Ruffo D, Capparelli F, Panfilo D, D'Acquisto G, Brusadelli A and Bella G, 'Synthetic Data under NIS2: From Regulatory Silence to an Evolutionary Standard of Design-Based Compliance' (2027) 1 EJPLT.
DOI: https://doi.org/10.57230/EJPLT271DRFCDPGDABGB
Corresponding author: A. Brusadelli. E-mail address: alessandro.brusadelli@aindo.com
P.zza del Mercato no. 115, Brescia (IT)

EJPLT
ISSN 2704-8012



Synthetic Data under NIS2: From Regulatory Silence to an Evolutionary Standard of Design-Based Compliance

Dati sintetici e NIS2: dal silenzio regolatorio a uno standard evolutivo di conformità by design

DAVIDE RUFFO
Teaching Fellow in European Union Law
Università Europea di Roma

FRANCESCO CAPPARELLI
President of the Italian Institute for Privacy

DANIELE PANFILO
PhD in Machine Learning and Artificial Intelligence
Università degli Studi di Trieste

GIUSEPPE D'ACQUISTO
Adjunct Professor of Artificial Intelligence and Law
Luiss Guido Carli University

ALESSANDRO BRUSADELLI
PhD Candidate in Artificial Intelligence in Medicine and Innovation in Clinical and Methodological Research
Università degli Studi di Brescia

GIAMPAOLO BELLA
Full Professor of Computer Science
Università degli Studi di Catania

© The author(s) 2027, published by Suor Orsola Benincasa Università Editrice.
This contribution is licensed under a Creative Common Attribution 4.0 International Licence CC-BY-NC-ND, all the details on the license are available at:
https://creativecommons.org/licenses/by-nc-nd/4.0/

Abstract
Directive (EU) 2022/2555 (NIS2) does not mention synthetic data and imposes no general duty to replace personal data with synthetic datasets. This article argues that such regulatory silence does not entail legal irrelevance: in high-risk environments, synthetic data may progressively become a safeguard to be considered under NIS2's evolutionary standard of appropriateness, proportionality and state of the art. The argument is developed through legal interpretation and tested against operational cybersecurity frameworks, with particular attention to the Belgian CyberFundamentals Framework and the Italian ACN baseline specifications. The article ultimately proposes a model of design-based compliance, where cybersecurity resilience, data minimisation, accountability and fundamental rights protection are embedded in system architecture.

La Direttiva (UE) 2022/2555 (NIS2) non menziona i dati sintetici e non impone alcun obbligo generale di sostituire i dati personali con dataset sintetici. Il contributo sostiene che tale silenzio regolatorio non equivale a irrilevanza giuridica: nei contesti ad alto rischio, i dati sintetici possono progressivamente diventare una garanzia da considerare alla luce dello standard evolutivo di adeguatezza, proporzionalità e stato dell'arte previsto dalla NIS2. L'argomento è sviluppato attraverso l'interpretazione giuridica e verificato rispetto a framework operativi di cybersecurity, con particolare attenzione al Belgian CyberFundamentals Framework e alle specifiche di base dell'ACN italiana. L'articolo propone infine un modello di design-based compliance, nel quale resilienza cybersecurity, minimizzazione dei dati, accountability e tutela dei diritti fondamentali sono integrate nell'architettura del sistema.

Keywords: synthetic data; NIS2 Directive; GDPR; cybersecurity law; privacy by design; re-identification risk; EHDS

Summary: 1. Introduction and State of the Literature. – 1.1 Existing Scholarship and the Gap This Article Addresses. – 2. Regulatory Silence and Open-Textured Obligations under NIS2. – 3. Synthetic Data in the Assessment of Incident Impact and Structural Risk Reduction. – 3.1 The Heterogeneity of Synthetic Outputs and the Limits of De-identification Claims. – 4. From Regulatory Silence to an Evolutionary Standard. – 5. Consequences for Articles 6 and 9 GDPR. – 5.1 Article 6(1)(c) GDPR: A Context-Sensitive Framework. – 5.2 Article 9 GDPR, the European Health Data Space, and the Evolving Legislative Horizon. – 6. Design-Based Compliance and the Future of Privacy and Cybersecurity Law. – 6.1 Limits and Conditions of the Model. – 7. Conclusion: Synthetic Data as an Emerging Standard of Adequacy.

1. Introduction and State of the Literature.

A recurring feature of contemporary digital regulation is that technological practices often acquire legal relevance before they are expressly named by legislation. Synthetic data provide a paradigmatic illustration of this phenomenon. They are increasingly deployed in health data governance, data sharing, analytics, testing environments and artificial intelligence

development,[1] yet Directive (EU) 2022/2555 ('NIS2') does not identify them as a distinct compliance instrument.[2] At first sight, this omission may appear decisive. If the Directive does not mention synthetic data, one might conclude that no specific legal consequences should follow from their deployment or omission. Such a conclusion would be too hasty and it would also overlook the fact that NIS2 compliance is increasingly operationalized through technical-control frameworks which translate open-ended legal obligations into auditable organizational and security measures.

NIS2 does not operate through an exhaustive catalogue of named technological measures. Rather, it relies on open-textured standards – most notably the obligation imposed on essential and important entities to adopt 'appropriate and proportionate technical, operational and organizational measures' to manage risks posed to the security of network and information systems, taking account of 'the state of the art and, where applicable, relevant European and international standards.'[3] The question is therefore not whether synthetic data were specifically contemplated by the legislature, but whether, within a dynamic and context-sensitive reading of adequacy, they may increasingly fall within the range of measures that regulated entities are expected to consider.

1.1 Existing Scholarship and the Gap This Article Addresses.

Three bodies of literature converge on the subject-matter of this article without having been brought into systematic dialogue. The first concerns open texture in legal norms and the dynamic interpretation of technology-neutral standards. H.L.A. Hart established that legal rules are necessarily open-textured because of the inherent vagueness of the language in which they are expressed: their application to borderline cases requires normative judgment that evolves with social and technological circumstances.[4] The philosophical roots of this concept trace back to Friedrich Waismann's notion of *die Porosität der Begriffe* (the porosity of concepts), which Hart transposed into legal philosophy to account for the inevitable indeterminacy of general terms when applied to unforeseen facts.[5] This insight has been applied extensively to EU digital regulation, where framework directives deliberately avoid rigid technological taxonomies in order to preserve adaptive capacity.[6]

---

[1] Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space and amending Directive 2011/24/EU and Regulation (EU) 2024/2847 [2025] OJ L, 2025/327, 5.3.2025 (EHDS Regulation); VB Vallevik and others, 'Processing of synthetic data in AI development for healthcare and the definition of personal data in EU law' (2026) 34 International Journal of Law and Information Technology eaag002.

[2] Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union [2022] OJ L333/80 (NIS2 Directive).

[3] ibid art 21(1), first and second subparas.

[4] HLA Hart, *The Concept of Law* (Clarendon Press 1961) ch VII; F Schauer, 'On the Open Texture of Law' (2013) 87 Grazer Philosophische Studien 197.

[5] F Waismann, 'Verifiability' (1945) 19 Proceedings of the Aristotelian Society, Supplementary Volumes 119.

[6] FMJ Teichmann, 'Platform governance under NIS2 and the Cyber Resilience Act: cybersecurity by design as social practice' [2026] Information, Communication & Society 1; C Djeffal, 'Law by

The second body of literature concerns the legal status of synthetic data under the GDPR. Scholarly contributions have examined the privacy-utility tradeoff that synthetic generation is intended to resolve, noting that the technique covers a heterogeneous spectrum of generation methods – including generative adversarial networks (GANs), variational autoencoders, rule-based simulation, and differential privacy mechanisms – each carrying different re-identification risk profiles.[7] The question of whether a synthetic output qualifies as non-personal data is no longer purely speculative: a growing body of technical literature has developed multidimensional evaluation frameworks that assess fidelity, utility and privacy protection through distinct metric families, enabling a structured, empirically grounded determination of when a synthetic dataset ceases to be reasonably linkable to the original population of individuals.[8]

A comprehensive scoping review of privacy and utility metrics in medical synthetic data published in *npj Digital Medicine* has confirmed that, while no universal consensus on evaluation methodology yet exists, the technical community is converging on standardized metrics for privacy risk assessment – notably membership inference resistance, attribute disclosure risk, and singularity scores – that together support a legally meaningful, though necessarily contextual and revisable, assessment of whether the resulting dataset is no longer reasonably likely to permit identification within the meaning of GDPR Recital 26.[9] Yet convergence is not closure. El Emam and co-authors have demonstrated that, where generative models are overfit to underlying training data, synthetic outputs can enable 'meaningful identity disclosure': an adversary who matches a synthetic record to a real individual may still learn something new about that person.[10] More recent work employing density-based membership inference attack frameworks – including the DOMIAS model and the GenLRA approach – has confirmed that such risks are most acute for underrepresented or atypical individuals in the original population, precisely those for whom re-identification would carry the greatest potential harm.[11] The implication for legal analysis is significant: a positive legal

---

design obligations: The future of regulating digital technologies in Europe?' (2025) 59 Computer Law & Security Review 106232.

[7] M Cairo, 'Synthetic Data and GDPR Compliance: How Artificial Intelligence Might Resolve the Privacy-Utility Tradeoff' (2023) 28 Journal of Technology Law & Policy 71.

[8] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data [2016] OJ L119/1 (General Data Protection Regulation) (GDPR), recital 26.

[9] B Kaabachi and others, 'A scoping review of privacy and utility metrics in medical synthetic data' (2025) 8 npj Digital Medicine 60; M Hernandez and others, 'Comprehensive evaluation framework for synthetic tabular data in health: fidelity, utility and privacy analysis of generative models with and without privacy guarantees' (2025) 7 Frontiers in Digital Health 1576290.

[10] K El Emam, L Mosquera and J Bass, 'Evaluating Identity Disclosure Risk in Fully Synthetic Health Data: Model Development and Validation' (2020) 22 Journal of Medical Internet Research e23139; K El Emam, S Rodgers and B Malin, 'Anonymising and sharing individual patient data' (2015) 350 BMJ h1139.

[11] B Van Breugel and others, 'Membership Inference Attacks against Synthetic Data through Overfitting Detection' (arXiv, 24 February 2023) https://arxiv.org/abs/2302.12580 accessed 3 July 2026; J Ward, C Wang and G Cheng, 'Privacy Auditing Synthetic Data Release through Local

classification of synthetic outputs as substantially anonymous cannot be asserted on the basis of the generation method alone; it requires documented, empirically grounded technical evaluation against the specific dataset, threat model, and deployment environment at issue.

Italian regulatory scholarship has argued that synthetization, when performed in accordance with state-of-the-art parameters and supported by documented residual-risk evaluation, may in certain circumstances be analytically treated as capable of producing outputs falling outside the GDPR framework for purposes of Recital 26, thereby removing the output from the scope of the GDPR.[12] These proposals align with the broader multidisciplinary analysis of synthetic data in healthcare settings, which has examined the compatibility of synthetic generation with European health law and the evidence-based design paradigm, proposing a systematic reading of the current legal order that legitimises synthetic data as an emerging form of evidence construction.[13]

The third body of literature concerns the structural relationship between NIS2 and the GDPR in personal data breach scenarios. Both instruments impose mandatory reporting duties, yet they operate through different notification timelines, different thresholds and different addressees. Scholarship on the coordination of these parallel obligations has argued that a convergent toolkit – rather than two separate compliance silos – is needed.[14] The present article bridges these three bodies of literature by asking how synthetic data fit within the evolving, converging framework of cybersecurity and data protection law, and what doctrinal consequences follow. It further contributes to this debate by connecting the doctrinal analysis to two operational frameworks of particular relevance in the European NIS2 landscape: the Belgian CyberFundamentals Framework and the Italian ACN baseline specifications.

The argumentative order adopted here – starting from the absence of any obligation and working towards a contextual, evolutionary relevance – is preferable for at least three reasons. First, it is textually more faithful to the structure of NIS2. Secondly, it avoids conflating the availability of a technological solution with its legal necessity. Thirdly, it makes possible a more rigorous dialogue with the GDPR, because it allows the legal basis question to emerge only after a prior assessment of context, proportionality, residual risk and architectural design.[15]

Likelihood Attacks' (arXiv, 28 August 2025) https://arxiv.org/abs/2508.21146 accessed 3 July 2026.

[12] G Finocchiaro and others, *Il futuro regolatorio dei dati sintetici. La sintetizzazione dei dati come risorsa per ricerca scientifica, innovazione e politiche pubbliche nel panorama giuridico europeo* (Istituto Italiano per la Privacy e la Valorizzazione dei Dati and Data Intermediaries Alliance, 15 July 2024) paras 2.1 and 4.2, DOI: 10.5281/zenodo.12751884.

[13] A Brusadelli and others, 'Ripensare l'evidenza: approccio multidisciplinare all'uso dei dati sintetici in ambito sanitario — Parte I' (2025) BioLaw Journal – Rivista di BioDiritto, DOI: 10.15168/2284-4503-3747.

[14] G Canto Moniz, 'GDPR and NIS 2 Reporting Duties in Personal-Data Breaches: Structural Coexistence and Coordination' (2025) 11 European Data Protection Law Review 452.

[15] European Data Protection Board, *Guidelines 4/2019 on Article 25 Data Protection by Design and by Default* (version 2.0, 20 October 2020) paras 3–4.

2. Regulatory Silence and Open-Textured Obligations under NIS2.

The legal architecture of NIS2 is characterised by technological neutrality. The Directive seeks to achieve a high common level of cybersecurity across the Union by imposing risk-management, reporting, governance and supervisory obligations across a wide range of sectors.[16] Its legal technique is not to enumerate every safeguard that might satisfy those obligations, but rather to articulate objectives and standards against which organisational choices can be evaluated. Among these standards, the obligation to adopt 'appropriate and proportionate technical, operational and organisational measures' is central, as is the reference to the 'state of the art' in Article 21(1) NIS2.[17]

In areas marked by rapid technical evolution, such a legislative technique performs an important function. It avoids the obsolescence that would accompany a rigid catalogue of measures, while allowing the content of the legal duty to evolve in light of technological maturity, implementation feasibility, risk exposure and expected harms. This is what Hart described as the 'open texture' of legal language: general terms in legal rules necessarily possess a penumbral area in which their application to new facts cannot be resolved by textual interpretation alone.[18] In the present context, silence as to a specific technology is therefore better read as an instance of technological neutrality than as an implicit exclusion.

This reading is reinforced by the wider context of European digital regulation. The GDPR itself relies heavily on open-textured concepts. Articles 24, 25 and 32 GDPR all operate through context-sensitive requirements of appropriateness, effectiveness, proportionality and state of the art.[19] The European Data Protection Board has clarified, in its Guidelines 4/2019 on Article 25, that data protection by design and by default requires controllers to integrate data protection principles into system design from the earliest stages, taking account of context, risks, available technologies and implementation choices.[20] Both NIS2 and the GDPR thus require reasoned, context-sensitive assessment rather than rigid taxonomy.

The ENISA Technical Implementation Guidance published in June 2025, although developed in connection with the technical and methodological implementation of specific NIS2-related cybersecurity risk-management requirements, confirms a broader interpretative point: the content of the 'state of the art' standard is operationalised through a process of documented and continuous assessment rather than static prescription.[21] The legal significance of synthetic data does not therefore derive from the mere existence of an advanced technological solution, but from the documented relationship between an identified risk, the insufficiency of less intrusive alternatives in the

[16] NIS2 Directive (n 2) recitals 1–4 and art 1.
[17] ibid art 21(1), second subpara.
[18] Hart (n 4) 124–36; Schauer (n 4).
[19] GDPR (n 8) arts 24, 25 and 32.
[20] EDPB (n 15).
[21] European Union Agency for Cybersecurity, *Technical Implementation Guidance on Cybersecurity Risk Management Measures* (version 1.0, June 2025).

specific context, and the capacity of data transformation to reduce the likelihood or impact of incidents affecting natural persons.

Determining when a technology becomes relevant under Article 21(1) NIS2 requires more than pointing to its technical availability. Four elements are particularly relevant: the maturity of evaluation frameworks; the proportional feasibility of implementation in the sector concerned; the existence of a materially elevated risk that cannot be adequately reduced by conventional safeguards alone; and the emergence of supervisory, judicial or institutional practice treating the technology as relevant to adequacy assessments. These elements do not create a separate legal duty to adopt synthetic data. They rather indicate when the omission of less exposure-intensive alternatives may require explanation. In healthcare and comparable high-risk environments, this threshold is becoming increasingly significant, especially where structured personal data circulate in development, testing, analytics or third-party support contexts. This evolutionary standard becomes operationally meaningful only when it is connected to the control architectures through which NIS2 compliance is implemented in practice. The Belgian CyberFundamentals Framework and the Italian ACN baseline specifications provide two useful examples of this translation from open-ended legal duties to auditable cybersecurity measures. Their relevance will be examined in the next section, where synthetic data are analysed as a structural measure of exposure reduction rather than as a separate mandatory control.

3. Synthetic Data in the Assessment of Incident Impact and Structural Risk Reduction.

The importance of synthetic data in the NIS2 framework becomes clearer once the role of personal data exposure is properly understood. Article 21(1) NIS2 expressly requires that measures be designed to 'prevent or minimise the impact of incidents on recipients of their services and on other services,' thereby embedding the consequences for natural persons within the proportionality assessment.[22] In sectors such as healthcare, the compromise of real personal data may generate not only technical disruption but also discrimination, dignity-related harm, loss of trust and long-term consequences for individuals.

A useful analytical distinction may therefore be drawn between measures that protect real data while those data remain present within a system and design choices that may reduce the extent to which real data need to be present in particular environments. The former include encryption, access controls, logging, network segmentation, monitoring and incident-response mechanisms. The latter concern the organisation of data flows and processing environments so that identifiable or reasonably linkable data are not used where their presence is not functionally necessary. Where technically mature and appropriate for the specific use case, the generation and use of synthetic datasets may support this second type of exposure-reducing design choice.

[22] NIS2 Directive (n 2) art 21(1), first subpara.

This does not mean that synthetic data constitute, in themselves, a distinct category of security measure under NIS2, nor that their use automatically alters the legal status of the source data or of the resulting dataset. Their relevance is contextual: they may enable certain activities to be conducted without maintaining the same degree of exposure to real personal data. Encryption, access control, network security, monitoring, segregation of duties, logging and incident response – all identified in Article 21(2) NIS2 and elaborated in the ENISA 2025 guidance – remain indispensable components of any mature cybersecurity framework.[23] Synthetic data matter because they may reduce the need for real personal data to circulate in environments where the use of identifiable or reasonably linkable data is not functionally necessary.

The legal significance of this possibility lies not in assigning synthetic data a predetermined classification, but in its potential effect on the proportionality assessment. Where a technically and operationally suitable alternative is available, the continued use of real data in non-essential environments may require a more specific justification. Synthetic data thus enter the Article 21 analysis as one possible design alternative whose maturity, effectiveness, costs and limitations must be assessed in relation to the particular risk profile and operational objective.

3.1 The Heterogeneity of Synthetic Outputs and the Limits of De-identification Claims.

The term 'synthetic data' covers a heterogeneous spectrum of generation methods. Outputs produced by GANs, VAEs, rule-based simulation and differential privacy mechanisms differ in their statistical fidelity, robustness and residual re-identification risk profiles. [24] This heterogeneity carries important legal consequences that literature has not always sufficiently emphasized. The label 'synthetic' does not establish any particular privacy property.

The risks specific to generative model outputs deserve careful attention. Where a model overfits to the underlying training data – a risk that is particularly acute for atypical or rare individuals in the original population – synthetic records may retain sufficient statistical proximity to the originals to enable membership inference: an adversary can determine, with meaningful probability, that a specific individual's data was used to train the model and draw additional inferences about that person.[25] Research applying the DOMIAS framework has demonstrated that conventional membership inference attack benchmarks underestimate this risk, as they fail to account for the local overfitting patterns characteristic of tabular generative models.[26] Advances in privacy auditing for synthetic data – including the recently proposed GenLRA 'no-box' attack, which requires neither internal model access nor white-box knowledge – have further confirmed that overconfident de-identification

[23] NIS2 Directive (n 2) art 21(2)(a)–(j); ENISA, Technical Implementation Guidance (n 21).
[24] Cairo (n 7); Vallevik and others (n 1).
[25] El Emam, Mosquera and Bass (n 10); Z Zhang, C Yan and BA Malin, 'Membership inference attacks against synthetic health data' (2022) 125 Journal of Biomedical Informatics 103977.
[26] Van Breugel and others (n 11).

claims are empirically unsustainable without rigorous post-generation evaluation.[27]

A related concern arises from the mosaic effect: the risk that synthetic datasets which appear anonymous in isolation may become linkable to real individuals when combined with auxiliary datasets available to a motivated adversary. [28] As AI-driven analysis tools lower the cost of cross-dataset inference, the threshold of 'means reasonably likely to be used' under GDPR Recital 26 becomes progressively more demanding.[29] The ICO has reaffirmed that anonymisation must be assessed dynamically, taking account of both the data itself and the population of potential adversaries – a standard that requires periodic re-evaluation rather than a one-time determination.[30]

Emerging multidimensional evaluation frameworks now make it possible to assess, with growing precision, where on this spectrum a given output can be considered substantially anonymous. Technical studies in the healthcare domain have demonstrated that, when synthetic outputs satisfy a defined constellation of privacy metrics – including low membership inference vulnerability, low attribute disclosure risk, and independence from the original individuals' identifiers – the resulting dataset is no longer reasonably linkable to the original data subjects, thereby satisfying the threshold of GDPR Recital 26 in an evidence-grounded manner. [31] The legal analysis proposed in this article applies exclusively to synthetic outputs that have been subjected to such a documented multidimensional technical evaluation and where that evaluation supports a positive conclusion of substantial anonymisation. Without such evaluation, the label 'synthetic data' has no independent legal significance; with it, the legal classification progressively shifts from uncertain to determinable.

This reduction in expected harm has direct significance under both NIS2 and the GDPR. The EDPB Guidelines 9/2022 on personal data breach notification identify confidentiality, integrity and availability as the three dimensions of a breach; synthetic data may specifically reduce the severity of confidentiality breaches by removing or substantially attenuating the personal dimension of the compromised dataset.[32] The convergence of NIS2 and GDPR notification obligations in breach scenarios affecting personal data has been identified in scholarship as a key governance challenge; the use of synthetic data in non-

[27] Ward and others (n 11).
[28] G LaFever, 'Beyond GDPR: Unauthorized Reidentification and the Mosaic Effect in the EU AI Act' (IAPP, 10 August 2023) https://iapp.org/news/a/beyond-gdpr-unauthorized-reidentification-and-the-mosaic-effect-in-the-eu-ai-act accessed 3 July 2026.
[29] GDPR (n 8) recital 26; J Hinkel, 'AI in Real-World Data Analytics: The Quiet Threats of Re-identification and Synthetic Data' (Scientist.com, 21 March 2025) https://www.scientist.com/blog/ai-in-real-world-data-analytics-the-quiet-threats-of-re-identification-and-synthetic-data accessed 3 July 2026.
[30] Information Commissioner's Office, 'How do we ensure anonymisation is effective?' (ICO) https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/anonymisation/how-do-we-ensure-anonymisation-is-effective/ accessed 3 July 2026.
[31] Kaabachi and others (n 9); Hernandez and others (n 9); L Pilgram and others, 'A consensus privacy metrics framework for synthetic data' (2025) 6 *Patterns* 101320; GDPR (n 8) recital 26.
[32] European Data Protection Board, *Guidelines 9/2022 on personal data breach notification under Regulation 2016/679* (version 2.0, 28 March 2023).

essential environments offers a systemic architectural response to that challenge.[33]

The Italian and Belgian implementation contexts are particularly instructive because they show how the open-ended duty of adequacy under NIS2 is translated into operational control architectures. Although neither the Italian ACN baseline specifications nor the Belgian CyberFundamentals Framework expressly identify synthetic data as a named cybersecurity measure, both frameworks are organized around control families that presuppose the protection, classification and controlled circulation of information assets. These include, in different formulations, governance arrangements, asset management, access control, protection of information, secure development and testing, incident handling, business continuity and supplier-risk management.[34]

Within this logic, synthetic data may operate as an enabling design measure. Their relevance does not derive from express textual recognition, but from their functional alignment with existing control objectives. Where identifiability is not necessary, the use of empirically validated synthetic datasets in development, testing, training, analytics, demonstration or third-party support environments may reduce the value of compromised assets, limit the impact of confidentiality incidents, restrict supplier access to real personal data and provide evidence that less exposure-intensive alternatives have been considered. The point is therefore not that synthetic data constitute a distinct mandatory control under either framework. Rather, they may function as a cross-cutting measure capable of strengthening several existing control families where the underlying risk is sufficiently elevated.

The point is not that synthetic data constitute a distinct mandatory control under either framework. Rather, they can operate as a cross-cutting measure capable of strengthening several existing control families. Their legal relevance therefore derives from functional alignment with the objectives of the framework, not from express textual recognition.

4. From Regulatory Silence to an Evolutionary Standard.

Synthetic data should not be regarded as the object of a general obligation under NIS2. Nothing in the text of the Directive warrants such a conclusion.[35] Yet the absence of a general obligation does not exhaust the matter. A more compelling interpretation is that synthetic data may progressively become part of the set of safeguards that regulated entities are expected to consider under an evolutionary standard of adequacy. Where a technology matures to the point that it can significantly reduce exposure-related harms in certain high-risk

[33] Canto Moniz (n 14).

[34] Agenzia per la Cybersicurezza Nazionale, 'Modalità e specifiche di base' (ACN) https://www.acn.gov.it/portale/nis/modalita-specifiche-base accessed 3 July 2026; Centre for Cybersecurity Belgium, 'CyberFundamentals Framework' (Safeonweb) https://atwork.safeonweb.be/tools-resources/cyberfundamentals-framework accessed 3 July 2026.

[35] NIS2 Directive (n 2) art 21.

environments, and where it is realistically implementable at proportionate cost, the omission of that technology is no longer self-evidently neutral. In high-risk contexts, and particularly where real personal data circulate in non-essential or weakly controlled environments, the organisation may increasingly have to explain why, despite the availability of less exposure-intensive alternatives, the continued use of real personal data remains proportionate.

The dynamic character of this standard reflects the hermeneutic logic that Hart identified in legal language: the core of a legal rule is settled, but its penumbra expands as new facts bring previously unclassified cases within its normative shadow.[36] An evolutionary standard is preferable to a universal obligation for several reasons. It preserves textual fidelity. It accommodates sectoral diversity and technological unevenness. It gives real legal content to the notion of 'state of the art' without transforming that notion into a slogan. And it aligns with the spirit of the ENISA Technical Implementation Guidance of June 2025, which frames the content of Article 21 NIS2 as a dynamic, evidence-driven process of continuous assessment.[37]

An evolutionary standard should be implemented through periodic review complemented by event-driven reassessment, rather than through a one-off determination or an obligation to anticipate every hypothetical future attack. Reassessment should remain bounded by the means reasonably likely to be used, taking account of relevant entities, available auxiliary information, access conditions, cost, time, technology and reasonably foreseeable developments. Within that scope, relevant feedback should include not only confidentiality incidents and near-misses, but also reproducible privacy attacks, newly available auxiliary information and material changes in access, recipients or reuse conditions. The analogy with civil aviation safety is instructive: experience from incidents and near-misses informs the evolution of safety arrangements, while proactive identification of hazards enables intervention before harm occurs.

Two connected processes should therefore be distinguished: organisational reassessment of the evidence supporting a particular deployment, and the evolution of technical standards and regulatory guidance informed by accumulated evidence. The former should not await the latter. This distinction is consistent with Article 21(2)(f) NIS2, which requires policies and procedures to assess the effectiveness of cybersecurity risk-management measures, and Article 21(4), which requires necessary, appropriate and proportionate corrective measures without undue delay where an entity finds that it does not comply with the measures provided for in paragraph 2. A new technical finding does not automatically establish non-compliance; it may, however, require an assessment of whether the evidence supporting the existing arrangements remains sufficient.

This evolutionary standard must be distinguished from two adjacent concepts in the growing literature on design obligations in digital regulation. 'Regulation by design' refers to the embedding of regulatory constraints directly into technical systems *ex ante*. 'Security by design' in the NIS2/Cyber Resilience Act literature refers specifically to the design of secure-by-default

[36] Hart (n 4) ch VII; Schauer (n 4).
[37] ENISA, *Technical Implementation Guidance* (n 21).

products and services.[38] Design-based compliance, as used here, describes a broader compliance methodology in which legal principles – resilience, minimisation, accountability, fundamental rights protection – are materially embedded within system architecture through documented design choices, rather than satisfied *post hoc* by discrete controls. Synthetic data can exemplify this logic because, where appropriately implemented, they may simultaneously reduce personal data exposure, limit the severity of incidents affecting individuals, support privacy by design, and create a more coherent governance model for distributed processing environments.

It is worth noting the limits of this evolutionary argument. The four conditions identified in Section 2 – convergence of evaluation frameworks, proportionate implementability, materially elevated risk profile, and engagement by supervisory practice – function as a threshold test, not a balancing exercise. Unless all four are satisfied in the specific context, the justificatory burden described above does not arise. Entities operating in low-risk environments, or processing data that is functionally inseparable from identifiability, remain subject to the standard adequacy analysis under Article 21(1) NIS2 without any presumption in favour of synthetic transformation.

This threshold is especially important in the context of operational frameworks such as CyberFundamentals and the ACN baseline specifications: the mapping of synthetic data to existing control families does not, by itself, transform them into a required measure. It merely strengthens the argument that their consideration may be expected where the underlying risk is sufficiently elevated.

5. Consequences for Articles 6 and 9 GDPR.

Once synthetic data are positioned within an evolutionary adequacy standard under NIS2, the question of legal basis under the GDPR can be approached more rigorously. It would be mistaken to begin by asserting that the generation of synthetic data is in itself 'necessary' and therefore directly grounded in Article 6(1)(c) GDPR. Article 6(1)(c) can become relevant, if at all, only at the end of a prior inquiry demonstrating that synthetic data have acquired concrete relevance within the controller's NIS2 compliance architecture.[39]

The strict character of the necessity test under Article 6(1)(c) GDPR is confirmed by the CJEU's case law. In *Orange România* (Case C-61/19), the Court held that the necessity of processing must be assessed specifically in relation to the identified legal obligation and cannot be inferred from the general desirability of a measure for the controller.[40] This principle applies *a fortiori* to

[38] Djeffal (n 6); Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements and amending Regulations (EU) No 168/2013 and (EU) 2019/1020 and Directive (EU) 2020/1828 [2024] OJ L, 2024/2847, 20.11.2024 (Cyber Resilience Act).
[39] EDPB (n 15).
[40] Case C-61/19 *Orange România SA v Autoritatea Națională de Supraveghere a Prelucrării Datelor cu Caracter Personal* EU:C:2020:901, paras 35–40.

a legal obligation whose very content is itself subject to contextual evolution: necessity cannot be established at the level of an abstract regulatory standard; it must be demonstrated with reference to a specific identified binding obligation, a specific identified risk, and a specific documented assessment showing that the measure in question is the least intrusive effective option available.

5.1 Article 6(1)(c) GDPR: A Context-Sensitive Framework.

While Article 6(1)(c) GDPR requires necessity in relation to a legal obligation, it would be doctrinally incorrect to treat this requirement as an impenetrable barrier in the cybersecurity context. At the same time, the open-textured character of the adequacy standard under Article 21 NIS2 means that the content of the relevant legal obligation is not exhausted by a fixed and technologically static catalogue of measures. Article 21 establishes a context-sensitive obligation structured around adequacy, proportionality, the state of the art, implementation costs and the seriousness of the relevant risks. As technical evaluation frameworks mature, as ENISA guidance further operationalises the state of the art, and as supervisory and judicial practice becomes more specific in relation to particular risk profiles, the practical content of that obligation may become correspondingly more determinate. Within this framework, the availability of technically mature synthetic-data methods may become relevant to the assessment of necessity. The argument is not that Article 21 NIS2 directly requires the use of synthetic data, or that synthetic-data generation constitutes a legal obligation in itself. Rather, where such methods are technically feasible and appropriate for the relevant use case, their availability may affect the justification for continuing to rely on real-data-intensive practices in environments where the operational objective could be achieved with a lower level of exposure.[41]

Necessity should therefore be understood as a contextual determination. It must be assessed in relation to the specific legal obligation, the processing environment, the risks involved, the available technical alternatives and the operational purpose pursued. The evolution of the state of the art may influence that assessment, but it does not dispense with the requirement to identify a sufficiently clear legal obligation and to demonstrate that the processing undertaken is objectively necessary for compliance with it.

This reading is reinforced by the broader trajectory of EU legislative developments, albeit with an important caveat of method. The European Commission's Digital Omnibus package, presented in November 2025, proposes a number of amendments to EU digital regulation that remain, at the time of writing, in an early stage of co-legislative negotiation. This article does not treat any element of the Digital Omnibus as settled law, and no conclusion advanced here depends on the adoption of those proposals. Their relevance is limited to what they may indicate about the direction of EU legislative intent. Among the GDPR-related components of the Digital Omnibus – which are separate from the AI Act amendments and have not yet been the subject of any

[41] Vallevik and others (n 1).

political agreement – the Commission proposed provisions designed to facilitate the processing of special categories of personal data in AI development contexts, as well as provisions that would identify the use of state-of-the-art privacy-preserving techniques as a relevant factor in assessments under Article 6(1)(f) GDPR. The precise scope, numbering and definitive wording of these provisions are not yet available; the text remains subject to amendment by the Parliament and the Council, and no political agreement on the GDPR amendments has been reached at the time of writing.[42] The provisional political agreement reached on 7 May 2026 was subsequently formally adopted as Regulation (EU) 2026/1744, which entered into force on 27 July 2026. The final act concerns targeted amendments to the AI Act, including the phased application of obligations for high-risk AI systems, and does not encompass any of the proposed GDPR amendments.[43] Should the legislative process eventually result in provisions broadly consistent with the Commission's proposal, this may provide an additional external indication that the availability of less exposure-intensive data-design alternatives is acquiring greater normative relevance in high-risk processing environments. Such developments would not, however, establish that synthetic data possess a predetermined legal status or that their use is generally required. At most, they would reinforce the doctrinal hypothesis that technically mature and appropriately evaluated alternatives may become increasingly relevant to assessments of necessity, proportionality and the state of the art. The argument developed in this article remains independent of the adoption of those proposals.

Where a controller seeks to rely on Article 6(1)(c) GDPR for processing associated with the generation or evaluation of synthetic datasets, several conditions must therefore be satisfied. First, the controller must identify the specific legal obligation whose performance requires the processing. Secondly, it must demonstrate, through a documented risk assessment, that the continued circulation of real personal data in the relevant environment creates a materially elevated risk profile. Thirdly, the controller must assess whether conventional safeguards under Article 21(2) NIS2, including encryption, access controls, segregation and monitoring, sufficiently reduce that risk. It must establish that synthetic transformation, assessed through multidimensional technical evaluation confirming substantial anonymisation,[44] is feasible and represents the least intrusive option capable of achieving the operational objective. And the processing must be embedded within an accountability structure that includes technical documentation and, where required, a data protection impact assessment.[45] These are demanding conditions, but they are

[42] European Commission, 'Simpler EU digital rules and new digital wallets to save businesses over EUR 5 billion' (Press Release IP/25/2718, 18 November 2025).
[43] Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI) [2026] OJ L 2026/1744.
[44] Kaabachi and others (n 9); Pilgram and others (n 31).
[45] Article 29 Working Party, *Guidelines on Data Protection Impact Assessment and Determining Whether Processing Is "Likely to Result in a High Risk" for the Purposes of Regulation 2016/679* (WP248 rev.01, 4 October 2017), as endorsed by the EDPB.

conditions whose satisfaction becomes increasingly plausible as the technical and regulatory environment matures. Far from being a near-closed door, Article 6(1)(c) in this context functions as a progressively widening gateway whose opening tracks the evolution of the adequacy standard itself.

In many cases, however, controllers will find that Article 6(1)(e) GDPR – processing for a task in the public interest – offers a more direct and less demanding route, particularly in healthcare and public sector environments. Article 6(1)(f) GDPR may also be relevant where applicable, provided that the necessity and balancing tests are satisfied. In that assessment, the availability of synthetic datasets may be considered as one of the possible design alternatives capable of reducing exposure and potential harm, without being characterised as either the direct object of a legal obligation or an autonomous legal category.[46] The key is that the reasoning proceeds in the right order: regulatory silence; evolutionary standard; contextual necessity – never in reverse.

5.2 Article 9 GDPR, the European Health Data Space, and the Evolving Legislative Horizon.

An even greater degree of analytical rigor is required under Article 9 GDPR. In health data contexts, Article 9 provides the gateway, and the conditions for processing are substantially more restrictive than under Article 6. Articles 9(2)(i) and 9(2)(j) GDPR permit processing of special categories of data respectively for public interest reasons in the area of public health and for scientific research purposes, but only where such processing is provided for by Union or Member State law and is subject to appropriate and specific safeguards.[47]

The normative landscape for health data has been substantially restructured by Regulation (EU) 2025/327 on the European Health Data Space (EHDS), which entered into force on 26 March 2025. The EHDS Regulation establishes a governance framework for the secondary use of electronic health data through a system of health data access bodies (HDABs), data permits and the HealthData@EU cross-border infrastructure. Under Articles 53 and 54 EHDS, HDABs are empowered to impose conditions on data access – including requirements to use privacy-enhancing technologies – as a condition of permit issuance.[48] Privacy-enhancing techniques including synthetic data generation are embedded in the design of this framework as tools for enabling secondary use while reducing the exposure of individual health data.[49] The EHDS confirms that in the health data sector the trajectory sketched in this article reflects the direction of positive law rather than doctrinal speculation. Any legal analysis of

[46] Cairo (n 7).
[47] GDPR (n 8) arts 9(2)(i) and 9(2)(j).
[48] EHDS Regulation (n 1) arts 53–54; TEHDAS2, *M5.2 Draft Guideline for Health Data Access Bodies on Minimum Categories and Limitations on the Reuse of Health Data* (16 September 2025).
[49] EHDS Regulation (n 1) recitals 57–62; S Van Drumpt and others, 'Secondary use under the European Health Data Space: setting the scene and towards a research agenda on privacy-enhancing technologies' (2025) 7 Frontiers in Digital Health 1602101.

synthetic data in the healthcare context that fails to engage the EHDS Regulation is doctrinally incomplete as of 2025.

The legal status of the resulting synthetic dataset is not automatic. Whether it qualifies as non-personal data depends on the residual risk of re-identification, assessed in light of GDPR Recital 26 and the means reasonably likely to be used.[50] As noted in Section 3.1, this determination must be revisited as generation methods and re-identification techniques evolve – a consideration that takes on particular urgency given the speed at which adversarial inference tools are advancing.[51] Where the multidimensional assessment confirms substantial anonymisation, the output exits the GDPR's scope entirely, and neither Article 6 nor Article 9 need be invoked for its subsequent use. This is the architecturally optimal outcome – and it is precisely the outcome that design-based compliance, properly structured, is capable of achieving.

6. Design-Based Compliance and the Future of Privacy and Cybersecurity Law.

The broader significance of this analysis lies in the model of compliance that it presupposes. In both cybersecurity and data protection law, compliance is increasingly assessed not only by reference to the presence of discrete controls, but also by reference to the architecture within which data circulate, risks are generated, and harms are prevented or mitigated. Design-based compliance differs from more traditional approaches in one important respect: it does not treat legal obligations as external constraints to be satisfied after technical choices have already been made. Instead, it asks whether the system has been designed so that legal principles – resilience, minimisation, accountability, confidentiality, reduction of harm – are materially embedded within its structure from the outset.[52]

Synthetic data can exemplify this logic because, where appropriately implemented, they may simultaneously reduce the exposure of personal data, limit the severity of incidents affecting individuals, support privacy by design under Article 25 GDPR, and create a more coherent governance model for distributed processing environments.[53] This is not a matter of technological optimism. Synthetic data remain fallible; not all outputs are equally robust; some contexts are unsuitable; and overconfidence in de-identification claims remains a serious and empirically documented regulatory danger, as the membership inference literature reviewed in Section 3.1 makes clear.[54] These caveats confirm rather than weaken the argument: they explain why an evolutionary standard is preferable to a universal obligation, and why the

[50] GDPR (n 8) recital 26; ICO (n 30).
[51] Hinkel (n 29); Ward and others (n 11).
[52] EDPB (n 15); Djeffal (n 6).
[53] EHDS Regulation (n 1) recitals 57–62; NIS2 Directive (n 2) art 21(1); GDPR (n 8) art 25.
[54] El Emam, Mosquera and Bass (n 10); Van Breugel and others (n 11); Zhang, Yan and Malin (n 25).

approach must be one of reasoned, context-sensitive engagement grounded in documented multidimensional technical evaluation.

The design-based compliance model connects to emerging scholarship on the governance of digital systems, which has increasingly emphasized the limitations of *ex post* regulatory controls and the necessity of embedding governance principles into system architecture at the design stage. The relationship between design obligations and fundamental rights impact assessment has been explored in the context of the AI Act; analogous design logic pervades Article 25 GDPR and the EHDS framework.[55] In that intellectual environment, this contribution should be understood as part of a wider transformation in European digital law: a movement from compliance as a checklist activity to compliance as an architecturally embedded methodology of risk governance.

In healthcare, this paradigm finds concrete operational expression across at least three dimensions. First, it translates the open-textured obligations of NIS2, together with the operationalization of the state-of-the-art standard in ENISA guidance, into architectures capable of structurally reducing data exposure – for instance, by replacing real datasets with technically validated synthetic outputs in non-essential environments such as development, testing and analytics. Secondly, it aligns with the GDPR and EHDS trajectory towards advanced minimisation, anonymisation and privacy-enhancing techniques as instruments of risk governance. Within this framework, the decision to keep real personal data in circulation cannot be assessed merely as a matter of technical convenience but must be justified in light of less intrusive alternatives where those alternatives are technically mature, proportionate and capable of preserving the relevant operational purpose. Thirdly, the same logic is increasingly visible in the operational standards of digital healthcare systems, including national guidance on telemonitoring, digital health infrastructures and health data ecosystems, where security, interoperability, data protection and accountability are progressively treated as integrated design requirements rather than isolated compliance layers.[56]

The legal consequence is that, in high-risk healthcare contexts, compliance is not measured only by the presence of individual security controls. It also depends on the organization's capacity to demonstrate that it has considered – and, where proportionate, adopted – data transformation solutions, including technically validated synthetic data generation, as components of overall adequacy under Article 21 NIS2 and Articles 25 and 32 GDPR.

6.1 Limits and Conditions of the Model.

A candid account of design-based compliance must acknowledge its limits. The argument advanced here does not establish a general justificatory burden

[55] Djeffal (n 6); Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 2024/1689 (Artificial Intelligence Act), arts 9, 27 and 43.

[56] See Agenzia Nazionale per i Servizi Sanitari Regionali, Linee di indirizzo percorso evolutivo. Sistemi Medicali per Telemonitoraggio (version 1.8, 2024); Decreto del Ministro della Salute 31 December 2024, Istituzione dell'Ecosistema Dati Sanitari, GU Serie Generale n 53, 5 March 2025.

applicable to all regulated entities, sectors or forms of processing. Its present legal force remains context-dependent and asymmetric.[57]

First, the model depends on the maturation of technical evaluation standards. Where such standards are contested, underdeveloped or poorly adapted to the relevant data type, the legal classification of a synthetic output cannot support a sufficiently robust necessity assessment. The strongest case therefore arises in healthcare and comparable high-risk environments involving structured tabular data, where privacy-utility metrics, membership inference testing, attribute disclosure analysis and post-generation audit methodologies are comparatively more mature. The argument is weaker for unstructured data, imaging, genomic data, free-text records and other categories for which fidelity, utility and residual identifiability remain harder to assess in a standardised and legally defensible manner.

Secondly, any justificatory burden should be understood as specific rather than general. At the present stage of technological and regulatory development, it may arise only where cumulative conditions are satisfied: the processing concerns highly sensitive or high-impact data; the relevant environment is non-essential to the primary service, such as development, testing, analytics, training, demonstration or third-party support; real personal data are not functionally indispensable; synthetic transformation is technically feasible; recognized post-generation evaluation can be performed; and the cost of implementation is proportionate to the risk addressed. In such circumstances, an entity that continues to circulate real personal data may increasingly be expected to explain why less exposure-intensive alternatives were not adopted.

Outside those comparatively mature contexts, the legal effect is more modest. Synthetic data should be treated as an emerging design option capable of strengthening compliance with NIS2, GDPR and sectoral cybersecurity frameworks, rather than as a measure whose omission automatically requires justification. The relevant expectation is one of reasoned consideration, not presumed adoption. A conclusion that synthetic data are unsuitable, immature or disproportionate may remain entirely defensible where it is properly documented.

The model also presupposes institutional capacity. Risk assessments, technical evaluations, architectural records, DPIAs, dataset inventories and periodic re-evaluations may impose costs that are disproportionate for smaller or less technically mature entities. This reinforces the need for sectoral guidance from national competent authorities and ENISA, capable of clarifying when synthetic data should be considered, preferred or merely monitored as an emerging safeguard.

Finally, where entities rely on synthetic outputs as substantially non-identifiable, the assessment cannot be treated as final. Periodic review and event-driven reassessment should determine whether the evidence supporting the relevant generation, access and reuse configuration remains sufficient. Such review should be supported by a proportionate record of the dataset version, intended use, access conditions, relevant threat-model assumptions

and evaluation results. Where new evidence materially undermines the previous assessment, proportionate responses may include strengthened access conditions, suspension of further distribution or regeneration of the dataset, followed by renewed evaluation. Where warranted, the entity should also reconsider the legal classification of the output and any applicable obligations.

Public dissemination requires particular caution because released copies may be practically irretrievable; post-release reassessment cannot therefore substitute for robust pre-release evaluation. The scope for subsequent corrective action should itself inform the initial choice of release conditions. Design-based compliance is consequently not a shortcut to compliance, but a disciplined method for embedding proportionality, minimisation and cyber-resilience into system architecture, including the capacity to revise operational decisions when their evidentiary basis changes.[58]

7. Conclusion: Synthetic Data as an Emerging Standard of Adequacy.

.Synthetic data are not expressly contemplated by NIS2 and the Directive imposes no general obligation to adopt them. That is the correct starting point of the analysis. It is not, however, the end of it.

Because NIS2 is built upon open-ended standards of appropriateness, proportionality and risk management – and because the impact of incidents on natural persons increasingly informs the legal assessment of cybersecurity harm – synthetic data may progressively become part of the set of safeguards that regulated entities are expected to consider where the circulation of real personal data in non-essential environments generates materially elevated risks. The ENISA Technical Implementation Guidance of June 2025 and the EHDS Regulation (EU) 2025/327 both indicate that this trajectory is being confirmed in operational guidance and hard law alike.

The proper legal formulation is, accordingly, evolutionary. In high-risk settings – particularly those involving large-scale health data processing, complex architectures and distributed non-essential environments – continued reliance on real personal data without genuine consideration of materially less harmful design alternatives may increasingly attract a substantial justificatory burden. That burden does not amount to a formal legal obligation derived from legislative command; it emerges, rather, from the structural logic of open-textured adequacy standards as interpreted against a maturing technological background. As multidimensional evaluation frameworks consolidate and as industry standards and supervisory guidance converge on specific privacy-metric thresholds, the weight of that burden will grow in proportion to the demonstrated availability and technical robustness of alternative system designs, including technically validated synthetic transformation where appropriate. Organisations that examine this option in good faith and conclude, on well-reasoned grounds, that it is not feasible or not suitable for their specific context can still remain in a defensible compliance

[58] ICO (n 30); EDPB (n 32).

position, provided that conclusion is supported by a specific and proportionate explanation of why real personal data continue to circulate. What becomes increasingly difficult to defend, in an evolutionary reading of adequacy, is not a negative conclusion as such, but the absence of any engagement with the issue at all: a failure to address the question seriously calls into question whether continued reliance on real personal data in non-essential environments can still be regarded as proportionate. Only after this prior inquiry does the question of legal basis properly arise. In certain carefully delimited contexts, Article 6(1)(c) GDPR may become available because synthetic transformation is shown – through a rigorous necessity assessment consistent with the CJEU's case law – to be required for compliance with binding cybersecurity obligations. In others, Articles 6(1)(e) or 6(1)(f) will remain more appropriate. In health data contexts, the EHDS governance framework provides the primary normative environment, within which the use of synthetic data must be assessed alongside Articles 9(2)(i) and 9(2)(j) GDPR and the data permit regime of Articles 53 and 54 EHDS. And where multidimensional technical evaluation confirms that synthetic outputs are no longer reasonably linkable to real individuals, the question of legal basis does not arise at all, because the data have exited the GDPR's scope entirely.

What matters, throughout, is that the legal reasoning proceed in the right order: from regulatory silence, through evolutionary standard, to contextual necessity – never in reverse. Synthetic data illuminate a wider transformation in European digital law: a movement toward regulation that operates through architecture, methodology and documented assessment, rather than through exhaustive technological catalogues. Design-based compliance is accordingly the framework within which the ongoing normative relationship between cybersecurity law, data protection, and emerging data practices is most likely to be worked out – not through prescriptive mandates, but through the progressive consolidation of standards, guidance, and institutional practice. It is on that terrain that the legal system is already moving, whether or not it has yet found the language to say so.