, whose sovereignty has been violated by Turkey's invasion in 1974 and since then the occupation of its northern territory, has an additional reason to follow the general trend and the declared official policy of the EU for European rearming (Council of the European Union, 2022; Skordeli, 2025). Militarisation and securitisation of the EU and its Member States has indeed been an issue that concerns society and various intellectual fora, however the decision to strengthen the EU's defence comes after decades of inaction, stagnant defence spending and the prevailing doctrine of the "end of history" (Biscop, 2024; Fiott, 2024b; Fukuyama, 1989).

[2] See more, for instance, at: https://uscpublicdiplomacy.org/users/simon_anholt.

[3] For more information on what a nation brand is, see, for instance: https://www.youtube.com/watch?v=agWYoEjCdJA.

[4] The term "soft power" was coined in the late 1980s by Joseph Nye, and can be described as the ability to attract/persuade. In antithesis to soft power—which arises from the attractiveness of a nation's culture, political ideals, as well as policies—hard power, which comprises the ability to coerce, grows out of a nation's military or economic strength. For more info, see: https://www.wcfia.harvard.edu/publications/soft-power-means-success-world-politics, https://www.hks.harvard.edu/faculty/joseph-nye and Nye Jr., J.S. (2005) *Soft Power: The Means to Success in World Politics*, Public Affairs Books, http://www.publicaffairsbooks.com/publicaffairsbooks-cgi-bin/display?book=1586482254.

The interplay between geopolitics and technology

This section is dedicated to geopolitics which, as illustrated in the ensuing text, is also driven by technology.

The interaction between geopolitics and technology is uncertain, with governments around the globe looking for more technological advantage and autonomy by using, for instance, export/investment controls. The future, moreover, belongs to those who shape it today. So, in proactively dealing with the complexity evident at the interconnection of geopolitics and advanced technology, governments need to assume a proactive, flexible as well as informed approach. A strategic mindset that blends geopolitical awareness with technological planning must be implemented to succeed in an ever more complex and dynamic world (Feakin, 2024).

Geopolitics, in simple terms, examines the interaction between geography and politics, whereas politics revolves around power (Ó Tuathail, 1996). Technology and knowledge are crucial factors that strengthen the capabilities of states and other groups in their effort to shape the geographical space (Agnew, 2003). Based on critical geopolitics' theory, this interaction materialises in a discursive way where those who seek to shape the spatial world use the existing resources and capabilities in their hands whilst, at the same time, striving to expand and enhance them, aiming to gain more power to affect their ends (Ó Tuathail and Agnew, 1992). Therefore, the development of a technological defence industrial ecosystem adds to the capabilities of the state whilst it is a useful tool for the various geopolitical practitioners that endeavour to impose their ideas.

Bearing in mind all the above, a credible defence industry also contributes to the nation branding effort as it strengthens and promotes the state's international reputation as a reliable, innovative and technologically capable partner (Anholt, 2007; Mashiah, 2024). Thus, in turn, nation branding acts as another key instrument in discursive practice for the formation of geopolitical imagination and perception.

A brief overview of the defence industry of small states: *The cases of Israel and Singapore*

Israel and Singapore are both involved in military-technological innovation in areas that are crucial to their sovereignty, strategic autonomy and resilience. Both these states have constantly supported high-levels of investment in military research and development, fostering their own defence industries. In addition, both these nations have, to different extents, backed the advance of their local science and technology sectors, frequently integrating commercial high-technology discoveries into their defence capabilities. Further to note, Israel has been more successful than Singapore in advancing military technology, mainly because it faces a more perilous strategic environment. In contrast to Israel, Singapore faces fewer existential threats, thus making its attempts to innovate military-technology more driven by desire than need (Bitzinger, 2021).

Israel

Given that Israel lacks considerable large-scale industrial sectors, it is worth noting how several of its leading arms industries have managed to secure a position among the globe's top one hundred arms-producing firms—with Israel's relatively large share of the world's arms market also being notable (Evron, 2025).[5] The excellence of the Israeli defence industry and its high position among the best worldwide nations has created ample and favourable opportunities for the enhancement of its diplomatic relations with many other countries, leading sometimes to strategic alliances, strengthening of economic and trade ties and eventually close relationships

[5] For more info regarding the market share of the leading exporters of major weapons between 2020 and 2024, by country see, for instance: https://www.statista.com/statistics/267131/market-share-of-the-leadings-exporters-of-conventional-weapons/?srsltid=AfmBOor5t5HoT4wq6n5LeikZcKb__jyW3gpOyHj5ugBDt5Tsr74fvqlD. The SIPRI Top 100 arms-producing and military services companies in the world (2023) offers additional insights: https://www.sipri.org/visualizations/2024/sipri-top-100-arms-producing-and-military-services-companies-world-2023.

and partnerships between business and societies. It is also important to highlight that a profound trust is built between nations when they cooperate closely in the defence domain, not least the development of defence technology and systems (Kinne, 2018).

Israel defence industry's prominence and calibre contribute to the bolstering of its nation branding. Through its potent defence industry, the state's power is enhanced, promoting and manifesting a firm and respected image abroad whilst Israeli society feels secure and proud (Evron, 2025; Grassiani, 2017). Thus, a pioneering defence industry undoubtedly serves the formation of a strong national branding.

Singapore

Singapore comprises an island city-state with the finest air force and navy in Southeast Asia, in addition to being one of the strongest militaries worldwide (Brimelow, 2018). Tan's study (2025) uses the small state of Singapore as a paradigmatic case to illustrate how the role of nation branding acts not merely as image-making, but also as a form of cultural governance. The small state of Singapore has used its advanced niche defence industry to assert itself as a country of high technological capabilities (Raska, 2024). Furthermore, it has connected the defence sector with the advancement and the prosperity of its entire society in the context of its Total Defence doctrine (Matthews and Bintang Timur, 2024). Thus, national resilience, economy, cutting-edge technology, reliability, innovation and competence are some of the key characteristics that promote Singapore's nation branding.

Cyprus

Cyprus' geopolitical position and status definitely impact on the entire spectrum of high-level political decisions and developments on the island. Being situated at the crossroads of three continents, at the centre of the Eastern Mediterranean – a notorious area for its volatility, instability and constant political tensions – close to North Africa and the Near East, regions of critical importance for the EU and a Member State itself, Cyprus holds a pivotal geopolitical position (Tziarras, 2025). This translates into critical defence and security issues. In addition, Turkey's occupation of Cyprus' northern territory since 1974 and its ultra-aggressive disposition pose an additional burden on the latter's security situation.

In relation to this situation, Cyprus is taking steps towards the creation of a small but robust cutting-edge defence ecosystem by taking advantage of the recent developments within the EU that favour the ramping-up of this sector (Hadjipavlis et al., 2026). Cyprus' significant relations with some countries of the region such as Israel and Egypt and their common vested interests are also conducive to this venture. In addition, the long, historical ties with France have also been strengthening, a fact that is also propitious. Cyprus' participation in the Security Action for Europe (SAFE) regulation and the allocation of the significant budget of approximately 1.2 billion euros for strengthening its defence capabilities constitutes a key opportunity to invest also in its own defence industry (Hadjipavlis et al., 2026). Regarding Cyprus, the Defence Minister, Vasilis Palmas, stated that the recent SAFE programme funding would see the Republic's National Guard acquire military hardware that will cover its priorities until 2030 (Cleaver, 2026).

The high-technology defence industrial sector in Cyprus, albeit small and fledgling, could also contribute amongst other sectors to the consolidation and promotion of the Cyprus nation brand. The high level of the Cypriot defence companies' reliability when cooperating with their European counterparts could boost nation branding, both internationally and internally. Hence, Cyprus SMEs could be an inextricable part of the European supply chain as regards the defence industrial sector (Hadjipavlis et al., 2026). In the upcoming years, as cooperation increases and tightens, Cyprus' European partners may improve their perception of the Cyprus' nation brand and appreciate the trustworthiness and professionalism of Cypriot companies. Moreover, Cypriot society may acknowledge this endeavour, which will boost its morale and impact positively on the nation brand.

Conclusion and recommendations

In this section, some of the limitations of this policy paper are presented before concluding and providing a few recommendations for policy makers.

Limitations of study

The lack of available time for conducting the survey as regards the emerging defence sector of Cyprus has been a significant constraint for this policy paper. Furthermore, the novelty of this topic has presented a number of obstacles as regards the scarcity of available sources whilst its sensitivity and confidentiality has posed additional barriers in the specific study.

Practical recommendations for policy makers

Based on the aforementioned analysis, this brief policy paper concludes with some recommendations as regards the development of the Cyprus defence industrial sector in relation to the contemporary geopolitical situation whilst also aiming to enhance the state's nation brand. It is necessary for Cyprus to:

- Restructure the state apparatus and legal and regulatory framework to be more favourable to the development of the defence sector.
- Develop a focused long-term defence industrial strategy accompanied by a robust action plan.
- Increase state investment in the defence domain, in order to indirectly boost the domestic defence industry.
- Increase state investment in research, technology and innovation in defence related fields, including dual-use systems.
- Enhance and intensify the interaction, understanding, cooperation and coordination between the state, academia and the private sector aiming to strengthen the domestic fledgling defence industry.
- Enhance strategic partnerships within the EU and other friendly countries. In this context, the exploitation of geopolitics is crucial for this achievement.
- Intensify efforts towards the attraction of foreign investment into the defence industrial domain.
- Exploit all EU funding schemes and tools for the development of the domestic defence industry.

Directions for future research

Future research could also focus on people, exports, investment and governance (i.e., Anholt's hexagon). The importance of human capital regarding the sound development of a robust defence industrial sector is understudied, so Cyprus could be the testing ground for this endeavour. In addition, further studies could focus on cases such as Sweden, and contribute further to knowledge by delving even deeper into how defence industry-related actors (public/private) employ nation branding to sell as well as legitimise their products/services to insiders (domestic constituents) and outsiders (the global security market) alike (Coetzee, Larsson and Berndtsson, 2024). Furthermore, it could compare it with another country/-ies of interest.

As regards the theoretical part of this study, the relation between the development of the defence industry, geopolitics and nation branding seems to be an understudied field. Thus, there is extensive room for new research in this domain seen through the lens of related inter-disciplinary theories and empirical cases such as that of Cyprus.

References

Agnew, J. (2003) *Geopolitics: Re-visioning world politics* (2nd ed.). Routledge. p. 75.

Anholt, S. (2007) *Competitive identity: The new brand management for nations, cities and regions*. Palgrave Macmillan. pp. 25–34.

Arteaga, F., Fiott, D. and Simón, L. (25 February 2025), 'All in? The revival of the Spanish and European defence industry', *Analyses of the Elcano Royal Institute*, Vol. 2025, pp. 1–4. Available at: https://www.realinstitutoelcano.org/en/analyses/all-in-the-revival-of-the-spanish-and-european-defence-industry/ (Accessed: 23 September 2025).

Attard, D. (28 November 2025) 'Malta's neutrality: A principle or a paradox?', *Times of Malta*. Available at: https://timesofmalta.com/article/malta-neutrality-principle-paradox.1120269 (Accessed: 11 January 2026).

Biscop, S. (2024) 'European defence: No Zeitenwende yet'. *Defence and Peace Economics*, 35(4), pp. 516–520. DOI: 10.1080/10242694.2023.2201739.

Biscop, S. (13 March 2025) 'How the EU can support Europe's NATO', *The Progressive Post* – FEPS-EUROPE. Available at: https://feps-europe.eu/how-the-eu-can-support-europes-nato/ (Accessed: 12 July 2026).

Bitzinger, R.A. (2021) 'Military-technological innovation in small states: The cases of Israel and Singapore', *Journal of Strategic Studies*, Vol. 44, No. 6, pp. 873–900. DOI: 10.1080/01402390.2021.1947252.

Brimelow, B. (8 April 2018) 'How a tiny city-state became a military powerhouse with the best air force and navy in Southeast Asia', Business Insider. Available at: https://www.businessinsider.com/singapore-military-best-air-force-navy-southeast-asia-2018-4 (Accessed: 03 September 2026).

Cleaver, T. (18 June 2026) 'Cyprus takes receipt of first Safe programme payment', *Cyprus Mail*. Available at: https://cyprus-mail.com/2026/06/18/cyprus-takes-receipt-of-first-safe-programme-payment (Accessed: 12 July 2026).

Coetzee, W. S., Larsson, S. and Berndtsson, J. (2024) 'Branding 'progressive' security: The case of Sweden', *Cooperation and Conflict*, Vol. 59, No. 1, pp. 86–106. Available at: https://doi.org/10.1177/00108367231184730.

Council of the European Union (2022) 'A strategic compass for security and defence: For a European Union that protects its citizens, values and interests and contributes to international peace and security'. Council document 7371/22. Available at: https://data.consilium.europa.eu/doc/document/ST-7371-2022-INIT/en/pdf (Accessed: 11 October 2026).

Evron, Y. (2025) 'Israel's defence industry: adaptation and growth in a changing arms market', *Defence Studies*, Vol. 25, No. 2, pp. 301–321. DOI: 10.1080/14702436.2025.2472720.

Feakin, T. (11 November 2024) 'Navigating the New Geopolitics of Tech', *Harvard Business Review*. Available at: https://hbr.org/2024/11/navigating-the-new-geopolitics-of-tech (Accessed: 03 October 2025).

Fiott, D. (2023) 'In every crisis an opportunity? European Union integration in defence and the war on Ukraine', *Journal of European Integration*, 45(3), pp. 447–462.

Fiott, D. (2024a) 'Strategic Domains: Developing the EU's Capabilities and Enablers for a Free and Secure Global Order', *Future Europe Journal*, 5(1), pp. 53–57.

Fiott, D. (2024b) 'The Challenges of Defence Spending in Europe', *Intereconomics*, 59(4), pp. 189–192.

Fukuyama, F. (1989) *The end of history?* The National Interest, (16), pp. 3–18.

Granitsas, A. (18 December 2024) 'The nation brand of Greece', ekathimerini.com. Available at: https://www.ekathimerini.com/opinion/1256483/the-nation-brand-of-greece/ (Accessed: 26 September 2025).

Grassiani, E. (2017) 'Commercialised occupation skills: Israeli security experience as an international brand'. In M. Leese and S. Wittendorp (Eds.), *Security/Mobility: Politics of Movement* (pp. 57–73). Manchester University Press.

Green, M.C. and Keegan, W.J. (2020) *Global Marketing*. Tenth Global Edition. Harlow: Pearson Education Limited.

Hadjipavlis, P., Hadjisavvas, C. and Dimitriou, G. (Eds.) (2026) *Towards a European Defence Union, The Deployment Decade 2026–2035.* Ministry of Defence of the Republic of Cyprus. pp. 57–64.

Kinne, B. J. (2018) 'Defense Cooperation Agreements and the Emergence of a Global Security Network', *International Organization*, 72(4), pp. 799–837.

Mashiah, I. (2024) '"We are a hub for tech, innovation, and entrepreneurship": How places use tech-driven storytelling for nation branding', *Place Branding and Public Diplomacy*, 20, pp. 451–467.

Matthews, R. and Bintang Timur, F. (2024) 'Singapore's 'Total Defence' strategy', *Defence and Peace Economics*, 35(5), pp. 638–658.

Murphy, P.T. (30 September 2022) 'Nation branding: beyond a cosmetic symbol', World Intellectual Property Organization (WIPO) – WIPO Magazine. Available at: https://www.wipo.int/web/wipo-magazine/articles/nation-branding-beyond-a-cosmetic-symbol-42841 (Accessed: 23 September 2025).

Ó Tuathail, G. (1996) *Critical Geopolitics: The Politics of Writing Global Space*. University of Minnesota Press. pp. 1–16.

Ó Tuathail, G. and Agnew, J. (1992) 'Geopolitics and discourse: Practical geopolitical reasoning in American foreign policy', *Political Geography*, 11(2), pp. 190–204.

Raska, M. (2024) 'Reimagining Defense Innovation: Defense AI in Singapore. In: Borchert, H., Schütz, T. and Verbovszky, J. (eds) The Very Long Game. Contributions to Security and Defence Studies. Springer, Cham. Available at: https://doi.org/10.1007/978-3-031-58649-1_25.

Rojas-Méndez, J.I. and Khoshnevis, M. (2023) 'Conceptualizing nation branding: the systematic literature review', *Journal of Product & Brand Management*, Vol. 32, No. 1, pp. 107–123. Available at: https://www.emerald.com/jpbm/article-abstract/32/1/107/246040/Conceptualizing-nation-branding-the-systematic?redirectedFrom=fulltext

Skordeli, M. (2025) 'Unlocking the European Defence Union: Independent and values-based', *European View*, 24(2), pp. 268–276. Available at: https://doi.org/10.1177/17816858251383257

Tan, K.P. (2025) 'From spectacle to reflexivity: rethinking nation branding through public pedagogy and foresight in Singapore, *Place Branding and Public Diplomacy*. Available at: https://doi.org/10.1057/s41254-025-00403-7 (Accessed: 23 September 2025).

Theophanous, A. (2025) *What really happened in Cyprus: A critical assessment since independence and the way forward*. Nicosia: Hippasus Publishing.

Tidey, A. (1 April 2025) 'European defence SMEs must ramp up but the struggle to access money is real. Here's why', Euronews. Available at: https://www.euronews.com/my-europe/2025/03/31/european-defence-smes-must-ramp-up-but-the-struggle-to-access-money-is-real-heres-why (Accessed: 03 October 2025).

Tziarras, Z. (2025) 'Revisiting the Foreign Policy of the Republic of Cyprus and Quasi-Alliances in the Eastern Mediterranean'. *The Cyprus Review*, 36(2), pp. 23–61.

Xuereb, K. (2017) 'Can nation branding be successful tool for the practice of public diplomacy? A Maltese case study', *Place Branding and Public Diplomacy*, Vol. 13, pp. 255–263. Available at: https://doi.org/10.1057/pb.2016.7.

Change Management Needs Surgery, Not Therapy: Why Radical Consistency Beats Gradual Transformation

By

Dr Stephan Meyer

Director, Almademey Limited, Cyprus & Change Management Expert

and Developer of the Sacred Cow Framework

Abstract

This article challenges the prevailing orthodoxy of "gentle transformation" in organisational change management, arguing that incremental, consensus-driven approaches often prolong uncertainty and amplify resistance. Drawing on empirical research with experienced change agents across multiple countries and industries, this paper demonstrates that radical consistency—characterised by clear deadlines, unambiguous communication, and minimal exceptions—yields superior outcomes in digital transformation initiatives. Through theoretical integration with established change management literature and practical case analysis, we establish that effective change management resembles surgical intervention more than therapeutic process, requiring precision, speed and decisiveness. The research identifies "exception proliferation" as a primary factor in change initiative failure and proposes a practitioner-oriented model for accelerated organisational transformation based on the Sacred Cow Framework.

Keywords: Change Management, Digital Transformation, Organisational Change, Radical Implementation, Sacred Cow Framework, Resistance to Change, Implementation Strategy.

Introduction

The change management industry has spent decades selling organisations a comfortable lie: that transformation should be gradual, consultative and, above all, gentle. Business gurus preach the gospel of "taking everyone along on the journey" whilst producing PowerPoint presentations adorned with butterflies and rainbows. "Change is a process, not an event," they philosophise, as organisations slowly deteriorate.

Here's a reality check from thirty years orchestrating digital transformation initiatives: When your house is burning, you don't convene a three-month workshop to discuss optimal evacuation sequences. You exit rapidly before suffering serious harm.

The author's doctoral research, involving interviews with experienced change agents across the globe, revealed a startling consensus: the well-known change management frameworks they had learned were largely ineffective in practice (Meyer, 2019). These practitioners—individuals bearing real scars from genuine transformation battles—consistently reported that radical, time-bound implementation outperformed incremental approaches by virtually every measurable outcome.

This article presents evidence for what I term "radical consistency" in change management: a surgical rather than therapeutic approach that prioritises clarity, speed and unwavering implementation over prolonged consensus-building and exception management. This perspective challenges dominant assumptions in change management literature whilst remaining grounded in both empirical research and established organisational theory.

The Mythology of Gentle Transformation

The Wellness Retreat Fallacy

Contemporary change management discourse saturates organisations with therapeutic language. We speak of "change journeys," "safe spaces" and "psychological readiness." Organisations invest heavily in change management consultants promising to "minimise disruption" and "maintain morale" throughout transformation initiatives (Burnes, 2004).

This approach fundamentally misunderstands the nature of organisational change. As Kotter (1996) observed decades ago, the primary barrier to successful change is not understanding but urgency. A sense of complacency or false urgency undermines transformation efforts more effectively than any technical barrier. Yet we persist in treating change as an educational deficit rather than an execution challenge.

The author's research with international change agents revealed a consistent pattern. When asked about their experiences with incremental versus radical approaches, practitioners reported that organisations approaching transformation as extended consultative processes rather than decisive interventions experienced:

- Significantly longer implementation timelines (often 50-70% beyond initial projections)
- Higher incidence of resistance and political maneuvering
- Greater budget overruns
- Lower staff satisfaction post-implementation

The last finding proves particularly counterintuitive. One might expect that "gentle" approaches would yield happier employees. Instead, prolonged uncertainty proved more damaging to morale than rapid, decisive action—a finding consistent with research on tolerance for ambiguity in organisational contexts (Rafferty & Simons, 2006).

The Science of Procrastinated Pain

From a psychological perspective, the preference for gradual change resembles the flawed logic of removing a plaster millimetre by millimetre. Research in pain perception consistently demonstrates that extended duration amplifies total suffering, even when peak intensity is reduced (Kahneman, Fredrickson, Schreiber, & Redelmeier, 1993). The principle of "duration neglect" suggests people remember pain experiences primarily by their peak intensity and end state, not by total duration—yet organisations design change programmes that maximise duration whilst moderating peaks.

Organisations persist in stretching change initiatives across quarters or even years, believing this reduces employee discomfort. In reality, this approach creates what I term "transformation fatigue"—a state of chronic uncertainty that erodes productivity, increases turnover, and paradoxically amplifies resistance. This phenomenon aligns with research demonstrating that ambiguity about organisational change predicts negative employee outcomes more strongly than the magnitude of change itself (Bordia et al., 2004).

Consider the typical digital transformation initiative: a two-year roadmap with multiple phases, extensive consultation and gradual rollout. At any given moment, some employees are preparing for change, others are experiencing it, and still others are reverting to previous systems. The result is organisational schizophrenia—multiple operating models competing simultaneously, each with its own advocates and processes. This directly contradicts Beer and Nohria's (2000) finding that successful change requires integration of economic and organisational capability perspectives, not fragmented implementation.

The Exception Proliferation Syndrome

Perhaps the most insidious enemy of successful change is what emerges within weeks of any transformation announcement: the exception.

"But IT can't do this." "But the Hamburg office operates differently." "But Herr Müller in accounting has always worked this way and retires next year."

Congratulations—you've just transformed your change initiative into a bureaucratic monster consuming more resources than necessary whilst undermining the transformation's fundamental purpose.

The Hidden Costs of Exception Management

The author's research identified what I call the "Exception Proliferation Syndrome"—a predictable pattern wherein initial transformation designs become progressively compromised by special cases, each seemingly reasonable in isolation but collectively fatal to change effectiveness. This phenomenon reflects what Pfeffer and Sutton (2006) term the "knowing-doing gap"—organisations possess knowledge of what needs changing but fail to implement it consistently.

Interview participants consistently reported that organisations permitting significant exceptions to new processes experienced:

- Substantially higher ongoing support costs (often 2-3x original projections)
- Longer time-to-adoption for remaining staff
- Greater likelihood of process reversion within 12 months
- Reduced credibility of change leadership

The logic is brutal but clear: every exception creates a precedent, every precedent invites comparison, and every comparison fuels resistance. Soon you're managing not one new system but dozens of variants, each with different rules, different training requirements, and different support needs (Repenning & Sterman, 2002).

Moreover, exceptions communicate a devastating message to the organisation: the new approach isn't actually required; it's negotiable. This undermines change leadership more effectively than any opposition campaign could. As Armenakis and Harris (2009) note in their research on readiness for organisational change, inconsistent messaging and application of change initiatives represent a primary factor in implementation failure.

The Loyalty Paradox

A particularly problematic justification for exceptions involves long-serving employees. "Maria has been here for 25 years; we can't ask her to change now." This sentiment, whilst emotionally appealing, reflects muddled thinking about loyalty and respect.

True respect for experienced employees involves:

1. Honest communication about why change is necessary for organisational survival
2. Adequate training and support for new methods
3. Recognition that their expertise remains valuable in the new context
4. Clear expectations applied consistently across all staff regardless of tenure

Excluding long-serving employees from change initiatives actually signals that leadership believes them incapable of adaptation—a far more insulting message than expecting them to learn new systems. This aligns with research on age and learning, which consistently demonstrates that older workers are fully capable of acquiring new skills when provided appropriate support (Warr & Fay, 2001).

Radical Consistency: A Surgical Approach

Core Principles

Effective change management operates more like a light switch than a dimmer: on or off, binary, decisive. You establish Day X, and from that moment, new rules apply. To everyone. Without exception. Without negotiation.

This approach—radical consistency—rests on four foundational principles supported by both the author's empirical research and established change management theory:

1. Temporal Precision

Not "sometime in Q2" but "15 March at 09:00." Precision communicates seriousness. Vagueness communicates ambivalence and enables procrastination.

This principle directly supports Kotter's (1996) first step in organisational transformation: establishing a sense of urgency. Vague timelines undermine urgency by suggesting the change isn't actually critical. Specific deadlines create what Cialdini (2009) terms "commitment and consistency"—once publicly committed to a specific date, organisational pressure builds to honour that commitment.

2. Unambiguous Communication

"We could perhaps possibly consider..." is change management suicide. Effective change communication employs the indicative mood exclusively: "We are implementing X. Training occurs on Y. Go-live is Z."

Patrick Lencioni (2002) identifies ambiguity and lack of commitment as central dysfunctions of teams. In change contexts, this dysfunction proves fatal. Conditional language signals that leadership itself questions the necessity or feasibility of the transformation, inviting subordinates to resist or delay.

Research on communication during organisational change suggests that message clarity and consistency predict employee acceptance more strongly than message frequency or channel selection (Elving, 2005). Leaders must communicate not just often, but unambiguously.

3. Prepared Resilience

Yes, complaints will emerge. Yes, some staff will threaten resignation. No, this is not organisational apocalypse. It's the beginning of transformation.

Resistance is frequently treated in change management practice as an obstacle to be eliminated. Ford, Ford, and D'Amelio (2008), however, challenge this view: resistance is not necessarily pathological, but may signal unresolved concerns, inadequate preparation, or inconsistencies in the change process itself.

Effective change leaders anticipate resistance, prepare responses, and most importantly, refuse to interpret pushback as implementation failure. Research demonstrates that the variable most strongly predicting successful change is not initial acceptance but leadership persistence through the inevitable period of elevated resistance occurring 2-6 weeks post-implementation (Higgs & Rowland, 2005).

4. Uncompromising Execution

If you yield at the first complaint, you might as well abandon the entire initiative. Consistency in execution matters more than perfection in design.

This principle proves particularly challenging for leaders with consultative leadership styles. The desire to accommodate reasonable concerns conflicts with the necessity of maintaining unified implementation. Yet research suggests that leaders who maintained strict adherence to new processes—whilst remaining empathetic to individual challenges—achieved faster adoption and higher ultimate satisfaction than those who granted exceptions (Herold, Fedor, Caldwell, & Liu, 2008).

The Market Leaders' Secret

Examine organisations that haven't merely survived disruption but mastered it. What do they share? They didn't discuss change for decades—they executed it. Rapidly, consistently and without excessive concern for temporary discomfort.

Netflix eventually made a decisive strategic pivot towards streaming, despite continuing to operate its DVD business for years (Keating, 2012). Apple didn't agonise endlessly about whether to introduce the iPhone. Steve Jobs made the decision, assembled the team, and redefined the mobile industry within 30 months of initial conception (Isaacson, 2011).

These examples illustrate what organisational scholars term "strategic agility"—the capacity to make and implement strategic decisions rapidly in response to market conditions (Doz & Kosonen, 2010). Strategic agility requires not just recognition of change necessity but willingness to execute decisively.

The Sacred Cow Framework: Identifying Barriers to Radical Consistency

The author's doctoral research produced the Sacred Cow Framework—a systematic approach to identifying and addressing organisational practices that impede transformation (Meyer, 2019). The framework emerged from asking experienced change agents a simple question: "What practices in organisations consistently undermine change initiatives, despite everyone knowing they're problematic?"

Sacred cows are established practices that persist despite lacking clear rationale or demonstrated value. They survive because "we've always done it this way" and challenging them feels transgressive. Yet they're precisely what radical consistency must address. This concept aligns with institutional theory's observation that organisations maintain practices for legitimacy rather than efficiency (DiMaggio & Powell, 1983).

The Four Categories of Sacred Cows

1. Process Sacred Cows

Workflows that made sense under previous conditions but persist despite changed circumstances. Example: manual approval processes requiring physical signatures in organisations that have otherwise digitalised.

These reflect what organisational scholars term "structural inertia"—the tendency for organisational forms to persist even when environmental conditions have changed (Hannan & Freeman, 1984).

2. Hierarchical Sacred Cows

Organisational structures that reflect historical rather than strategic logic. Example: geographic divisions in global, digital organisations where physical location is increasingly irrelevant.

Research on organisational design suggests that structural misalignment with strategic imperatives creates coordination costs that increase exponentially with organisational size (Burton & Obel, 2004).

3. Cultural Sacred Cows

Unexamined assumptions about "how things work here" that constrain innovation. Example: "Engineers don't do sales" in technology companies, despite evidence that technical founders often make excellent salespeople.

These represent what Schein (2010) terms "basic underlying assumptions"—the deepest level of organisational culture, typically unconscious and taken for granted.

4. Technical Sacred Cows

Legacy systems maintained beyond reasonable lifecycle because change seems too disruptive. Example: bespoke software written by long-departed developers that everyone fears touching.

Technical debt accumulates when organisations defer necessary technology updates, ultimately creating systems so fragile that any change risks catastrophic failure (Kruchten, Nord, & Ozkaya, 2012).

Radical consistency requires systematic sacred cow identification and elimination. Each represents a potential exception-generation source. Address them preemptively or face death by a thousand accommodations.

Practical Implementation: The 90-Day Model

Based on synthesis of successful change initiatives from the author's research, the author proposes a 90-Day Model for radical transformation. This timeframe aligns with research suggesting that 90 days represents an optimal period for initiating organisational change—long enough for adequate preparation but short enough to maintain urgency and focus (Watkins, 2003).

Phase 1: Weeks 1-2 – Diagnosis and Design

- Identify sacred cows blocking transformation using the framework
- Design target state with minimal compromise
- Establish unambiguous success metrics
- Set Day X (90 days forward) and communicate publicly

Phase 2: Weeks 3-6 – Preparation

- Develop comprehensive training programmes
- Create support infrastructure
- Communicate incessantly: what, why, when, how
- Identify and proactively address sources of likely resistance

Phase 3: Weeks 7-10 – Intensive Training

- Mandatory, comprehensive, role-specific training
- Hands-on practice with new systems/processes
- No exceptions for seniority, expertise or claimed time constraints
- Assessment to ensure readiness

Phase 4: Weeks 11-12 – Final Preparation

- Confirm Day X repeatedly through multiple channels
- Final system testing and contingency planning
- Position support teams for intensive post-implementation period
- Intensify communication frequency

Phase 5: Day X – Surgical Cutover

- Previous system OFF, new system ON
- No parallel operation permitted
- All-hands support for initial 72 hours minimum
- Rapid response protocols for emerging issues

Phase 6: Weeks 13+ – Reinforcement

- Daily metrics monitoring with transparent dashboards
- Rapid correction of any deviations
- Public celebration of quick wins
- Absolute consistency on exceptions (Grant fewer than 1% of exception requests)

Critical Success Factors

Leadership Visibility: Change sponsors must be intensely present during the cutover period. Absence communicates lack of conviction more powerfully than any verbal message (Higgs & Rowland, 2005).

Support Intensity: Concentrated support resources during weeks 13-16 prevent regression and build confidence. Research demonstrates that perceived organisational support during change predicts employee acceptance (Rafferty & Griffin, 2006).

Metrics Transparency: Public dashboards showing adoption and performance eliminate ambiguity about progress and create social pressure for compliance.

Exception Management: Establish a high-level review panel for claimed exceptions. Set the bar prohibitively high. Grant fewer than 1% of exception requests. Each exception granted undermines every subsequent implementation effort.

Addressing Predictable Objections

"This Approach Is Too Risky"

Interestingly, analysis suggests incremental approaches carry greater risk. Prolonged change initiatives accumulate more exposure to leadership changes, budget cuts, competing priorities, staff turnover, technology evolution and market disruption (Repenning & Sterman, 2002).

The 90-Day Model concentrates risk into a manageable timeframe with maximum organisational focus. This aligns with research on project management indicating that shorter, focused initiatives have higher success rates than extended programmes (Shenhar & Dvir, 2007).

"Our Culture Couldn't Handle This"

Translation: "We have sacred cows we're unwilling to address." Every organisation believes it's uniquely change-resistant. The author's research reveals this is universally false. Resistance levels correlate weakly with industry or national culture but strongly with leadership clarity and consistency (Meyer, 2019).

Research on organisational culture and change demonstrates that culture can either enable or constrain change, but cultural barriers are not insurmountable—they simply require explicit attention and management (Cameron & Quinn, 2011).

"We Need More Consultation"

Consultation serves two legitimate purposes: identifying genuine technical barriers and ensuring communication clarity. It does not include negotiating whether change occurs or perpetually refining the solution.

Set clear boundaries: "We're consulting on how to implement X, not whether to implement X." This distinction aligns with research on participative decision-making, which suggests that employee input on implementation details increases commitment, whilst consultation on strategic necessity often increases resistance (Sagie & Koslowsky, 2000).

"What About Employee Well-Being?"

Genuine employee well-being emerges from clarity, competence and completion—not prolonged uncertainty disguised as consideration. Research consistently demonstrates that ambiguity about organisational change predicts stress and reduced well-being more strongly than change magnitude (Bordia et al., 2004).

The kindest thing leadership can do is make change swift and certain. Furthermore, the alternative to transformation is usually organisational decline, with its accompanying job losses, budget cuts, and demoralisation. Radical consistency serves employee well-being by securing organisational viability.

Conclusion: Courage Over Comfort

In an era where markets evolve rapidly and business models face constant disruption, hesitation represents a luxury available only to monopolists. Everyone else must master the surgical approach to change.

The author's research with experienced change agents revealed a consistent message: the frameworks they learned emphasised process over outcome, consultation over execution, and comfort over effectiveness. The practitioners who achieved genuine transformation success were those who had learned—often through painful experience—that radical consistency outperforms gentle gradualism.

The question isn't whether your organisation needs transformation. The question is whether you possess sufficient courage to implement it effectively. Will you continue clinging to comfortable illusions about gentle change, or will you embrace the surgical precision required for genuine transformation?

The choice is yours. But choose quickly—your competition isn't waiting for your decision. They're already implementing theirs.

References

Armenakis, A. A., & Harris, S. G. (2009). Reflections: Our journey in organizational change research and practice. *Journal of Change Management*, 9(2), 127–142.

Beer, M., & Nohria, N. (2000). Cracking the code of change. *Harvard Business Review*, 78(3), 133–141.

Bordia, P., Hunt, E., Paulsen, N., Tourish, D., & DiFonzo, N. (2004). Uncertainty during organizational change: Is it all about control? *European Journal of Work and Organizational Psychology*, 13(3), 345–365.

Burnes, B. (2004). *Managing change: A strategic approach to organisational dynamics* (4th ed.). Pearson Education.

Burton, R. M., & Obel, B. (2004). *Strategic organizational diagnosis and design: The dynamics of fit* (3rd ed.). Springer.

Cameron, K. S., & Quinn, R. E. (2011). *Diagnosing and changing organizational culture: Based on the competing values framework* (3rd ed.). Jossey-Bass.

Cialdini, R. B. (2009). *Influence: Science and practice* (5th ed.). Pearson Education.

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147–160.

Doz, Y. L., & Kosonen, M. (2010). Embedding strategic agility: A leadership agenda for accelerating business model renewal. *Long Range Planning*, 43(2-3), 370–382.

Elving, W. J. (2005). The role of communication in organisational change. *Corporate Communications: An International Journal*, 10(2), 129–138.

Ford, J. D., Ford, L. W., & D'Amelio, A. (2008). Resistance to change: The rest of the story. *Academy of Management Review*, 33(2), 362–377.

Hannan, M. T., & Freeman, J. (1984). Structural inertia and organizational change. *American Sociological Review*, 49(2), 149–164.

Herold, D. M., Fedor, D. B., Caldwell, S., & Liu, Y. (2008). The effects of transformational and change leadership on employees' commitment to a change: A multilevel study. *Journal of Applied Psychology*, 93(2), 346–357.

Higgs, M., & Rowland, D. (2005). All changes great and small: Exploring approaches to change and its leadership. *Journal of Change Management*, 5(2), 121–151.

Isaacson, W. (2011). *Steve Jobs*. Simon & Schuster.

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, 4(6), 401–405.

Keating, G. (2012). *Netflixed: The epic battle for America's eyeballs*. Portfolio/Penguin.

Kotter, J. P. (1996). *Leading change*. Harvard Business School Press.

Kruchten, P., Nord, R. L., & Ozkaya, I. (2012). Technical debt: From metaphor to theory and practice. *IEEE Software*, 29(6), 18–21.

Lencioni, P. (2002). *The five dysfunctions of a team: A leadership fable*. Jossey-Bass.

Meyer, S. (2019). *Killing the sacred cow: Radical organisational change as seen through the eyes of change agents*. [Doctoral dissertation, University of Gloucestershire]. University of Gloucestershire Research Repository. https://jo.my/doktorarbeit

Pfeffer, J., & Sutton, R. I. (2006). *Hard facts, dangerous half-truths, and total nonsense: Profiting from evidence-based management*. Harvard Business School Press.

Rafferty, A. E., & Griffin, M. A. (2006). Perceptions of organizational change: A stress and coping perspective. *Journal of Applied Psychology*, 91(5), 1154–1162.

Rafferty, A. E., & Simons, R. H. (2006). An examination of the antecedents of readiness for fine-tuning and corporate transformation changes. *Journal of Business and Psychology*, 20(3), 325–350.

Repenning, N. P., & Sterman, J. D. (2002). Capability traps and self-confirming attribution errors in the dynamics of process improvement. *Administrative Science Quarterly*, 47(2), 265–295.

Sagie, A., & Koslowsky, M. (2000). *Participation and empowerment in organizations: Modeling, effectiveness, and applications*. Sage Publications.

Schein, E. H. (2010). *Organizational culture and leadership* (4th ed.). Jossey-Bass.

Shenhar, A. J., & Dvir, D. (2007). *Reinventing project management: The diamond approach to successful growth and innovation*. Harvard Business School Press.

Warr, P., & Fay, D. (2001). Age and personal initiative at work. *European Journal of Work and Organizational Psychology*, 10(3), 343–353.

Watkins, M. (2003). *The first 90 days: Critical success strategies for new leaders at all levels*. Harvard Business School Press.

The CIS Strategy as a Driver of Self-Esteem and Political Popularity: Posting Selfies to Promote Needs and Values

By

Dr Nicos Antoniades

Professor, Department of Marketing, St. John's University, USA

antonian@stjohns.edu

Abstract

Taking and posting a "selfie" has become a frequent activity for many people globally. However, insufficient research can be found on the impact of selfies in politics. The researcher utilizes the CIS Strategy (Create, Inform, and Support) as the bridge to politicians' social media campaigns. Using regression analysis, the study examines the relationship between specific parameters behind a political selfie (i.e., the improvement of a politician's self-esteem, increased satisfaction and higher popularity) and politicians' attempts to make their values and needs known to the public. The sample included 150 U.S. citizens over 18 years old. According to the results, politicians who use selfies to make their needs and values known improve their self-esteem and increase their popularity. In turn, the use of the CIS strategy has a significant positive impact on politicians' attempts to promote their needs and values to the public.

Keywords: Political Selfies; Self-Esteem; Popularity; Needs and Values; The CIS Strategy.

Introduction

Marketing in Politics

Political players implement marketing strategies to build stronger relationships with their target markets (Lees-Marshment, 2019). Politicians have no other choice than to clearly understand their people's needs and put people at the center of their agenda (Antoniades, 2025; Antoniades and Buoye, 2024; Antoniades, 2023a; 2023b; Antoniades, 2020). This is why political actors are called upon to create political products that the voters-citizens will value positively (Henneberg, 2003).

Political Selfies

In recent decades, a selfie describes an individual's photo captured with a mobile device to post on social media. According to Google, Snapchat and Instagram entered the social media world in 2012, with ninety-three million selfies taken daily (DXOMARK, 2015). Facebook, Snapchat, TikTok, Instagram and other social media platforms are the vehicles of selfies. It is worth noting that one out of two users take a selfie once a month and one out of four take a selfie once a week (Diefenbach and Christoforakos, 2017). Sung et al. (2016) report that selfies facilitate people's understanding of who and what they are, whereas Rutledge (2013) states that taking and posting a selfie is a way to impress. Karadimitriou and Veneti (2016) note that politicians use selfies to promote personal and professional activities. Šimunjak's study (2023) shows that selfies are a useful communication tool that strengthens political performance. Senft and Baym (2015) find that selfies develop a two-way communication between a politician and their target audience and they refer to the prime ministerial campaigns of Justin Trudeau in Canada, Boris Johnson in the United Kingdom, and Narendra Modi in India.

The CIS Strategy

According to Antoniades (2023a, 2023b), the CIS Strategy (i.e., Create, Inform, and Support) shows that the creation of an idea cannot stand alone. Politicians must also give enough information to people (in the form of solutions) and support their ideas. The CIS Strategy has been tested for issues including branding, global warming, cybersecurity, the US economy, and Covid-19.

The Gap

International research has not yet shown (a) whether and how a political selfie promotes a politician's needs and values and (b) whether the CIS Strategy may also impact these needs and values. To fill this gap, the current study aims to examine the relationship between specific parameters behind a political selfie (i.e., the

improvement of a politician's self-esteem, increased satisfaction, and higher popularity) and politicians' attempts to make their values and needs known to the public. The conceptualization of constructs is presented in Appendix A.

Research Hypothesis

To fill this gap, the researcher developed the following hypotheses:

H1: Politicians who take selfies to make their needs and values known improve their self-esteem.

H2: Politicians who take selfies to make their needs and values known increase their satisfaction.

H3: Politicians who take selfies to make their needs and values known strengthen their popularity.

H4: The CIS Strategy (CIS) has a positive impact on a politician's attempts to make their needs and values known.

Research Methodology

One thousand and thirty email invitations were randomly sent by email via a SurveyMonkey link to U.S. citizens over 18 years old. The one hundred and fifty respondents comprised this study's sample (15% response rate). To measure each one of the survey modules, the researcher used closed questions on a 7-point Likert rating scale ranging from 1 (Strongly Disagree) to 7 (Strongly Agree), capturing the strength of the respondents for each given statement, i.e., "Selfies improve self-esteem" (Appendix A). Statistical significance was determined at *P*-value <0.05 (95% confidence level). Pearson correlations (*r*) measured the intensity of each correlation. Shortell (2001) indicates that a correlation exceeding 0.4 is strong, between 0.2 and 0.4 is moderate, and failing shorter than 0.2 is weak.

Presentation of Results

Demographic Information

The sample consisted of U.S. citizens over 18 years old coming from all regions of the U.S. A detailed demographic information list can be found in Appendix B.

Research Hypothesis Testing

According to Table 1:

H1: Politicians who take selfies to make their needs and values known improve their self-esteem. This hypothesis is accepted (P-value <0.05), showing a moderate positive significance (r=0.32) and confirming that the more politicians use selfies to make their needs and values known the higher their self-esteem.

H2: Politicians who take selfies to make their needs and values known increase their satisfaction. This hypothesis is rejected (P-value >0.05).

Table 1 can be found on the next page.

SUMMARY OUTPUT

Regression Statistics	
Multiple R	0.934218
R Square	0.872763
Adjusted R Square	0.864229
Standard Error	2.115198
Observations	150

ANOVA

	df	*SS*	*MS*	*F*	*Significance F*
Regression	3	4511.313	1503.771	336.1087143	2.54E-65
Residual	147	657.687	4.474061		
Total	150	5169			

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*	*Lower 95.0%*	*Upper 95.0%*
Intercept	0	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A
SLF1	0.761151	0.127102	5.988519	1.55505E-08	0.509969	1.012334099	0.509968586	1.012334
SLF2	-0.10967	0.158279	-0.6929	0.489466108	-0.42247	0.203125284	-0.4224685	0.203125
SLF3	0.583109	0.120025	4.858217	3.00205E-06	0.345911	0.820307581	0.345911207	0.820308

Table 1: Selfies - Needs and Values Known

H3: Politicians who take selfies to make their needs and values known strengthen their popularity. This hypothesis is accepted (P-value <0.05), showing a moderate positive significance (*r*=0.34) and confirming that the more politicians use selfies to make their needs and values known the higher their popularity.

According to Table 2:

H4: The CIS Strategy (CIS) has a positive impact on a politician's attempts to make their needs and values known. This hypothesis is accepted (P-value <0.05), showing a moderate positive significance (*r*=0.34) and confirming that the more the use of the CIS Strategy (Create, Inform, and Support), the more effectively politicians promote their needs and values to the public.

Table 2 can be found on the next page.

SUMMARY OUTPUT

Regression Statistics	
Multiple R	0.914497
R Square	0.836305
Adjusted R Square	0.829594
Standard Error	2.383017
Observations	150

ANOVA

	df	*SS*	*MS*	*F*	*Significance F*
Regression	1	4322.863	4322.863	761.2319466	3.25E-60
Residual	149	846.1371	5.678772		
Total	150	5169			

	Coeffi-cients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*	*Lower 95.0%*	*Upper 95.0%*
Intercept	0	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A
CIS	1.132529	0.041048	27.59043	1.98803E-60	1.051418	1.213640094	1.051417764	1.21364

Table 2: The CIS Strategy - Needs and Values Known

Conclusion

The current study shows that politicians who use selfies to make their needs and values known improve their self-esteem and increase their popularity. In turn, the use of the CIS Strategy has a significant positive impact on politicians' attempts to promote their needs and values.

Limitations

The biggest limitation of this study is the relatively low response rate (15%). This may be a result of today's limited political interest, disappointment and political apathy.

Future Research

Further research could use this study's results to examine how Gen Z responds to the above questions. New studies could also test the impact of groupies in politics by using the CIS Strategy as the heart of a political social media campaign.

Disclosure Statement

There is no conflict of interest to declare.

References

Antoniades, N. (2025). *Opposition research: forget it and focus on people.* SAM Advanced Management Journal. Emerald Publishing Limited. 2996-6078. DOI: 10.1108/SAMAMJ-07-2024-0034

Antoniades, N., & Buoye, A. (2024). *The 4Ps Marketing Strategy as a Driver of Dynamic Capabilities: Path to Consumer/Voter Satisfaction.* Review of Marketing Science. https://doi.org/10.1515/roms-2023-0125

Antoniades, N. (2023a). *Packaging Global Warming Products: The CIS Strategy as a Driver of Gen Z's Satisfaction and Loyalty.* Review of Marketing Science, 2023. https://doi.org/10.1515/roms-2022-00884

Antoniades, N. (2023b). *The CIS strategy as a driver of nation branding*. International Journal of Export Marketing, Vol. 5, No. 3-4. Published Online: April 20, 2023, pp. 345–353. https://doi.org/10.1504/IJEXPORTM.2022.1305000

Antoniades, N. (2020). *Political Marketing Communications in Today's Era: Putting People at the Center*. Society. Vol. 57, No. 6. December 2020, Springer U.S.

Diefenbach, S., & Christoforakos, L. (2017). *The Selfie Paradox: Nobody Seems to Like Them, Yet Everyone Has Reasons to Take Them. An Exploration of Psychological Functions of Selfies in Self-Presentation*. Frontiers in Psychology, Volume 8. DOI: 10.3389/fpsyg.2017.00007. ISSN: 1664-1078

DXOMARK (2015). *The evolution of the selfie.* https://www.dxomark.com/the-evolution-of-the-selfie/

Henneberg, S. (2003). *Generic Functions of Political Marketing Management*. University of Bath. School of Management. Working Paper Series, 19.

Karadimitriou, A., Veneti, A. (2016). *Political Selfies: Image Events in the New Media Field*. In: Karatzogianni, A., Nguyen, D., Serafinelli, E. (eds) The Digital Transformation of the Public Sphere. Palgrave Macmillan, London. https://doi.org/10.1057/978-1-137-50456-2-16

Lees-Marshment, J. (2019). *Marketing Scholars and Political Marketing: The Pragmatic and Principled Reasons Why Marketing Academics Should Research the Use of Marketing in the Political Arena*. Customer Needs and Solutions. Vol. 6 (pp. 41–48). Available at: https://doi.org/10.1007/s40547-019-0091-0.

Rutledge, P. B. (2013). "*#Selfies: Narcissism or Self-Exploration?*" in: Psychology Today: Health, Help, Happiness. Available at: https://www.psychologytoday.com/blog/positively-media/201304/selfies-narcissism-or-self-exploration

Senft, T. M., & Baym, N. K. (2015). *What does the selfie say? Investigating a global phenomenon. International Journal of Communication, 9*(Feature), 1588–1606. http://ijoc.org/index.php/ijoc/article/view/4067/1387

Šimunjak, M. (2023). *Chapter 12: Me, myself and I: selfies as vehicles of personalized politics in the social media era.* Political Science and Public Policy 2023. P 153–166. DOI: https://doi.org/10.4337/9781800376939.00020

Sung, Y., Lee, J.-A., Kim, E., & Choi, S. M. (2016). *Why we post selfies: understanding motivations for posting pictures of oneself*. Personal. Individ. Differ. 97, 260–265. DOI: 10.1016/j.paid.2016.03.032

Appendix A

Conceptualization of Constructs

Selfie/Groupie	Adapted From
Selfies improve self-esteem (SLF1). Systematic selfie viewing leads to increased satisfaction (SLF2). Selfie viewing leads to increased popularity (SLF3).	Wang et al. (2017)
The CIS Strategy	
My favorite politicians create proposals that satisfy my needs (CIS-C). My favorite politicians provide useful information regarding their proposals (CIS-I). My favorite politicians support their proposals efficiently (CIS-S).	Antoniades (2023a; 2023b)
Making Needs and Values Known	
By taking and posting a selfie/groupie, I am making my needs and values known (SVK).	

Note: This research used a structured approach with closed statements based on a 7-point Likert rating scale (1932) ranging from 1 (Strongly Disagree) to 7 (Strongly Agree).

Appendix B: Demographic Profile

Group	Participants (%)
Gender	
Men	70
Women	30
X	0
Age	
18-29	22
30-44	36
45-60	32
>60	10
Household Income	
$0-$24,999	14
$25,000-$49,999	18
$50,000-$74,999	20
$75,000-$99,999	14
$100,000-$149,999	20
$149,000-$199,999	4
>$200,000	6
Prefer not to answer	4
Education	
College/University	72
High School	16
Less than High School	10
Prefer not to say	2
Party Affiliation	
Republican Party	50
Democratic Party	40
Other Affiliations	10
Region	
East North Central	22
East South Central	4
Middle Atlantic	12
Mountain	6
New England	4
Pacific	24
South Atlantic	16
West North Central	4
West South Central	8
Device Type	
iOS Phone / Tablet	60
Android Phone /	40
Desktop	0

Safety Campaigns in the Maritime Industry

By

Captain Valentin Mavrinac

Director of Shipping Studies | Senior Lecturer

CIM-Cyprus Business School, Limassol Campus

Senior Research Fellow

Cyprus Centre for Business Research – CIM

Abstract

Safety campaigns are widely used in the maritime industry to improve safety awareness, reinforce compliance, and support behavioural change; however, their effectiveness is rarely evaluated as a standalone intervention. This study reviews peer-reviewed maritime safety literature to assess the role, mechanisms, and potential benefits of campaign-based interventions within shipping operations. The findings indicate strong conceptual convergence across empirical, review-based, and qualitative studies, demonstrating that improvements in safety culture, communication, competence, and behavioural awareness are consistently associated with enhanced safety performance.

The literature suggests that safety campaigns are most effective when embedded within broader safety management systems and aligned with leadership practices, organisational culture, and operational controls. Campaigns contribute to proactive safety management by reinforcing learning, supporting hazard reporting, and sustaining attention on priority risks. Increasingly, predictive approaches based on leading indicators and data analytics offer new opportunities to target campaign interventions more effectively.

From an economic perspective, improved safety performance is linked to reduced incident costs, increased operational efficiency, enhanced organisational resilience, and stronger competitive positioning. However, the review identifies significant methodological limitations, including the absence of campaign-specific evaluation, limited use of longitudinal research designs, insufficient integration of objective performance indicators, lack of comparative analysis across campaign types, and minimal economic valuation of safety interventions.

The paper concludes that while safety campaigns represent an important component of maritime safety practice, their effectiveness depends on integration within a systematic, risk-based, and leadership-driven framework. Future research should focus on developing rigorous empirical methods, including longitudinal and mixed-method approaches, to evaluate campaign outcomes and quantify their contribution to both safety performance and organisational value.

Keywords: Maritime safety; Safety campaigns; Safety culture; Human factors; Proactive and predictive safety; Leading indicators; Operational efficiency.

Introduction

Safety campaigns in the maritime industry are generally designed to improve workforce awareness, reinforce compliance, shape safe behaviour, and support the development of a stronger organisational safety culture. Although relatively few peer-reviewed studies evaluate safety campaigns as a standalone intervention, a substantial body of maritime research examines closely related measures such as safety culture development, training, safety communication, and broader safety management initiatives. Collectively, this literature provides a credible basis for assessing the likely effects and benefits of safety campaigns in maritime operations, further strengthened by significant industry expertise and extensive empirical knowledge contributed by multiple maritime professionals and practitioners.

Various tools used to strengthen safety culture are particularly important in the maritime industry, where safety performance depends not only on technical systems and regulatory compliance, but also on organisational culture, crew competence, communication quality, and the management of human factors. In this respect, campaign-based interventions are best understood as part of a wider system designed to strengthen the behavioural and organisational conditions that support safe operations. In this context, safety campaigns can be regarded as both reactive and proactive tools, depending on their design, and also depending on impetus being either by lagging indicators (e.g. incident analysis and lessons learned) or by leading indicators aimed at preventing incidents and accidents before they materialise.

Safety Culture as a Core Outcome of Safety Campaigns

A central theme in the contemporary maritime safety literature is the importance of safety culture as both an outcome and a mechanism of safety intervention. Campaigns typically aim to reinforce shared values, strengthen communication, signal management commitment, and encourage collective ownership of safe behaviour. These are all core dimensions of safety culture identified in maritime research.

Xi et al. (2025) provide strong recent empirical support for this relationship. Using structural equation modelling with survey data from 229 ship deck officers and captains, the authors found that safety culture has a statistically significant positive effect on safety behaviour. Safety material culture, safety system culture, safety management culture, and safety value fit all contributed to stronger safety compliance and more proactive safety-related behaviour onboard. This is directly relevant to safety campaigns because such interventions commonly seek to influence exactly these dimensions of organisational life.

At a broader level, Xu et al. (2023), in a systematic review of 186 studies on maritime transportation safety management, similarly identify safety culture, communication, and safety management processes as central foundations of resilient maritime operations. Their findings suggest that safety improvement in shipping depends not only on formal rules and procedures, but also on the behavioural and cultural conditions that support implementation in everyday work. This strengthens the case that safety campaigns are most likely to be effective when embedded within the wider safety management system rather than delivered as isolated awareness events.

Behavioural Change and Safety Compliance

A principal objective of safety campaigns is to influence day-to-day behaviour. The maritime literature strongly supports the proposition that safety-focused organisational interventions can improve both compliance and proactive conduct. Xi et al. (2025) found that stronger safety culture was associated not only with routine compliance but also with prosocial safety behaviour, including actions taken in support of collective safety. This is significant because maritime incidents are often preceded by unsafe acts, unsafe behaviour, and near misses rather than by technical failure alone.

The role of learning, awareness, and reinforcement is also supported by Chowdhury et al. (2024), whose review of human factors in maritime engine-room safety highlights the importance of collaboration, training, and competence development in reducing operational error. Their analysis concludes that human-factor-related accident risk can be mitigated through improved safety knowledge, stronger competency development, and closer attention to preventive measures. Safety campaigns that combine communication, learning reinforcement, and behavioural expectations can therefore be understood as practical tools for supporting these forms of risk reduction.

Overall, the literature suggests that campaigns are most effective when they do more than provide information. To influence behaviour sustainably, they need to reinforce an organisational environment in which safe choices are expected, supported, and recognised.

Human Factors and Error Reduction

The importance of behaviour change is intensified by the central role of human factors in maritime accidents. Across the maritime safety literature, human error and human-factor-related failures are repeatedly identified as major contributors to incidents at sea. Chowdhury et al. (2024) review the root causes and categories of human-factor risk and conclude that these factors make a significant contribution to accidents in maritime operations, especially in technically demanding environments such as engine rooms.

This finding is directly relevant to safety campaigns because many such initiatives are designed to address human-factor vulnerabilities such as complacency, fatigue, communication breakdown, procedural drift, and weak risk perception. Xi et al. (2025) support this interpretation by demonstrating that stronger safety culture improves officers' safety behaviour, thereby addressing the conditions under which errors are more likely to occur.

The implication is that safety campaigns can generate safety and business value when they address the cognitive, social, and organisational precursors of accidents rather than merely repeating generic compliance messages.

Proactive Safety Management and Organisational Learning

Contemporary maritime safety management has increasingly shifted from reactive, incident-driven approaches towards more proactive and resilience-oriented models. Xu et al. (2023) note that recent maritime safety research increasingly addresses resilience, scenario representation, emergency management, and other advanced methods aimed at strengthening prevention.

In this context, safety campaigns can be interpreted as one mechanism of proactive safety management. They help maintain attention on priority risks, encourage hazard reporting, keep lessons learned visible, and provide repeated reinforcement of expected standards. The findings of Xi et al. (2025) suggest that such reinforcement is particularly valuable when it strengthens safety value fit and management culture, since these dimensions were associated with broad improvement in safety behaviour.

In practice, this means that campaigns are likely to be most effective when integrated within a structured organisational learning cycle, including incident review, communication of lessons learned, supervisory follow-up, and continuous feedback. However, the maritime industry must take a further step by advancing towards predictive safety management as a natural evolution of these practices. The utilisation of big data represents a critical enabler in this transition, allowing organisations to systematically analyse trends, identify emerging patterns, and anticipate potential incidents. By leveraging predictive insights to more precisely target the focus of safety campaigns, organisations can move beyond reinforcement and actively eliminate or mitigate hazards before incidents occur.

Economic Benefits of Safety Campaigns in the Maritime Industry

Reduction in Direct and Indirect Incident Costs

The most immediate economic benefit of better safety performance is the reduction of incident-related losses. In maritime operations, such losses can include impacts affecting life, the environment, the vessel, property, cargo, legal costs, and the cost of incident investigation. There are also substantial indirect costs, including delay, off-hire time, reputational damage, schedule disruption, loss of customer confidence, and management time diverted away from productive operations.

Raharjo et al. (2026), in their systematic literature review on safety performance and competitiveness in the shipping industry, argue that stronger safety performance reduces avoidable incidents and supports lower operational costs, improved efficiency, and stronger long-term sustainability. Because safety campaigns seek to prevent unsafe behaviour and reinforce risk awareness, they can contribute to these savings by reducing the frequency or severity of incidents.

Operational Efficiency and Reliability

Safety performance and operational performance are closely linked. Safer operations are generally more stable, predictable, and efficient because they experience fewer interruptions, fewer emergency responses, and less unplanned repair work. Raharjo et al. (2026) explicitly associate improved safety performance with efficiency gains and competitive advantage in shipping. Chowdhury et al. (2024) likewise show that strengthening competence and addressing human factors can improve operational conditions by reducing error-prone situations.

In practical terms, fewer mistakes and incidents mean fewer delays, less equipment damage, fewer disruptions to voyage schedules, and better utilisation of assets and crew time. These are economically important gains in an industry where margins are often narrow and contractual performance is commercially significant.

Reputation, Competitiveness and Market Position

A strong safety record can also support improved market standing. Raharjo et al. (2026) emphasise that safety performance can become a source of competitive advantage through improved reputation, efficiency, and sustainability. In the shipping market, charterers, cargo owners, insurers, and regulators routinely use safety performance as an indicator of operational reliability and management quality. A company that demonstrates visible commitment to safety may therefore enjoy stronger commercial credibility and lower barriers in safety-sensitive markets.

Although the available literature does not provide a direct campaign-to-charter-rate model, it strongly supports the wider proposition that improved safety performance enhances market competitiveness.

Workforce Stability and Human Capital Value

The human and economic dimensions of maritime safety are closely intertwined. Shan (2022), examining occupational safety and health challenges for seafarers during the COVID-19 period, highlights the importance of communication, wellbeing, and safe working conditions in maintaining seafarer health and operational resilience. Where safety communication is poor or risk controls are weak, organisations may face higher fatigue, lower morale, poorer retention, and reduced productivity.

Campaigns that focus on wellbeing, fatigue awareness, communication, or safe work practices may therefore create value not only by reducing injuries but also by improving trust, motivation, and workforce stability.

Compliance and Avoidance of Business Interruption

Campaigns also reinforce awareness of procedural requirements and support adherence to company rules and regulatory expectations. Xu et al. (2023) show that maritime safety management is shaped by formal

systems, risk analysis, and emergency management processes, while Xi et al. (2025) show that safety culture is positively related to compliance behaviour. Better compliance reduces the likelihood of detentions, failed inspections, costly corrective action, or business interruptions linked to safety deficiencies.

Strategic Return on Investment

Viewed strategically, safety campaigns should not be understood solely as short-term administrative costs. The literature suggests that they are relatively low-cost interventions compared with the financial consequences of accidents, operational disruption, reputational damage, or workforce instability. Taken together, the findings of Raharjo et al. (2026), Xi et al. (2025), and the wider safety management literature support the argument that well-designed campaigns can offer a positive long-term return on investment when embedded within broader safety management practice.

Industry platforms such as HiLo posit that predictive safety management leads to cost savings. This view is supported by academic literature. While authors commonly define outcomes in terms of reduced downtime, improved reliability, and enhanced operational efficiency, these effects collectively constitute measurable economic benefits through avoided failures and optimised resource utilisation (Ilavarasi et al., 2023; Kalafatelis et al., 2025; Simion et al., 2024). Predictive maintenance and data-driven safety management have been shown to reduce failures, optimise resource allocation, and improve profitability, thereby confirming that early intervention is more cost-effective than post-incident response.

Safety Communication in Crisis and High-Risk Contexts

The value of campaign-style communication is especially visible during periods of heightened operational stress. Shan (2022) demonstrates that maritime workers experienced serious occupational safety and health pressures during the COVID-19 crisis, including conditions that affected wellbeing and safe performance. The study highlights the importance of communication and institutional support in protecting crews in a rapidly changing risk environment.

Although this study focuses on a crisis context rather than routine campaign delivery, it nevertheless provides useful evidence that safety-related communication can become a critical control during disruption. This strengthens the broader case for maintaining regular, visible, and workforce-centred safety campaigns in both normal operations and emergency conditions.

Critical Appraisal of the Literature

Strengths of the Existing Evidence Base

One clear strength of the literature is strong conceptual convergence across different study types. Empirical, review-based, and qualitative studies all point in a similar direction: improvements in safety culture, communication, competence, and behavioural awareness are associated with stronger safety performance in maritime operations (Arslan et al., 2016; Hasanspahić et al., 2021; Karakasnaki et al., 2018; Kim and Gausdal, 2020; Kirwan et al., 2021; Mbong and Bygvraa, 2021; Thai and Grewal, 2006; Veluplay et al., 2020). Xi et al. (2025) provide direct quantitative evidence that safety culture positively influences both compliance and proactive safety behaviour, while Xu et al. (2023) identify safety culture and safety management processes as recurring foundations of maritime safety improvement.

A second strength is the practical relevance of the literature to the operational realities of the maritime industry. The reviewed studies focus on real maritime problems such as crew behaviour, human error, occupational safety pressures, and organisational resilience (Mok, D'Agostini and Ryoo, 2023; Rothblum, n.d.; Sha, 2020). Chowdhury et al. (2024) examine operational risks linked to human factors in engine-room environments, while Shan (2022) explores the occupational safety and health challenges faced by seafarers during a period of extreme disruption.

A third strength is that the evidence base spans multiple levels of analysis, from individual behaviour to system-wide safety management and strategic outcomes such as competitiveness and efficiency. This is useful because maritime safety campaigns are rarely effective in isolation; their success depends on alignment between behaviour, supervisory practice, organisational culture, and senior management priorities (Curcuruto et al., 2015; Banda and Goerlandt, 2018). Jalonen and Salmi (2009) highlight the importance of safety indicators, suggesting that alignment must be measurable and continuously monitored, while Pun, Yam and Lewis (2003) indicate that formal Safety Management System (SMS) adoption does not guarantee effectiveness, reinforcing that alignment must be functional, not only structural. Karanikas (2014) suggests safety systems require planning, feedback, and control mechanisms, not just alignment between levels.

Limitations of the Current Literature

The most significant limitation is that very few studies explicitly evaluate safety campaigns as a standalone variable. In most cases, campaigns are subsumed within broader categories such as safety culture initiatives, training programmes, or communication practices. The literature supports the underlying logic of campaigns indirectly, and rarely isolates whether a specific campaign caused an observed improvement (Arslan et al., 2016; Banda and Goerlandt, 2018; Curcuruto et al., 2015; Hasanspahić et al., 2021; Karanikas, 2014; Maurino, 2017; Mok, D'Agostini and Ryoo, 2023; Thai and Grewal, 2006). This gap highlights the limited attention given to assessing the independent effectiveness of safety campaigns and points to a clear need for future research employing rigorous empirical approaches capable of isolating campaign effects and establishing their causal contribution to safety performance.

A further limitation is the dominance of cross-sectional and review-based methodologies. Xi et al. (2025) rely on cross-sectional self-reported survey data, which is valuable for identifying relationships but limited for establishing causality over time. Chowdhury et al. (2024) and Raharjo et al. (2026) are review studies, and while such studies are valuable for synthesis, they depend on the quality and scope of included evidence.

Another limitation concerns context specificity and generalisability. Xi et al. (2025) draw on respondents from a specific company setting in Shanghai, while Shan (2022) focuses on seafarers' experiences during the COVID-19 crisis. These studies are highly relevant, but their findings may reflect organisational, cultural, or crisis-specific conditions that do not map equally across all vessel types, flag states, or maritime sectors.

There is also a limitation related to outcome measurement. Much of the literature uses intermediate constructs such as attitudes, perceptions, safety culture scores, or reported behavioural intentions rather than hard operational outcomes such as lost-time injury frequency, near-miss rates, detention data, insurance claims, or quantified cost savings.

Methodological Gaps

The most notable methodological gap is the lack of longitudinal intervention studies. To assess campaign effectiveness with greater confidence, future studies should measure outcomes before, during, and after implementation. Without such designs, it remains difficult to distinguish between the effects of campaigns and those of wider organisational or regulatory changes occurring at the same time (Banda and Goerlandt, 2018; Karakasnaki et al., 2018; Mbong and Bygvraa, 2021; Thai and Grewal, 2006; Xi et al., 2025; Xu et al., 2023).

A second methodological gap is the limited use of objective performance indicators. Much of the existing literature relies on behavioural, perceptual, or qualitative measures, with relatively few studies incorporating operational data such as incident frequency, lost time injury frequency (LTIF), total recordable case frequency (TRCF), near-miss reporting rates, utilisation of STOP work authority, unsafe act, unsafe conditions, port state control findings and detentions, contained spills and actual spills, audit non-conformities, third party observations, machinery failures, off-hire time, and claims history (Banda and Goerlandt, 2018; Jalonen and Salmi, 2009; Karanikas, 2014; Mbong and Bygvraa, 2021; Veluplay et al., 2020; Xi et al., 2025).

A third gap concerns comparative research design. Existing studies tend to examine specific dimensions of safety improvement—such as leadership, safety culture, or behavioural interventions—in isolation, with little effort to compare the relative effectiveness of different campaign types or intervention strategies. As a result, there is limited empirical evidence to guide the selection of appropriate approaches, such as visual communication campaigns, leadership-led engagement initiatives, behaviour-based observation programmes, near-miss reporting drives, or wellbeing and fatigue awareness efforts. This gap is particularly important given that the effectiveness of such interventions is likely to vary across operational contexts, yet the current literature provides limited insight into which approaches are most effective under specific conditions (Banda and Goerlandt, 2018; Hasanspahić et al., 2021; Karakasnaki et al., 2018; Kim and Gausdal, 2020; Xu et al., 2023).

Finally, there is a major gap in economic evaluation methodology. While the literature strongly suggests that better safety performance supports efficiency, competitiveness, and reduced disruption, few studies calculate campaign return on investment in monetary terms. As a result, the economic value of safety initiatives—particularly campaign-based interventions—remains largely unmeasured. Future research should therefore focus on developing robust cost–benefit and return-on-investment models that integrate safety performance data with financial outcomes, enabling a clearer understanding of the economic impact of safety interventions (Kalafatelis et al., 2025; Mok, D'Agostini and Ryoo, 2023; Raharjo et al., 2026; Simion et al., 2024; Thai and Grewal, 2006).

Implications for Maritime QSHE (Quality, Safety, Health, Environment and Energy) Practice

For maritime QSHE professionals, the literature indicates that safety campaigns should not be implemented as isolated communication exercises. They should be integrated into the organisation's wider management system, learning processes, and leadership practices. Campaigns are likely to be most effective when they reinforce existing risk controls, procedures, and supervisory expectations rather than simply repeating slogans.

The literature also implies that campaigns should target specific behavioural and operational risks. Human-factor-related findings suggest that campaigns are likely to create more value when they address concrete precursors of incidents such as poor communication, fatigue, weak situational awareness, procedural deviation, empowerment to STOP work when unsafe, or under-reporting of near misses.

A further implication is the importance of internal measurement and evaluation. Because the literature currently lacks strong campaign-specific causal evidence, QSHE practitioners should build their own evidence

base by monitoring campaign-related indicators such as reporting patterns, improved learning and understanding, improved effectiveness, observation quality, office personnel involvement, behaviour change, permit compliance, and safety perception over time.

Leadership visibility is also critical. Xi et al. (2025) show that safety management culture and safety value fit are influential in shaping behaviour, implying that campaign messages require reinforcement by masters, senior officers, superintendents and senior shore management. This perspective is further supported by the principles of human and organisational performance, which emphasise that leadership behaviour shapes the operational system within which individuals act (Conklin, 2019). Without leadership follow-through, campaigns risk becoming symbolic rather than operationally meaningful, with limited impact on actual safety performance.

Targeted Improvements for QSHE Practice

Move from Generic Campaigns to Risk-Based Campaign Design

QSHE teams should replace generic awareness campaigns with risk-based interventions informed not only by retrospective analysis but also by proactive and predictive safety insights. While incident investigations, near-miss trends, audit findings, and high-risk task analysis provide valuable reactive inputs, increasing use should also be made of leading indicators, data analytics, and predictive risk modelling to anticipate emerging hazards before they materialise. Campaign themes should therefore be dynamically selected and prioritised based on both current and forecasted operational risks, such as navigation, mooring, enclosed space entry, hot work, work at height, emergency preparedness, line-of-fire exposure, or fatigue management.

Embed Campaigns Within the Safety Management System

Campaigns should be integrated into permit-to-work systems, toolbox talk routines, risk assessments, audit plans, and vessel visit agendas. Each campaign should be linked to a specific procedure, checklist, verification step, and accountabilities so that messages are translated into operational controls.

Introduce Measurable Leading KPIs for Campaign Effectiveness

Organisations should establish measurable leading indicators. These indicators should be tracked before, during, and after campaigns in order to establish whether meaningful change has occurred.

Strengthen Leadership and Supervisory Involvement

Masters, senior officers, superintendents, and senior shore management should take visible roles in campaign launch, discussion, reinforcement, and follow-up. Leadership participation is essential if campaign messages are to be seen as credible, relevant, and aligned with operational priorities.

Shift from Awareness to Behavioural Reinforcement

QSHE practice should focus not only on sharing information but also on reinforcing specific behaviours. Practical measures include peer observation, short behavioural checklists, real-case discussions, STOP work authority exercises, and explicit recognition of safe conduct.

Incorporate Human Factors and Non-Technical Skills

Campaigns should explicitly address human factors such as fatigue, communication breakdown, workload, situation awareness, complacency, risk normalisation, procedural violations (work as done versus work as imagined), supervision, training and competence gap, human-technology interaction, psychological safety, reporting culture, teamwork, organisational pressure, and decision-making under pressure. These topics should be linked to vessel operations through scenarios, toolbox discussions, and supervisor-led reflection on recent events.

Develop Continuous Campaign Cycles Rather Than One-Off Events

Campaigns should be delivered through a structured cycle consisting of launch, reinforcement, verification, feedback, and improvement. A campaign duration of several weeks, supported by repeated engagement rather than a single communication release, is more likely to produce sustained behavioural change.

Enhance Crew Engagement and Bottom-Up Feedback

Crew members should be encouraged to contribute examples, lessons learned, and suggestions during campaigns. Short surveys, suggestion tools, and vessel-led safety case sharing can increase ownership and improve the practical fit of campaign content.

Align Campaigns with Commercial and Operational Objectives

Safety campaigns should be linked to broader operational goals such as reliability, voyage performance, cargo integrity, and customer satisfaction. Presenting campaign value in both safety and business terms can improve management support and resource commitment.

Build a Structured Campaign Governance Framework

Maritime organisations should establish a standard campaign governance model covering risk selection, campaign design, implementation responsibilities, monitoring arrangements, outcome evaluation, and post-campaign learning. This will help transform campaigns from ad hoc safety communications into evidence-informed risk control interventions.

Conclusion

The literature indicates that safety campaigns in maritime settings are best understood as part of a broader system of behavioural and organisational intervention. Although direct evidence on campaigns as a discrete variable remains limited, current peer-reviewed research strongly supports the mechanisms through which campaigns are expected to work: strengthening safety culture, improving compliance and proactive behaviour, addressing human-factor-related risk, enhancing communication, and supporting proactive safety management.

The evidence further suggests that these interventions deliver benefits beyond safety alone. By reducing accident-related costs, improving operational efficiency, strengthening workforce resilience, and supporting competitive advantage, safety campaigns may generate substantial economic value for maritime organisations. For this reason, they should be treated not merely as compliance tools, but as strategic components of maritime safety and organisational performance. The most effective QSHE approach is therefore likely to be one in which campaigns are proactive and predictive, risk-based, measurable, behaviour-focused, leadership-supported, and fully integrated into the management system.

Recommendations for Future Research

The review identifies several important directions for future research on maritime safety campaigns. Greater use of longitudinal and intervention-based research designs would strengthen the evidence base by allowing safety outcomes to be measured before, during, and after campaign implementation. Such approaches would provide a firmer basis for establishing causality and distinguishing campaign effects from concurrent organisational, operational, or regulatory influences.

A further priority is the development of more robust measurement frameworks. Much of the existing literature relies on behavioural, perceptual, or cultural indicators, with comparatively limited use of objective operational data. Future studies should integrate behavioural measures with performance indicators such as incident frequency, near-miss reporting, permit-to-work compliance, safety observations, audit findings, port state control outcomes, equipment reliability, and claims history in order to provide a more comprehensive assessment of campaign effectiveness.

The literature would also benefit from comparative evaluation of different campaign formats and intervention strategies. Current evidence provides little guidance on the relative effectiveness of approaches such as leadership-led engagement initiatives, behaviour-based safety programmes, near-miss reporting campaigns, visual awareness campaigns, and wellbeing-focused interventions. Comparative studies conducted across different vessel types, fleet segments, and operational environments would help identify the conditions under which particular approaches are most successful.

The growing availability of leading indicators, predictive analytics, and big-data methodologies presents another promising avenue for research. Future investigations should examine how proactive and predictive safety information can be used to identify emerging risks, trigger targeted interventions, and optimise campaign design before incidents occur. Such work would contribute to the ongoing transition from reactive to predictive safety management within the maritime sector.

There is also a clear need for more rigorous economic evaluation. Although existing studies consistently suggest that improved safety performance contributes to operational efficiency, organisational resilience, and competitive advantage, little research has quantified the financial returns associated with campaign-based interventions. Future research should therefore develop cost-benefit and return-on-investment methodologies capable of measuring the economic value generated through reductions in incidents, operational disruption, equipment failures, claims, and other safety-related losses.

Collectively, these research priorities would strengthen the empirical foundation of the field, support more evidence-based decision-making, and improve understanding of how safety campaigns contribute to both safety performance and organisational effectiveness within maritime operations.

References

Arslan, V., Kurt, R.E., Turan, O. and De Wolff, L. (2016) 'Safety culture assessment and implementation framework to enhance maritime safety', *Transportation Research Procedia*, 14, pp. 3895–3904. Available at: https://doi.org/10.1016/j.trpro.2016.05.477

Banda, O.A.V. and Goerlandt, F. (2018) 'A STAMP-based approach for designing maritime safety management systems', *Safety Science*, 109, pp. 109–129. Available at: https://doi.org/10.1016/j.ssci.2018.05.003

Chowdhury, M.N., Shafi, S., Arzaman, A.F.M., Teoh, B.A., Kadhim, K.A., Salamun, H., Kadir, F.K.A., Said, S. et al. (2024) 'Navigating human factors in maritime safety: A review of risks and improvements in engine rooms of ocean-going vessels', *International Journal of Safety and Security Engineering,* 14(1), pp. 1–14. Available at: https://doi.org/10.18280/ijsse.140101

Conklin, T. (2019) '*The 5 Principles of Human Performance*'. Los Alamos, NM: Pre-Accident Investigation Media.

Curcuruto, M., Mariani, M.G., Conchie, S.M. and Violante, F.S. (2015) 'The role of prosocial and proactive safety behaviors in predicting safety performance', *Safety Science*, 80, pp. 317–323. Available at: https://doi.org/10.1016/j.ssci.2015.07.032

Hasanspahić, N., Frančić, V., Vujičić, S. and Mandušić, M. (2021) 'Safety leadership as a means for safe and sustainable shipping', *Sustainability*, 13(14), 7841. Available at: https://doi.org/10.3390/su13147841

Ilavarasi, S., Kaliappan, V.K., S., R. and Dugki, M. (2023) 'AI-driven predictive maintenance in maritime vessels'. *International Journal of Robotics and Mechatronics*, 10(3). Available at: https://ojs.unsysdigital.com/index.php/ijrm/article/view/1286

Jalonen, R. and Salmi, K. (2009) 'Safety performance indicators for maritime safety management', Helsinki University of Technology. Available at: http://appmech.tkk.fi/fi/julkaisut/TKK-AM-9.pdf

Kalafatelis, A.S., Nomikos, N., Giannopoulos, A., Alexandridis, G., Karditsa, A. and Trakadas, P. (2025) 'Towards predictive maintenance in the maritime industry: A component-based overview'. *Journal of Marine Science and Engineering*, 13(3), 425. Available at: https://doi.org/10.3390/jmse13030425

Karakasnaki, M., Vlachopoulos, P., Pantouvakis, A. and Bouranta, N. (2018) 'ISM Code implementation: An investigation of safety issues in the shipping industry', *WMU Journal of Maritime Affairs*, 17, pp. 461–474. Available at: https://doi.org/10.1007/s13437-018-0153-4

Karanikas, N. (2014) 'A comprehensive contemporary safety management systems framework including planning and monitoring guidance', *MERC Global International Journal of Management*, 2(3), pp. 69–96. Available at: http://ijm.mercglobal.org

Kim, T.E. and Gausdal, A.H. (2020) 'Leaders' influence tactics for safety: An exploratory study in the maritime context', *Safety*, 6(1), 8. Available at: https://doi.org/10.3390/safety6010008

Kirwan, B., Bettignies-Thiebaux, B., Cocchioni, M., Baumler, R. and Carrera Arce, M. (2021) 'Towards a safety learning culture for the shipping industry', World Maritime University. Available at: https://commons.wmu.se/lib_documents/

Mbong, T.S.-S. and Bygvraa, D.A. (2021) 'Analysis of the implementation of the International Safety Management Code using motivation theory: The seafarer's views', *International Maritime Health*, 72(3), pp. 172–178. Available at: https://doi.org/10.5603/IMH.2021.0033

Mok, I.-S., D'Agostini, E. and Ryoo, D.-K. (2023) 'A validation study of ISM Code's continual effectiveness through a multilateral comparative analysis of maritime accidents in Korean waters', *The Journal of Navigation*, pp. 1–14. Available at: https://doi.org/10.1017/S0373463322000571

Pun, K.F., Yam, R.C.M. and Lewis, W.G. (2003) 'Safety management system registration in the shipping industry', *International Journal of Quality & Reliability Management*, 20(6), pp. 704–721. Available at: https://doi.org/10.1108/02656710310482140

Raharjo, B.A., Utami, H.N., Afrianty, T.W. and Ruhana, I. (2026) 'The influence of safety performance on competitive advantage in the shipping industry: A systematic literature review study', F1000Research, 14, Article 490. Available at: https://doi.org/10.12688/f1000research.163658.4

Rothblum, A.M. (n.d.) 'Human error and marine safety', U.S. Coast Guard Research and Development Center. Available at: https://www.google.com/search?q=Human+error+and+marine+safety+Rothblum

Sha, M. (2020) 'A study of the human factors in maritime safety', MSc Dissertation – World Maritime University. Available at: https://commons.wmu.se/msem_dissertations/335/

Shan, D. (2022) 'Occupational safety and health challenges for maritime key workers in the global COVID-19 pandemic', *International Labour Review*, 161(2), pp. 267–287. Available at: https://doi.org/10.1111/ilr.12220

Simion, D., Postolache, F., Fleacă, B. and Fleacă, E. (2024) 'AI-driven predictive maintenance in modern maritime transport—enhancing operational efficiency and reliability', *Applied Sciences*, 14(20), 9439. Available at: https://doi.org/10.3390/app14209439

Thai, V.V. and Grewal, D. (2006) 'The maritime safety management system (MSMS): A survey of the international shipping community', *Maritime Economics & Logistics*, 8, pp. 287–310. Available at: https://doi.org/10.1057/palgrave.mel.9100161

Veluplay, G.K., Tsimplis, M.N., Shenoi, R.A. and Abdul Rahman, N.S.F. (2020) 'Insights of safety practices in the shipping industry: A qualitative assessment', *International Journal of Scientific & Technology Research*, 9(8), pp. 441–443. Available at: https://www.ijstr.org

Xi, Y., Wang, Z., Hu, S., Han, B. and Yin, J. (2025) 'The effect of safety culture on the safety behavior of ship deck officers: Empirical evidence from shipping industry', *Frontiers in Marine Science*, 12, Article 1599455. Available at: https://doi.org/10.3389/fmars.2025.1599455

Xu, M., Ma, X., Zhao, Y. and Qiao, W. (2023) 'A systematic literature review of maritime transportation safety management', *Journal of Marine Science and Engineering*, 11(12), 2311. Available at: https://doi.org/10.3390/jmse11122311

* Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author used Microsoft Copilot to assist with topic development, structuring of the literature review, language refinement, and improvement of readability. After using this tool, the author reviewed, verified, edited, and approved all content as needed. The AI tool was not used to generate original research data, was not listed as an author, and did not replace the author's critical analysis, interpretation, or professional judgement. The author takes full responsibility for the content of the publication.

AUTHORS

Christiana Charalambidou

Constantinos Constantinou

Demetris Tillyris

Elena Constantinou

Ioannis Sidiropoulos

Marios Georgiou

Maryam Kazemi Manesh

Nicos Antoniades

Panayiotis Hadjipavlis

Stephan Meyer

Valentin Mavrinac

Yotam Werzansky-Orland

Zano Tyrannis

Christiana Charalambidou is an economist specialising in labour market and education research, applied econometrics and large-scale data analysis. Her research has examined issues including over-education and skills mismatch, the relationship between education and labour market outcomes, worker behaviour and on-the-job search, and the changing distribution of jobs in response to technological change. She has extensive experience using large datasets and advanced econometric methods to generate evidence that can inform public policy, support strategic business decision-making and contribute to a better understanding of labour market dynamics.

Her research interests also extend to behavioural economics, where she has worked with experimental and laboratory data to study the behaviour and decision-making of economic agents. More recently, her work has included contributions to the Cyprus Tax Reform project and impact evaluation studies assessing the effectiveness of labour market policies. She has also worked on multidisciplinary projects involving economic analysis, feasibility studies and the assessment of the socioeconomic impact of environmental, sustainability and circular economy policies.

Dr Charalambidou has extensive teaching and supervision experience across British and Cypriot universities and higher education institutions, and has supervised more than 40 postgraduate dissertations. She teaches economics, applied econometrics using Python, econometrics and statistics, data analysis and research methods at both undergraduate and postgraduate level.

She is also an HRDA Certified Vocational Trainer and has delivered corporate training in business and managerial economics, business analytics, executive decision-making and negotiation skills.

Dr Charalambidou holds a PhD in Economics from the University of Sheffield, an MSc in International Employment Relations and Human Resource Management from the London School of Economics and Political Science (LSE), and a BA (Honours) in Economics with French from the University of Nottingham. She has also studied Economics at the University of Social Sciences in Toulouse, France, and is a member of the European Association of Labour Economists (EALE).

At CIM-Cyprus Business School, she holds the position of Associate Professor and contributes as a lecturer, MBA dissertation supervisor and student mentor.

Contact info:
Email: christiana@cim.ac.cy.

Constantinos Constantinou is the Director of Research and Head of Undergraduate Studies, Nicosia at CIM-Cyprus Business School. He is also the Editor-in-Chief of *The Market: International Journal of Business* and has previously served as a post-doctoral Research Fellow and Associate Director of the Cyprus Centre for Business Research – CIM. He has also worked as a Research Fellow at the Centre for European and International Affairs at the University of Nicosia (UNIC), and in the Media and Communication Department of the European Defence Agency in Brussels. Further, he is a Senior Research Fellow at the CCBR – CIM and a reviewer for journals such as the *Journal of Marketing Communications*, the *Journal of the Academy of Marketing Science*, the *Review of Marketing Science*, and the *SAM Advanced Management Journal*.

At CIM-Cyprus Business School, he holds the position of Associate Professor of Marketing, acting as a lecturer for CIM/UWL undergraduate and postgraduate modules. In addition, he acts as a supervisor for CIM/UWL Masters students' dissertations, student mentor and has served as the Project Manager for the recent launch

of CIM's online MBA. Constantinos holds a PhD in Business Administration from the European University Cyprus (EUC), an MA in International Relations and European Studies from UNIC and an MBA/BA in European Studies from EUC. His research focuses on political marketing and communication.

ORCID ID: 0000-0001-6595-0355.

Contact info:
Email: constantinos@cim.ac.cy
LinkedIn: https://www.linkedin.com/in/findcon/.

Demetris Tillyris is a leadership scholar and award-winning educator, specialising in leadership, ethics, management and governance, strategic communication, and international affairs. He studied at the London School of Economics and Political Science (LSE), where he completed an MSc in Management, Organisations and Governance, and earned his PhD in Politics and International Relations from the University of Leeds, where his doctoral thesis received a Research Excellence Award. He also holds a First-Class BA (Hons) in Politics and Economics from the University of Manchester, a Postgraduate Certificate in Academic Practice from Canterbury Christ Church University and was elected a Fellow of the Higher Education Academy (FHEA) in recognition of his commitment to teaching excellence.

Dr Tillyris's research focuses on leadership, ethics, management, governance and communication, and has been published in leading Q1 international journals, including the *British Journal of Politics and International Relations*, *Journal of Ethics*, *Ethical Theory and Moral Practice*, *Res Publica*, *Critical Review of International Social and Political Philosophy*, *Theoria*, *Philosophia*, and *Contemporary Politics*.

With more than 15 years of experience in higher education, Dr Tillyris has taught at the University of Manchester, the University of Leeds and Canterbury Christ Church University, where he served as Reader (Associate Professor). During his academic career, he has designed and directed numerous undergraduate and postgraduate programmes across business, leadership, politics, international relations and sustainability. He is the recipient of multiple Awards for Teaching Excellence and Outstanding Lectures, together with repeated nominations for Outstanding Leadership Awards.

Beyond academia, Dr Tillyris is an active contributor to public debate through keynote addresses, policy commentary, and media interviews for outlets including BBC Radio Kent, KMTV, Kent Online, RTL, RFI, and Mediapart. He was also the founder and director of the Making Politics Matter public lecture series, which has brought together leading politicians, policymakers, business leaders and internationally recognised scholars to discuss contemporary societal challenges.

Dr Tillyris currently teaches at the University of Cyprus, and CIM-Cyprus Business School, where he serves as Director for Postgraduate Studies. At CIM, Dr Tillyris brings a distinctive combination of internationally recognised scholarship, award-winning teaching, leadership expertise, and real-world strategic experience, helping students bridge cutting-edge research with the practical challenges of contemporary business, governance and public life.

Contact info:
Email: demetris.t@cim.ac.cy.

Elena Constantinou is a senior partner and lead trainer at Evolving Minds, a firm specialising in corporate training, team coaching and HR consulting. She has designed and delivered numerous training programs on leadership and skills development to large corporations in Cyprus and has also trained groups virtually across the globe. For the last four years, she has been teaching Entrepreneurship, Business Communication and Negotiations at the University of Cyprus. She will be defending her PhD shortly at the School of Education at Frederick University. Her research interests include how critical thinking and reflection improve metacognition and self-regulated learning in university students and adults. As part of her PhD research, she has developed an online application which can be used for reflective journaling in university courses or coaching engagements.

Ms Constantinou studied Economics and Finance at Princeton University and earned her MBA from the London Business School with a concentration in Entrepreneurship and Strategy. She worked for more than two decades in finance and business development for groups of companies such as JPMorgan Chase in New York and London, Lanitis Group of Companies, Bank of Cyprus, Hellenic Bank and B2Kapital (Cyprus). She also holds coaching certifications from Coachmasters Academy in Singapore and Team Coaching International in the USA. She is a certified trainer through the Cyprus Human Resource Development Authority.

In addition, Ms Constantinou is the President of the Nicosia Toastmasters Club and has been distinguished in two global public speaking competitions. Her contributions to the community include chairing the Cyprus American Alumni Network, supporting young people pursue their dreams and aspirations as a university interview coach and through mentoring teams participating in the Junior Achievement Company Program and the Cyprus Entrepreneurship Competition at the University of Cyprus.

Contact info:
LinkedIn: https://www.linkedin.com/in/elena-constantinou-279967180
Email: elena.constantinou@evolving-minds.net.

Ioannis Sidiropoulos is LL.M (Dual), MA, Lawyer (Athens Bar, Cyprus Bar) and Legal Analysis Executive at Chrysses Demetriades & Co LLC. He is also a Senior Fellow at Strategy International Think Tank, in addition to a Law Lecturer at CIM-Cyprus Business School. Further, he is a PhD Candidate in AI Governance, Sovereignty and National Security at the European University Cyprus.

Contact info:
Email: Ioannis.Sidiropoulos-lecturer@cim.ac.cy.

Marios Georgiou has been the MBA Director and Professor at CIM-Cyprus Business School since 2015. He studied Chemical Engineering at the University of Manchester Institute of Science and Technology (UMIST) and obtained the degree of BSc. Having received a scholarship from UMIST, he has undertaken research in the field of distillation modelling and obtained the degrees of MSc and PhD.

Dr Georgiou worked for BP Oil International for 32 years in various management and leadership positions, with the most recent position as a sales and marketing manager for East Med and North Africa.

During his employment with BP, he completed an MBA degree with Warwick University and later on, a Doctorate in Business Administration (DBA), with Gloucestershire University, with his submitted thesis on Foreign Market Entry.

Contact info:
Email: marios.g@cim.ac.cy.

Maryam Kazemi Manesh is a business professional known for her expertise in managing large-scale projects, driving operational efficiency, and delivering high-value outcomes in the technology and innovation sectors. Throughout her career in international tech companies and consultancies, she has led the development and delivery of enterprise software solutions for diverse markets. As Senior Technical Project Manager at GILAWA Ltd, she oversees strategic initiatives, leads cross-functional teams to achieve organisational goals. She holds a PhD in Financial Economy from the University of Mannheim, Germany and a BSc in Applied Mathematics. With over 20 years of academic and professional background, Dr. Kazemi Manesh continues to inspire and empower others through her leadership and dedication to excellence.

Contact info:
Email: mkm@gilawa.com
LinkedIn: https://www.linkedin.com/in/dr-maryam-kazemi-manesh/.

Nicos Antoniades, Ph.D., is a Professor of Marketing and Political Marketing in New York (St. John's University, CUNY/College of Staten Island, DeVry/Keller Graduate School of Management). He acted as an adviser to politicians and organizations. He published several academic papers and developed the AMPCOPE Test, an online tool measuring political performance. He is an award-winning scholar (The CIS Model, University of Greenwich, 2019), the Teaching Excellence Award (2021 & 2022), the Award of Distinction (Communicator Awards 2024, Column-Editorial), and the "Values & Inclusion" award (SJU, 2025). Nicos is a reviewer for top marketing journals (i.e., American Marketing Association, Journal of Marketing Communications, and Academy of Marketing Science). He is an editor at Cyprus' newspapers and an accredited reporter at the UN. In 2019, St. John's University invited Nicos to design and teach the Political Marketing course, a course that he has been teaching since then.

Contact info:
Email: antonian@stjohns.edu

Panayiotis Hadjipavlis is a defence and aviation specialist with more than 30 years of experience in defence strategic analysis, military aviation, airspace geopolitics and flight safety issues. He is characterised by a comprehensive and interdisciplinary approach to analysing political and security related issues, a virtue attributed to his diverse academic background. Panayiotis graduated as a jet pilot from the Hellenic Air Force Academy, where he also earned a BSc in Aviation Science. He also holds a BA in European Studies (European University Cyprus), an MBA (University of Sunderland, UK), an MSc in Aviation Safety Management (City University, UK), a Master II in Religions, Cultures and Politics in Modern and Contemporary Europe (École Pratique des Hautes Études, France) and a Professional Masters in Superior Military Studies from the French War College. In July 2021, he earned a doctoral degree in Geopolitics of the Airspace from Newcastle University, UK.

Panayiotis is a visiting lecturer at the Centre for Security and Intelligence Studies (BUCSIS) of the University of Buckingham, UK, a Research Associate of the Centre for European and International Affairs of the University of Nicosia, a Research Fellow of the Defence and Security Research Institute (DSRI) of the University of Nicosia and a Research Fellow at CERIDES Research Centre of the European University Cyprus. At the same time, he has been an associate editor of several academic publications and has published several academic articles in various journals and books. Panayiotis' research interests include the geopolitics of the airspace, techno-geopolitics, the geopolitics of the Eastern Mediterranean, Near East and Sahel regions, issues of state sovereignty and European defence cooperation.

As regards his professional career, Dr Panayiotis Hadjipavlis has been a Defence Director at the Cyprus Ministry of Defence since October 2023, Head of the General Procurement and Defence Capabilities' Development. Under his remit, among other matters, are the European Defence Agency affairs and all the Defence Capabilities' Developments within the EU, the development of the local Defence Industry and research into and innovation of the Cyprus Armed Forces. As a retired Colonel, Dr Hadjipavlis has served in several critical posts at the Cyprus Ministry of Defence. In 2019 he was posted as a strategic analyst for the Sahel region in Africa at the EU Military Staff/ European External Action Service, while in August 2021 he assumed the post of the Head of the Africa Section of the EUMS Strategic Analysis Section.

Contact info:
Email: phadjipavlis@windowslive.com.

Stephan Meyer is a renowned expert in business transformation and change management with over 30 years of experience supporting visionary leaders and organisations in implementing meaningful change. As Director of Almademey Ltd. and internationally sought-after keynote speaker, he specialises in making complex transformations simple and sustainable.

With a doctorate in Business Administration focusing on change management and a diploma in Business Psychology, Dr. Meyer brings both academic rigour and practical expertise to organisational transformation. As Senior Lecturer at CIM-Cyprus Business School, he shares his knowledge with the next generation of business leaders. His innovative "Sacred Cow Framework" helps companies identify and overcome outdated practices that hinder progress and profitability.

Throughout his career, Dr. Meyer has accompanied major corporations such as Volkswagen, Deutsche Telekom, and PWC through successful transformation initiatives. He is co-author of books including *People, Culture & Communication* and *Digital Project Practice for New Work and Industry 4.0*.

As a strategic advisor and psychological counselor to CEOs and executive boards, Dr. Meyer works primarily with leaders of medium-sized companies, supporting them in modernising their organisations and digital transformation. His unique methodology combines rigorous analysis with practical implementation, transforming abstract visions into actionable realities.

Contact info:
Email: office@stephanmeyer.com.
Web: https://stephanmeyer.com
LinkedIn: https://jo.my/smli

Valentin Mavrinac is Director of Shipping Studies, Senior Lecturer, and Senior Research Fellow at CIM-Cyprus Business School and the CIM Research Centre, Limassol, Cyprus. He lectures in Shipping Management and coordinates undergraduate and postgraduate shipping programmes. He holds a Diploma in Maritime Traffic Engineer – Nautical Course from the Faculty of Maritime Studies in Rijeka, and an MSc in Human Resource Management and Organisational Behaviour from the Cyprus International Institute of Management.

Captain Mavrinac began his maritime career at sea and progressed to command merchant vessels before transitioning to shore-based ship management. He currently serves as QSHE Manager and Designated Person Ashore (DPA) within a major international ship management organisation. His professional and research interests include maritime safety management, safety culture, human factors, human and organisational

performance (HOP), behavioural safety, organisational resilience, regulatory compliance, and proactive and predictive approaches to safety management.

His work focuses on bridging academic research and industry practice through the application of evidence-based safety management principles to maritime operations.

Contact info:
Email: valentin@cim.ac.cy | valentin.mavrinac@gmail.com.

Yotam Werzansky-Orland, Esq., is an IP strategist, scholar and advisor working at the intersection of IP strategy, innovation management and AI governance. He is Head of Innovation and Intellectual Property at the Cyprus Centre for Business Research, teaches at several academic institutions, and advises executive boards.

He is the only individual from Cyprus ever to have been named to the exclusive IAM Strategy 300, a global ranking to leading IP strategists. His work focuses on the legal, strategic and governance implications of emerging technologies, particularly for authorship, creativity and human–machine collaboration.

He holds LL.B. and LL.M. degrees and is pursuing doctoral research at Adam Mickiewicz University.

Contact info:
Email: yotam@cim.ac.cy.

Zano Tyrannis is a Performance Partner and organisational effectiveness professional and consultant with extensive experience across the healthcare and financial sectors. He has led strategic initiatives in private healthcare, particularly within chiropractic and wellness services, and has played a key role in developing international investment and residency programs through Eleon Capital Fund Managers in Cyprus.

With a background that bridges clinical practice and business strategy, Zano brings a unique perspective to leadership, marketing, and organisational development. His consulting work as a Performance Partner focuses on helping businesses enhance operational efficiency, expand market reach, and integrate innovative technologies such as AI-driven marketing and health analytics.

Dr Tyrannis is passionate about combining academic insight with real-world application, contributing to the professional development of future leaders in business and healthcare management. The aforesaid academic insight owes much to his rich educational background, including an M.Tech. Chiro (University of Johannesburg), M.Sc. Med. (University of the Witwatersrand), and an M.B.A. (University of Pretoria - GIBS). His current interests include digital transformation, wellness entrepreneurship, and sustainable investment models.

Contact info:
Email: zanott@gmail.com.

The Market

International Journal of Business

www.cim.ac.cy/market

market@cim.ac.cy

A publication of the Cyprus Centre for Business Research