Do the EHDS, GDPR and AI Act Protect Against the Risk of Discrimination in DQUL-Based Health Technology Assessment?

Sabire Sanem Yilmaz, Andrea Parziale*

ABSTRACT: This article examines whether the European Health Data Space (EHDS), the General Data Protection Regulation (GDPR), and the Artificial Intelligence Act adequately prevent discriminatory outcomes in DQUL-based health technology assessment. Focusing on the Data Quality and Utility Label (DQUL) and its interaction with Quality-Adjusted Life Year (QALY) cost-effectiveness models, the article argues that regulatory compliance does not eliminate the risk of structural disadvantage. It introduces the concept of numerical discrimination to describe inequality arising from threshold-based evaluative architectures that systematically undervalue underrepresented populations. Through doctrinal analysis of EU equality, data protection, and AI governance law, the article demonstrates that current frameworks struggle to capture distributive harms arising from hierarchical data governance and proposes reforms that integrate distributive impact assessment and population representativeness into health data architecture.

KEYWORDS: European Health Data Space (EHDS); Data Quality and Utility Label (DQUL); numerical discrimination; Health Technology Assessment (HTA); Quality-Adjusted Life Years (QALYs)

1. Introduction

The European Health Data Space (EHDS)[1] introduces the Data Quality and Utility Label (DQUL)[2] as a governance mechanism intended to enhance transparency, reliability, and trust in the secondary use of health data, including its use in health technology assessment (HTA) mechanisms structured around cost-effectiveness analysis, particularly through Quality-Adjusted Life Years

* *Sabire Sanem Yilmaz: PhD in Health Science and Technology, Scuola Sant'Anna School of Advanced Studies. Mail: sabire.yilmaz@santannapisa.it; Andrea Parziale: Assistant Professor of Comparative Private Law, Sant'Anna School of Advanced Studies. Mail: andrea.parziale@santannapisa.it.* *This article was subject to double blind peer review process.*

[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.

[2] The Data Quality and Utility Label (DQUL) is regulated under Article 78 of the EHDS Regulation.

(QALYs).[3] These appraisal systems are progressively data-driven and, in some contexts, algorithmically supported.

This article argues that even full compliance with the EHDS, the GDPR, and the AI Act does not eliminate the risk of numerical discrimination in DQUL-based systems. The interaction between hierarchical data quality labelling and threshold-based cost-effectiveness models may systematically disadvantage populations that are statistically underrepresented, clinically heterogeneous, or difficult to standardise, including people with disabilities, rare disease patients, elderly populations, and socioeconomically marginalised groups.

To conceptualise this phenomenon, the article introduces the notion of numerical discrimination. Numerical discrimination refers to a form of structural disadvantage that arises from threshold-based evaluative architectures, in which formally neutral numerical criteria systematically undervalue statistically underrepresented or heterogeneous populations, without explicit discriminatory intent and despite regulatory compliance.[4] This is fundamentally a problem of threshold-based evaluation, the lack of neutrality in the normative tools relied upon, and the distinction between formal equality and substantive equality.

Statistical discrimination, as theorised in foundational economic scholarship by Edmund Phelps and Kenneth Arrow, is a mechanism whereby decision-makers operating under incomplete information rely on group-average statistics to predict individuals' outcomes, generating biased outcomes without explicit discriminatory intent.[5] In Phelps's original formulation, a decision-maker lacking individualised performance data uses group membership as a proxy for expected outcomes, systematically disadvantaging members of groups associated with lower averages. Arrow's complementary model demonstrated that such stereotyping generates self-reinforcing dynamics: statistical discrimination tends to persist structurally because it may depress investment incentives for disadvantaged group members, partially validating the original assumption.[6] As Moro's comprehensive survey shows, policy responses diverge depending on whether bias is information-based or structural: information-asymmetry models predict that supplying better individualised data will correct discrimination. In contrast, structural models indicate deeper institutional roots requiring systemic reform.[7]

---

[3] Quality-Adjusted Life Years (QALYs) calculate the extent of full, healthy life years a medical intervention or treatment grants to an individual. 1 QALY represents one full year spent in perfect health, while 0 represents death; the threshold fluctuates on a scale between 0 and 1.

[4] V. CHARLTON, M. DISTEFANO, *An empirical ethics study of the coherence of NICE technology appraisal policy and its implications for moral justification*, in *BMC Med Ethics*, 25(28), 2024.

[5] E.S. PHELPS, *The Statistical Theory of Racism and Sexism*, in *American Economic Review*, 62(4), 1972, 659-661. This foundational paper formalises the mechanism by which decision-makers relying on group averages as proxies for individual ability systematically disadvantage members of groups associated with lower average outcomes, even in the absence of animus.

[6] K.J. ARROW, *The Theory of Discrimination*, in O. ASHENFELTER, A. REES (eds.), *Discrimination in Labor Markets*, Princeton, 1973, 3-33. Arrow's complementary model highlights the self-reinforcing dynamics of statistical discrimination: systematic underinvestment by disadvantaged groups may partially validate the employer's original stereotype, creating structural persistence.

[7] A. MORO, *Theories of Statistical Discrimination and Affirmative Action: A Survey*, in J.B. BENHABIB, A. BISIN, M. JACKSON (eds.), *Handbook of Social Economics*, 1A, North-Holland, Amsterdam, 2011, 115-155. For a comprehensive updated synthesis see also A. MORO, *Statistical discrimination*, in *The New Palgrave Dictionary of Economics*, London, 2018, 12996-13001.

Proxy discrimination occurs when an algorithm produces discriminatory outcomes by utilising variables that are highly correlated with protected attributes (such as postal codes, healthcare utilisation patterns, or nutritional records) even when those protected categories are not explicitly included in the dataset.[8] Under the EU indirect discrimination doctrine, analysed in Section 4, proxy discrimination is cognizable as an apparently neutral provision or practice that produces a disproportionate disadvantage to a protected group, unless objectively justified.[9]

Numerical discrimination, as introduced in this article, shares structural features with both concepts but is analytically distinct from each. Unlike statistical discrimination, the focus is not on the epistemic assumptions of individual decision-makers but on the institutional architecture of threshold-based evaluation systems and the hierarchical data-quality frameworks that provide their evidentiary basis. Unlike proxy discrimination, the mechanism is not the substitution of a correlated variable for a protected characteristic, but the systematic undervaluation of entire populations through formally neutral quantitative thresholds that, by privileging large, homogeneous, and standardised datasets, structurally assign lower epistemic weight to heterogeneous or clinically complex populations. The disadvantage does not originate in variable selection or in the decision-maker's inferential shortcuts, but in the normative assumptions embedded in the threshold architecture itself.

This distinction carries important doctrinal implications. Statistical discrimination is amenable, in principle, to informational remedies: better individualised data reduces reliance on group averages. Proxy discrimination can be addressed by constraining the use of correlated variables. Neither remedy is structurally adequate for numerical discrimination because the source of the disadvantage lies in the threshold design and the data quality hierarchy. The appropriate regulatory response – articulated in the proposals in Section 7 – must therefore target the threshold architecture itself: its representativeness requirements, transparency obligations, and the normative assumptions it embeds in ostensibly technical evaluative frameworks.

The central claim of the article is that discrimination in this context is not the result of unlawful intent or regulatory breach, but of a structural misalignment between legal equality norms and quantitative decision-making architectures. The article does not argue that the EHDS, GDPR, or AI Act are deficient in their protective ambition. Rather, it contends that their protective logic operates at a different normative level than the structural mechanisms through which disadvantage emerges in data-driven appraisal systems. The problem is the architectural incompatibility between formally neutral regulation and threshold-based evaluative logic.

2. DQUL

Under Article 78 of the EHDS, datasets made available by Health Data Access Bodies (HDABs) may be accompanied by a Data Quality and Utility Label (DQUL). DQUL is designed as a standardised mechanism

[8] G.M. JOHNSON, *The hard proxy problem: proxies aren't intentional; they're intentional*, in *Philosophical Studies*, 182, 2025, 1383-1411.

[9] K. LIPPERT-RASMUSSEN, *"We are all Different": Statistical Discrimination and the Right to be Treated as an Individual*, in *J Ethics*, 15, 2011, 47-59.

for characterising data quality attributes, including completeness, accuracy, consistency, and maturity.[10] In the EHDS, DQUL is defined as a graphical diagram, including a scale, that represents both the data quality and the conditions of use of a dataset.[11]

DQUL is structured along two key elements: data quality and conditions of use. The vertical element represents the graded assessment of data quality, ranging from low to high, based on indicators such as accuracy, completeness, consistency, maturity, and population representativeness. The horizontal element captures the spectrum of use conditions, ranging from highly restrictive access regimes to more open secondary-use environments. Each dataset is positioned within this matrix, making visible not only its technical robustness but also the normative constraints governing its deployment.[12]

These two key elements ensure that data governance is not reduced to compliance alone. A dataset with high quality but restrictive use conditions occupies a different regulatory and distributive position than a similarly robust dataset that is openly reusable.[13] In contrast, poor-quality datasets (especially those lacking representativeness) pose distinct risks when used in open evaluative environments such as cost-effectiveness analysis. By situating datasets within this analytical framework, the system clarifies how technical characteristics and legal constraints jointly shape downstream decision-making processes, particularly in HTA contexts.[14]

The approach we describe as a hierarchical data quality labelling system[15] requires that, at the moment datasets are labelled, the conditions of use are not only normatively justified but also practically reflected in the dataset itself. These systems will predominantly operate in integration with AI, which will automatically determine the datasets' maturity matrix and label characteristics.[16] Human intervention remains necessary, given that using datasets in AI systems is a high-risk activity. While the labelling process itself is not inherently high-risk under the AI Act, the risk arises when the resulting labelled data is utilised in a high-risk AI system, such as diagnostic software. Therefore, data governance obligations under Article 10 of the AI Act should be integrated into the EHDS framework.[17]

---

[10] Quantum Consortium, *The Health Data Quality Label, Deliverable 1.1: Specification Dataset's Quality and Utility Label*, 8.11.2024, https://zenodo.org/record/14937423 (last visited 22/06/2026).

[11] See Article 2(2)(aa) Regulation (EU) 2025/327.

[12] J. Daumas, A. Schäfer, L. Dolanski-Aghamanoukjan, K. Weishäupl, M.M. Kuhrn, N. Schutte, C. Proietti Mercuri, E. Bernal-Delgado, *Deliverable 1.1 Specification of the data sets' quality and utility label (2.3.0)*, Zenodo, 2025.

[13] Á. Sánchez-García, C. Proietti Mercuri, N. Schutte, F. Estupiñán-Romero, C. Tellería-Orriols, A. Doñate-Martínez, J.M. García-Gómez, E. Bernal-Delgado, C. Sáez, on behalf of QUANTUM, *Landscape analysis towards data quality and utility labelling in the European Health Data Space*, in *European Journal of Public Health*, 2026.

[14] M. Mashoufi, H. Ayatollahi, D. Khorasani-Zavareh *et al.*, *Data quality in health care: main concepts and assessment methodologies*, in *Methods of Information in Medicine*, 62, 2023, 5-18.

[15] While there is no official terminology designating the Data Quality and Utility Label mechanism as such within the EHDS framework, we have characterised it in our study as a hierarchical data quality and utility label. The primary rationale is the decisive influence that quality and utility indicators exert in governing secondary data use, establishing an implicit hierarchy among datasets and shaping data users' selection behaviour.

[16] J. Mariño, E. Kasbohm, S. Struckmann, L.A. Kapsner, C.O. Schmidt, R. *Packages for Data Quality Assessments and Data Monitoring: A Software Scoping Review with Recommendations for Future Developments*, in *Applied Sciences*, 12(9), 2022, 4238.

[17] P. Cervera De La Cruz, J.W. Hazel, M. Shabani, I.G. Cohen, *Driving AI Health Innovation through the European Health Data Space: Opportunities and Challenges for Non-EU Country Participation*, in *NEJM AI*, AIpc2501270, 2026.

If the DQUL process is automated via AI, the system must first be trained on expert-annotated datasets to model the relationship between data characteristics and their normative scores. Once deployed, the AI can be integrated into the data lifecycle to enable automated labelling of new datasets.[18] Validation data monitors the AI model's generalisation; testing data provides independent performance evaluation.[19] Once trained, AI can process new datasets to automatically generate DQUL. While AI may reliably reproduce technical quality indicators, the normative justification of use conditions cannot be reduced to pattern recognition. Therefore, human oversight remains essential not merely for risk mitigation under the AI Act, but for preserving the normative character of the DQUL itself.[20]

The risk that DQUL-based HTA may lead to discrimination arises if the label is inaccurate. Under the EHDS, inaccurate DQULs may be revoked following HDAB supervision. However, the primary issue is not the physical existence of the data, but its legal circulation within the health data ecosystem.[21] The revocation of a label does not delete the data, but strips the dataset of its legal validity for secondary use.[22] It is also debatable whether revocation by a central body has practical binding effect across all national SPE systems without cross-border synchronisation mechanisms.[23]

The main concern is not the weakness of mathematical measurement itself, but the contestability of its normative justification.[24] The language used to conceptualise the normative aspects of HTA is often imprecise and applied inconsistently, thereby undermining transparency.[25] The label is formally agnostic as to purpose.[26] However, in practice, hierarchical representations of quality and utility impose an implicit ordering among datasets. Although DQUL is formally purpose-agnostic under Article 78 of the EHDS, its hierarchical structure generates de facto epistemic ordering.[27]

In HTA environments, where evidentiary robustness directly influences reimbursement and coverage decisions, this ordering may translate into a structural preference for certain datasets. Large-scale, homogeneous, standardised, and statistically robust datasets are more likely to meet favourable quality indicators.[28] To truly overcome the risk of numerical discrimination, the DQUL must evolve beyond a mere tech-

[18] T. Petročnik, *Health data between health(care) and the data economy*, in *Technology and Regulation*, 2022, https://techreg.org/article/view/13284 (last visited 22/06/2026).

[19] D. Fåhraeus, J. Reichel, S. Slokenberga, *The European Health Data Space: Challenges and Opportunities*, Stockholm, 2024, https://urn.kb.se/resolve?urn=urn:nbn:se:su:diva-227156 (last visited 22/06/2026).

[20] J. Daumas *et al.*, *op. cit.*

[21] L. Ehrlinger, W. Wöß, *A Survey of Data Quality Measurement and Monitoring Tools*, in *Frontiers in Big Data*, 5, 2022, 850611.

[22] E. Hyvonen, R. Albertoni, A. Isaac, *Introducing the Data Quality Vocabulary (DQV)*, in *Semantic Web*, 12(1), 2021, 81-97.

[23] J. Hutton, P. Trueman, K. Facey, *Harmonization of evidence requirements for health technology assessment in reimbursement decision making*, in *International Journal of Technology Assessment in Health Care*, 24(4), 2008, 511-517.

[24] L. Cabrera, *Grounded normative theory and moral justification*, in *Journal of Global Ethics*, 16(1), 2020, 110-115.

[25] M. Johannesson, B. Jönsson, *Economic evaluation in health care: is there a role for cost-benefit analysis?*, in *Health Policy*, 17(1), 1991, 1-23.

[26] Quantum Consortium, *Deliverable 1.1*, cit. According to Hofmann, claiming that a system is "agnostic" can further diminish transparency by rendering its inherent ethical choices invisible.

[27] V.S.V. Pulla, C. Varol, M. Al, *Open source data quality tools: revisited*, in S. Latifi (ed.), *Information Technology: New Generations*, Cham, 2016, 893-902.

[28] World Health Organisation, *Health Technology Assessment of Medical Devices*, Geneva, 2011, https://www.who.int/publications/i/item/9789241501620 (last visited 22/06/2026).

nical scorecard into a robust normative filter. A dataset may be technically flawless – exhibiting high completeness and consistency – yet remain socially biased by overlooking vulnerable populations or rare pathologies. If the DQUL grants a high-quality status to such biased datasets without a normative filter, it risks institutionalising numerical discrimination under the guise of technical objectivity.[29]

This is where the success of DQUL lies: overcoming the heterogeneity and lack of representation that persist in existing datasets, ensuring the use of appropriate datasets with this standardised tag mechanism, and ensuring that quality datasets circulate in secondary-use areas and that accurate outcomes are produced.[30] However, a concrete question persists: can DQUL effectively address and overcome the heterogeneity, fragmentation, and difficulty of standardisation that characterise datasets concerning rare diseases, complex comorbidities, and socially vulnerable groups?[31] Such datasets are often smaller in scale, fragmented, heterogeneous, and difficult to standardise. When such hierarchically labelled evidence feeds into threshold-based decision systems, structural inequalities risk being reproduced through seemingly neutral numerical processes.[32]

3. Cost-Effectiveness and Threshold Logic

Cost-effectiveness analysis, particularly via QALYs and Incremental Cost-Effectiveness Ratios (ICERs), operates through threshold-based logic. Technologies are recommended if their cost per QALY falls below certain thresholds.[33] Cost-effectiveness analysis compares the costs and effectiveness of one policy with those of doing nothing and determines the average outcome. Health policies are often measured in terms of QALYs or disability-adjusted life years (DALYs).[34] While formally neutral, this structure has distributional consequences.[35] First, individuals with lower baseline health status generate smaller marginal QALY gains. Secondly, rare conditions often produce higher cost-per-patient ratios. Thirdly, heterogeneous populations introduce uncertainty that negatively affects modelling outputs.

QALYs combine life expectancy and quality-of-life utility to produce a single numerical value that compresses diverse lived experiences into standardised metrics.[36] This compression may systematically un-

[29] V. Charlton, M. Di Stefano, *op. cit.*

[30] S. Rousset, L. Charrier, M. Bersia *et al.*, *Enhancing representativeness in population-based surveys to improve data quality and decision-making*, in *Scientific Reports*, 15, 2025, 31605.

[31] N. Micic, D. Neagu, F. Campean, E.H. Zadeh, *Towards a Data Quality Framework for Heterogeneous Data*, in *IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom)*, Exeter, UK, 2017, 155-162.

[32] S. Kent, E. Burn, D. Dawoud *et al.*, *Common Problems, Common Data Model Solutions: Evidence Generation for Health Technology Assessment*, in *PharmacoEconomics*, 39, 2021, 275-285.

[33] C. Donaldson, G. Currie, C. Mitton, *Cost effectiveness analysis in health care: contraindications*, in *BMJ*, 325(7369), 2002, 891-894.

[34] A.C. Klimchak, L.E. Sedita, E.M. Perfetto, K.L. Gooch, D.C. Malone, *Discriminatory Properties of the Quality-Adjusted Life Year Based Cost-Effectiveness Analyses for Patients With Disabilities: A Duchenne Muscular Dystrophy Case Study*, in *Value in Health*, 27(12), 2024, 1641-1647.

[35] L.K. Newby, E.L. Eisenstein, R.M. Califf, T.D. Thompson, C.L. Nelson, E.D. Peterson *et al.*, *Cost-effectiveness of early discharge after uncomplicated acute myocardial infarction*, in *New England Journal of Medicine*, 342, 2000, 749-755.

[36] L. Prosser, J. Koplan, P. Neuman, M. Weinstein, *Barriers to using cost-effectiveness studies in managed care decision making*, in *American Journal of Managed Care*, 6, 2000, 173-179.

dervalue conditions that do not fit dominant modelling assumptions. Some authors[37] argue that QALY calculations, as determined by economists, can inherently lead to discrimination. This is because a score of 1 is assigned to "perfect health", while the health of people with disabilities, chronic illnesses, or older adults is always rated below 1.[38] As a result, their lives are implicitly considered less valuable, which can influence healthcare resource allocation in ways that systematically favour those with higher baseline health. This apparent formal equality (the premise that "a life is a life") is fundamentally undermined by the quality-of-life multiplier.[39]

As a standardised model development effort, Distributional Cost-Effectiveness Analysis (DCEA)[40] aims to minimise health inequalities by incorporating health equity considerations into the economic evaluation of healthcare interventions. The framework proceeds on the basis of social value judgments regarding health inequalities without presupposing any single normative stance. Crucially, however, the success of DCEA is closely linked to the normative filter that the DQUL must provide. If data quality labelling is confined to technical metrics such as mere completeness or accuracy, even advanced distributional models will inevitably be built on biased data.[41]

The principal challenge concerns the interpretation of inequalities in mortality and morbidity metrics. Empirical observations indicate that lower-income groups tend to have shorter life expectancy and higher levels of comorbidity compared to higher-income groups. As a result, when the same intervention is applied, the health gains achieved in disadvantaged populations are often smaller than those observed in more advantaged groups. In practical terms, the health gain (e.g., QALYs) may be lower, and the ICER may be higher, leading the intervention to appear less cost-effective in economic evaluations.[42] If access to new treatments increases among disadvantaged groups, existing utilisation data may fail to capture this change adequately; consequently, the intervention's inequality-reducing effects may be underestimated.[43]

Moreover, quality-of-life measurements can vary across individuals. A specific treatment may provide significant quality of life gains for one patient but not for another. However, such differences may not be adequately represented by numerical data, which can ultimately lead to unequal outcomes across groups.[44] In such contexts, discrimination is neither explicit nor intentional: it is structural. Because algorithms are presented as objective systems, their occurrence may go entirely unnoticed. All technologies

[37] J. HARRIS, *QALYfying the value of life*, in *Journal of Medical Ethics*, 13, 1987, 117-123; E. NORD, *The trade-off between severity of illness and treatment effect in cost-value analysis of health care*, in *Health Policy*, 24, 1993, 227-238.

[38] H.A. DAKIN, N.J. DEVLIN, I.A.O. ODEYEMI, *"Yes", "No" or "Yes, but"? Multinomial modelling of NICE decision-making*, in *Health Policy*, 77(3), 2006, 352-367.

[39] Z. OBERMEYER, B. POWERS, C. VOGELI, S. MULLAINATHAN, *Dissecting racial bias in an algorithm used to manage the health of populations*, in *Science*, 366(6464), 2019, 447-453.

[40] R. COOKSON, S. GRIFFIN, O.F. NORHEIM, A.J. CULYER, K. CHALKIDOU, *Distributional Cost-Effectiveness Analysis Comes of Age*, in *Value in Health*, 24(1), 2021, 118-120.

[41] *Ibid*.

[42] J. RICHARDSON, M. SCHLANDER, *Health technology assessment (HTA) and economic evaluation: efficiency or fairness first*, in *Journal of Market Access & Health Policy*, 7(1), 2019, 1557981.

[43] J.M. MUIR, A. RADHAKRISHNAN, I. OZER STILLMAN, G. SARRI, *Health Equity Considerations in Cost-Effectiveness Analysis: Insights from an Umbrella Review*, in *ClinicoEconomics and Outcomes Research*, 16, 2024, 581-596.

[44] A. ANGELIS, A. LANGE, P. KANAVOS, *Using health technology assessment to assess the value of new medicines: results of a systematic review and expert consultation across eight European countries*, in *European Journal of Health Economics*, 19(1), 2018, 123-152.

are evaluated under the same threshold. Yet, uniform thresholds applied to structurally unequal populations may produce indirect disadvantage. This is where numerical discrimination emerges.[45]

This risk may be further intensified when threshold-based appraisal systems operate alongside hierarchical data quality labelling mechanisms such as DQUL.[46] Although DQUL is formally agnostic to purpose, its graded structure implicitly orders datasets by technical robustness and usability. In practice, large-scale, standardised, and statistically coherent datasets are more likely to receive higher-quality labels, while datasets concerning rare diseases, socially vulnerable populations, or complex comorbidities may be positioned lower due to fragmentation, heterogeneity, or limited representativeness.[47] When HTA bodies accord greater epistemic weight to high-labelled datasets within cost-effectiveness modelling, a compounding effect emerges: structurally advantaged populations benefit from both more robust data representation and more favourable threshold outcomes.[48] Conversely, disadvantaged groups may be doubly penalised: first through poor data representation and second through threshold logic that converts poor measurable gains into lower priority.

The NICE technology appraisal programme is a particularly instructive case. Many countries use NICE HTA decisions while producing their own evidence in HTA processes.[49] Especially since 2019, with NICE planning its technology methods to be future-oriented and aligned with the latest technologies, the evidence produced by NICE's technology appraisal programme has continued to be widely adopted. In the same year, NICE adopted a trend toward lowering the threshold for the required accuracy of evidence in decisions related to rare diseases, paediatric evaluations, and emergencies.[50] However, even when the appraisal threshold is elevated or eased for rare diseases, the data quality hierarchy introduced by mechanisms such as the DQUL can render this flexibility functionally ineffective in practice. While a policy might allow for higher costs per QALY, the hierarchical labelling of the underlying data may still devalue the evidence base of a rare disease due to its inherent fragmentation, thereby obstructing the technology's path to approval regardless of threshold flexibility. The most important factor is adopting a QALY-based budget.[51]

In practice, rigid threshold values may fail to reflect real-world complexities and patient heterogeneity. As a result, decisions based solely on predefined numerical limits risk overlooking important contextual data and the specific needs of vulnerable populations. Some studies[52] fail to explain the heterogeneity in decisions made using QALYs and the increasingly lottery-like character of those decisions. Undoubtedly,

---

[45] S. Mohammed, L. Ehrlinger, H. Harmouch, F. Naumann, D. Srivastava, *The Five Facets of Data Quality Assessment*, in *SIGMOD Record*, 54(2), 2025, 18-27.

[46] Y. Sei, H. Okumura, A. Ohsuga, *Re-Identification in Differentially Private Incomplete Datasets*, in *IEEE Open Journal of the Computer Society*, 3, 2022, 62-72.

[47] R. Zhang, M. Indulska, S. Sadiq, *Discovering Data Quality Problems*, in *Business Information Systems Engineering*, 61, 2019, 575-593.

[48] E. Gustavsson, L. Lindblom, *Justification of principles for healthcare priority setting: the relevance and roles of empirical studies exploring public values*, in *Journal of Medical Ethics*, 2023.

[49] H.A. Dakin, N.J. Devlin, I.A.O. Odeyemi, *op. cit.*, 352-367.

[50] S. Howard, L. Harrison, *NICE guidance implementation tracking: data sources, methodology and results, report commissioned by NICE*, Bicester, 2004.

[51] N. Devlin, D. Parkin, *Does NICE have a cost-effectiveness threshold and what other factors influence its decisions? A binary choice analysis*, in *Health Economics*, 13(5), 2004, 437-452.

[52] H.A. Dakin, N.J. Devlin, I.A.O. Odeyemi, *op. cit.*; N. Devlin, D. Parkin, *op. cit.*

countries' budget constraints and inconsistencies in reimbursement decisions lead to inequities and injustices in patient access, even within the same geographical region. When the DQUL is evaluated alongside appraisal technologies such as NICE, it may lead to numerical discrimination, requiring particular attention to the structure of the algorithms used, especially when assessing new technologies.[53]

The Charlton and Di Stefano study[54] is particularly promising for understanding how algorithmic discrimination occurs. Their study is the first to empirically assess the coherence of NICE's formal approach and, in doing so, generates evidence-based conclusions about the extent to which this approach is morally justified.[55] While emphasising that it is not easy to determine the priorities of the entire population in health services, their study states that independently basing justice on concise principles will not eliminate injustice. The data feeding these systems may be processed in full compliance with legal requirements, within a framework nominally committed to social justice; yet the decisions the algorithm produces from that data cannot themselves be justified on social justice grounds: the normative foundation and the operational output are systematically misaligned.[56]

Societal value judgments may favour giving greater priority to younger populations or to individuals with rare diseases. However, embedding such priorities within threshold-based systems may implicitly deprioritise older patients. While societies may agree in principle on who deserves priority, they may be less consistent when determining the degree of that prioritisation, as this depends heavily on underlying normative sensitivities.[57]

When hierarchical data quality labelling mechanisms such as DQUL operate alongside fixed cost-effectiveness thresholds, the risk of numerical discrimination becomes structurally embedded rather than incidental. The issue does not arise from explicit exclusion or technical malfunction, but from the layered interaction between data hierarchies and evaluative thresholds.[58] As established in Section 2, datasets that are large-scale, standardised, and statistically robust are more likely to yield higher-quality labels and, consequently, exert greater epistemic authority within HTA modelling. When such hierarchically labelled evidence feeds into threshold-based decision systems, structural inequalities risk being reproduced through seemingly neutral numerical processes.

4. EU Anti-Discrimination Law

EU anti-discrimination law traditionally distinguishes between direct and indirect discrimination. Direct discrimination shall be taken to occur where one person is treated less favourably than another is, has been or would be treated in a comparable situation, on any of the grounds referred to in Article 1.[59] Indi-

[53] *Ibid*.
[54] V. CHARLTON, M. DISTEFANO, *op. cit.*
[55] A. RID, *Justice and procedure: how does "accountability for reasonableness" result in fair limit-setting decisions?*, in *Journal of Medical Ethics*, 35(1), 2009, 12-16.
[56] V. CHARLTON, M. DISTEFANO, *op. cit.*
[57] N. DANIELS, *Just Health Care*, New York, 1985.
[58] M. GOMES, A.J. TURNER, C. SAMMON, D. DAWOUD, S. RAMAGOPALAN, A. SIMPSON, U. SIEBERT, *Acceptability of Using Real-World Data to Estimate Relative Treatment Effects in Health Technology Assessments: Barriers and Future Steps*, in *Value in Health*, 27(5), 2024, 623-632.
[59] According to Article 2(2)(a) of Directive 2000/43/EC of the Council of 29 June 2000 implementing the principle of equal treatment between persons irrespective of racial or ethnic origin, OJ L 180, 19.7.2000, 22-26.

rect discrimination arises where an apparently neutral provision places certain groups at a disadvantage unless objectively justified.[60]

Before applying these concepts to DQUL-based HTA systems, it is necessary to identify which EU anti-discrimination instruments are potentially relevant and to acknowledge the genuine doctrinal uncertainty attending their application to healthcare resource allocation. EU anti-discrimination law operates principally through two secondary instruments.[61] Directive 2000/43/EC (the Racial Equality Directive) prohibits discrimination on grounds of racial or ethnic origin across a materially broad scope that extends, beyond employment, to social protection, social security, social advantages, education, and access to goods and services available to the public – a scope that, in principle, reaches national healthcare systems insofar as they constitute social protection or social advantages. Directive 2000/78/EC (the Employment Equality Directive) prohibits discrimination on grounds of religion, belief, disability, age, and sexual orientation, but is primarily limited to employment and occupation.[62] Article 21 of the EU Charter of Fundamental Rights contains an open-ended list of prohibited grounds (including disability, age, and social origin) that goes beyond the directive's numerus clausus of enumerated characteristics. The Charter applies to EU institutions and Member States acting within the scope of EU law; its horizontal application to purely national healthcare resource decisions remains contested.

This structural uncertainty requires transparent argumentation. The most defensible doctrinal position is that insofar as HTA decisions implement EU instruments (including the EHDS, the HTA Regulation (EU) 2021/2282, or national law giving domestic effect to EU equality directives) the Charter's equality guarantee is directly engaged.[63] The directives' grounds – particularly disability and age – are directly relevant to the populations this article identifies as vulnerable to numerical discrimination. Where HTA decisions fall outside EU implementing contexts, challenges must operate through national constitutional equality guarantees and the ECtHR framework discussed below.

The case law of the European Court of Human Rights and various EU directives define the prohibition of discrimination from the perspective of differential treatment arising from comparable situations. The pro-

[60] S. CHOUDHRY, *Distribution vs. Recognition: The Case of Anti-Discrimination Laws*, 2000. For foundational EU anti-discrimination theory, see S. FREDMAN, *Discrimination Law*, 3rd ed., Oxford, 2022; M. BELL, *The concept of indirect discrimination*, in *Research Handbook on European Anti-Discrimination Law*, Cheltenham, 2025.

[61] Directive 2000/43/EC, cit., extends its material scope beyond employment to social protection, social security, social advantages, education, and access to goods and services, including housing (Article 3(1)(e)–(h)). Directive 2000/78/EC of the Council of 27 November 2000 establishing a general framework for equal treatment in employment and occupation, OJ L 303, 2.12.2000, 16-22, is primarily limited to employment and occupation (Article 3(1)). Article 21 of the Charter of Fundamental Rights of the European Union, OJ C 364, 18.12.2000, 1, contains a non-exhaustive open-ended list of prohibited grounds extending beyond the directives' enumerated categories, including social origin and generic discrimination. On the relationship between Article 21 Charter and the directives, and the contested horizontal direct effect of the Charter in private law, see M.A. WÓJCIK, Algorithmic Discrimination in Health Care: An EU Law Perspective, Health and Human Rights Journal, 24(1), 2022, 93-103.

[62] For the EU equality law framework see: S. Fredman, *op. cit.*; M. Bell, *op. cit.*; M.A. WÓJCIK, *op. cit.*, 93-103; F.J. ZUIDERVEEN BORGESIUS, *Discrimination, Artificial Intelligence and Algorithmic Decision-Making*, Council of Europe, Strasbourg, 2018. On the general theoretical framework of big data and disparate impact, US scholarship provides parallel (though doctrinally distinct) insights: S. BAROCAS, *A.D. Selbst, Big Data's Disparate Impact*, in *California Law Review*, 104, 2016, 671-732; the US categories of disparate treatment/impact do not map directly onto EU indirect discrimination doctrine and should be read with that caveat.

[63] D. HELLMAN, *When is Discrimination Wrong?*, Cambridge MA, 2008, 31, 33, 35.

nounced divergence between EU anti-discrimination law, the ECtHR's jurisprudence, and the principles enshrined in the EU Charter of Fundamental Rights is noteworthy; it exposes the limitations of the'CJEU's exhaustive list model in the face of contemporary algorithmic discrimination. In particular, the Court's effort to create space for new categories within a model of exhaustive lists that cannot be judicially expanded (and its failure to define AI-driven and algorithmic discrimination as autonomous categories) represents a structural gap that secondary legislation has yet to address.

The Test-Achats ruling (Case C-236/09) serves as a crucial example. In its ruling, the Court noted that the right to equal treatment between men and women is enshrined in the TFEU, and examined how the utilisation of statistical data to solidify attributes regarding women affects premiums paid by both men and women.[64] Since DQUL-based HTA systems operate inherently through algorithmic decision-making, the traditional elements of direct and indirect discrimination under EU anti-discrimination law have become analytically strained. As these systems are defined as systematic, statistical, and opaque, the defence that they lack discriminatory intent will inevitably arise. Viewing individuals strictly as individuals (and ensuring that their unique characteristics within groups are not disregarded) would reduce statistical uncertainty. However, this is not sufficient; reducing statistical uncertainty at an individual level may infringe upon the individual's right to self-determination.

The foundational proportionality framework for indirect discrimination in CJEU case law was established in Bilka-Kaufhaus GmbH v. Weber von Hartz (Case 170/84).[65] The Court held that an apparently neutral measure producing a disproportionate disadvantage to members of a protected group is permissible only where it corresponds to a real need of the undertaking or authority concerned, is appropriate to achieving the objective pursued, and is necessary to that end. This tripartite test – real need, appropriateness, and necessity – has since been applied across employment, pay, and social protection contexts and constitutes the operational content of proportionality in EU equality law. Applied to DQUL-based HTA, the framework asks: (1) does cost-effectiveness optimisation through QALY thresholds correspond to a real need? (2) is the uniform application of rigid thresholds to structurally unequal populations appropriate? (3) are such thresholds necessary, or would representativeness-weighted alternatives or distributional adjustments achieve the same aim with less discriminatory impact? The article's core argument is that where DQUL evidence reveals material representational gaps, a rigid QALY threshold cannot easily satisfy the necessity criterion. In Seymour-Smith (Case C-167/97), the Court confirmed that statistical evidence of differential impact alone is sufficient to establish a prima facie case of indirect discrimination, shifting the burden of justification to the respondent.[66] This is directly significant for the DQUL framework: where data quality

[64] E. GOESSENS, G. DE BAERE, *Gender Differentiation In Insurance Contracts After The Judgment In Case C-236/09, Association Belge Des Consommateurs Test-Achats ASBL v. Conseil Des Ministres*, in *Columbia Journal of European Law*, 2012, 339.

[65] Court of Justice, Case 170/84, *Bilka-Kaufhaus GmbH v. Karin Weber von Hartz*, ECLI:EU:C:1986:204, ECR 1607, paragraph 36. The Bilka test requires that a measure producing a disproportionate disadvantage to a protected group must: (i) correspond to a real need of the undertaking or authority concerned; (ii) be appropriate to achieving the objective pursued; and (iii) be necessary to that end. This tripartite proportionality framework remains the cornerstone of the objective justification analysis under EU equality law. See also Court of Justice, Case C-127/07, *Arcelor Atlantique et Lorraine and Others*, ECLI:EU:C:2008:728, paragraph 47, applying comparable proportionality reasoning in a complex regulatory setting.

[66] Court of Justice, Case C-167/97, *Regina v. Secretary of State for Employment, ex parte Nicole Seymour-Smith and Laura Perez*, ECLI:EU:C:1999:60, ECR I-623, paragraph 60. The Court confirmed that statistical data showing that a

labels systematically reflect lower scores for datasets representing disabled or elderly populations due to structural fragmentation rather than genuine quality deficiency, the statistical evidence of that representational gap could constitute the prima facie showing required to engage the indirect discrimination framework. Additionally, in Coleman v. Attridge Law (Case C-303/06), the Court affirmed the influence of the UN Convention on the Rights of Persons with Disabilities (CRPD) on EU disability law, reinforcing the positive obligation of individualised accommodation that sits in tension with the group-averaging logic of QALY-based appraisal.[67]

As stated in Article 5(1) of the Gender Equality Directive,[68] the implementation of the rule requiring gender-neutral premiums and benefits clearly stipulates that consideration of gender as an actuarial factor must not lead to discrepancies in individual premiums and benefits. Such statistical discrimination constitutes a breach of the principle of equal treatment, and an objective justification is fundamentally lacking where the very grounds upon which such justification is based are themselves discriminatory.[69]

The European Court of Human Rights has developed a complementary body of jurisprudence on structural inequality and positive obligations that reinforces the doctrinal arguments outlined above. In Thlimmenos v. Greece, the Court held that Article 14 ECHR is violated not only when States treat differently persons in analogous situations without objective justification, but also when they treat similarly persons in relevantly different situations.[70] This positive dimension of the non-discrimination obligation (requiring differentiated treatment where situations differ in relevant respects) is directly engaged by DQUL-based HTA systems that apply uniform thresholds to populations in fundamentally unequal epistemic and health-status positions. In D.H. and Others v. Czech Republic (Grand Chamber), the Court accepted that statistical evidence of differential exclusion suffices to establish a presumption of indirect discrimination, confirming that a general policy having prejudicial effects on a group may constitute indirect discrimination even where it is not specifically aimed at that group.[71] This acceptance of aggregate statistical data as indirect

considerably smaller percentage of women than men is able to satisfy a requirement constitutes evidence of apparent indirect discrimination. The statistical threshold for prima facie discrimination need not be absolute; relative statistical disadvantage is sufficient to shift the burden of justification.

[67] Court of Justice, Case C-303/06, *S. Coleman v. Attridge Law and Steve Law*, ECLI:EU:C:2008:415, ECR I-5603. The Court affirmed that the prohibition of disability discrimination under Directive 2000/78/EC is not limited to persons who are themselves disabled, extending the scope of protection and confirming the influence of the UN Convention on the Rights of Persons with Disabilities (CRPD), ratified by the EU, on EU disability discrimination law. On the CRPD's influence on EU equality law, see further M. Bell, *op. cit.*

[68] As stipulated in Article 5(1) of the Gender Equality in Goods and Services Directive (2004/113/EC) of the Council of 13 December 2004 implementing the principle of equal treatment between men and women in the access to and supply of goods and services, OJ L 373, 21.12.2004, 37-43.

[69] A. Moro, *Statistical Discrimination*, cit., 12996-13001.

[70] European Court of Human Rights, *Thlimmenos v. Greece*, Application no. 34369/97, Judgment of 6 April 2000. The Court held, at paragraph 44, that the right to enjoy Convention rights without discrimination is violated when States, without an objective and reasonable justification, fail to treat differently persons whose situations are significantly different. This positive dimension of the equality obligation is directly engaged by uniform threshold-based systems applied to populations in structurally unequal positions.

[71] European Court of Human Rights (Grand Chamber), D.H. and *Others v. Czech Republic*, Application no. 57325/00, Judgment of 13 November 2007. The Grand Chamber confirmed, at paragraphs 177-195, that statistical evidence of a difference in treatment between two groups is sufficient to raise a presumption of indirect discrimination, shifting the burden to the respondent State to provide an objective and reasonable justification. The Court accepted that structural and societal disadvantage, rather than intentional targeting, can ground an indirect discrimination finding.

discrimination evidence aligns precisely with the article's proposal that DQUL representational data be cognizable as evidence of structural inequality before equality bodies and courts. In Nitecki v. Poland, the Court addressed healthcare resource allocation decisions under Article 2 ECHR, confirming that while States retain a significant margin of appreciation in healthcare resource management, that margin is not unlimited and must be exercised consistently with the Convention's fundamental rights framework.[72]

While EU law classifies algorithmic bias under the rubric of indirect discrimination, this categorisation often grants deployers an unwarranted objective justification defence. By framing proxy-based discrimination as circumvention of the prohibition of direct discrimination, this article argues for a stricter scrutiny level, effectively denying the availability of such justifications for technically engineered biased outcomes. At this early stage, the fundamental problem is how we can apply the theory of direct discrimination when we are unable to prove discriminatory intent.[73] While relying on the theory of indirect discrimination in AI-based systems appears to be the best available doctrinal framework, the level of objective justification will remain contested for some time. In this process, the court may bypass the discussion of discrimination altogether and limit the scope to proving how the AI programme operates, effectively accepting the presumption that discrimination is inherent in the nature of AI.[74]

The proportionality analysis outlined above raises a deeper constitutional question that the article must address: to what extent can courts or equality bodies supervise socio-economic policy choices, such as HTA thresholds, without substituting judicial for legislative or administrative judgment? This question of institutional competence determines the practical viability of the doctrinal proposals advanced here. The CJEU's approach to proportionality in discrimination cases involves calibrated scrutiny that does not require courts to optimise policy but does require that discriminatory impacts be demonstrably justified by legitimate aims pursued through appropriate and necessary means.[75] In complex socio-economic contexts (where resource allocation, budgetary constraints, and technical expertise are implicated), the Court has traditionally exercised institutional restraint, granting national decision-makers a margin of appreciation.[76] As Baker has analysed, this deference follows a graduated model: courts are more restrained in complex socio-economic policy areas but will intervene where discriminatory impacts are severe, systematic, or inadequately reasoned. However, deference is not absolute. In cases where the discriminatory impact affects access to essential healthcare services, the Court's deference contracts. The constitutional

[72] European Court of Human Rights, *Nitecki v. Poland*, Application no. 65653/01, Decision of 21 March 2002. The Court addressed state obligations under Article 2 ECHR in relation to healthcare resource allocation, confirming that the margin of appreciation enjoyed by States in health policy does not immunize decisions from fundamental rights scrutiny, particularly where resource allocation choices produce life-threatening inequalities.

[73] C. NARDOCCI, *Artificial Intelligence-based Discrimination: Theoretical and Normative Responses. Perspectives from Europe*, in *DPCE online*, 60(3), 2023.

[74] Tribunal of Bologna, 31 December 2020, with comment by D. TESTA, *La discriminazione degli algoritmi: il caso Deliveroo, Trib. Bologna, 31 dicembre 2020*, in *IusinItinere.it*, 26 January 2021.

[75] Court of Justice, Case C-127/07, *Arcelor Atlantique et Lorraine and Others v. Premier Ministre*, ECLI:EU:C:2008:728, ECR I-9895, paragraph 23.

[76] M. BAKER, *Proportionality, the CJEU, and Discrimination*, in *ERA Forum*, 2013, available at http://www.era-comm.eu/oldoku/Adiskri/14_Other%20topics/2013_Baker_EN.pdf (last visited 22/06/2026). Baker's analysis shows that proportionality review in the CJEU follows a graduated deference model: courts are more restrained in complex socio-economic policy areas but will nonetheless intervene where discriminatory impacts are severe, manifestly disproportionate, or inadequately reasoned. This model directly informs the analysis of HTA threshold review in the text.

dimension of this article's proposal is therefore best understood not as a maximalist claim to judicial supremacy over HTA methodology, but as a procedure-focused argument: where DQUL evidence reveals material representational gaps, the burden shifts to HTA authorities to demonstrate through substantive justification – not merely technical opacity – that their threshold design satisfies the Bilka criteria. Courts do not thereby become HTA designers; they require that the designers render their normative choices legible and proportionate.

If we are to accept that certain things are inherently discriminatory, those things should be neither algorithms nor economic foundations alone. When data is employed as a proxy, the objective is to measure, control, and manage resources; yet the compression of human complexity into numerical metrics raises profound normative questions about the delegation of representational authority in health governance.[77] DQUL-based HTA models prioritise technical quality, yet they inherently risk marginalising certain protected groups. This exclusionary effect is further exacerbated by the rigid application of cost-effectiveness analysis; such economic frameworks are fundamentally driven by resource maximisation, often at the expense of equitable treatment.[78]

Is the QALY concept simultaneously tasked with ensuring non-discrimination? Because when resources are limited, QALYs inherently prioritise the health needs of certain groups over others, potentially leading to the neglect of other groups' needs. This process can reinforce social inequalities by limiting access to healthcare for the most vulnerable groups, who may benefit less from health interventions. For instance, certain demographic groups, such as the elderly or low-income populations, may have less access to healthcare because opportunity-cost calculations in health economics may restrict funding for technologies aimed at these groups. The Court has held that the absence of clearly articulated criteria for authorisation renders the exercise of discretionary power arbitrary and therefore incompatible with EU law.[79] This principle is directly applicable to DQUL-based HTA models: when these systems operate as "black boxes" lacking transparent, publicly known criteria for their outcomes, they mirror the arbitrary discretionary power that the Court rejected. By cloaking resource allocation decisions in technical methodologies that evade ex-ante transparency, such models effectively circumvent the requirement for objective justification.[80]

Threshold-based HTA models resemble neutral provisions. The comparator problem, however, becomes complex. The disadvantaged group is not explicitly excluded, but is statistically less likely to meet the cost-effectiveness criteria.[81] When principles such as non-discrimination, balanced treatment processes, fair distribution of treatment, life-saving requirements, and the prioritisation of life-threatening conditions are considered together, would positive discrimination for individuals with a higher chance of survival or a longer life expectancy be regarded as a violation of anti-discrimination provisions?[82] A key limitation

[77] A. ROMEI, S. RUGGIERI, *A multidisciplinary survey on discrimination analysis*, in *The Knowledge Engineering Review*, 29(5), 2014, 582-638.

[78] M. BULTERMAN, *European Court of Justice*, in *European Journal of Health Law*, 13(2), 2006, 173-181.

[79] M. PEETERS, *Free Movement of Patients: Directive 2011/24 on the Application of Patients' Rights in Cross-Border Healthcare*, in *European Journal of Health Law*, 19(1), 2012, 29-60.

[80] L.Z. RAND, A.S. KESSELHEIM, *op. cit.*

[81] D. HELLMAN, *op. cit.*

[82] L. PUGH, *Disability and Discrimination in Triage Frameworks: A Commentary on Mathematical Approaches to Reducing Discrimination*, in *Johns Hopkins University School of Medicine*, 2021.

emerges: EU equality law is primarily designed to address identifiable grounds and relational comparisons; numerical discrimination may arise from aggregation effects, statistical modelling assumptions, or the epistemic hierarchy of data. These mechanisms operate beyond traditional comparator-based frameworks.[83] Consequently, structural disadvantage may remain legally invisible unless explicitly connected to protected characteristics.

This reveals a structural limitation of the comparator-based framework itself. Where disadvantage results not from explicit differential treatment but from statistical aggregation, modelling assumptions, or threshold design, identifying a suitable comparator becomes conceptually strained. In such cases, inequality does not arise between two clearly distinguishable groups, but is embedded in the architecture of the decision-making model. The traditional relational logic of equality law struggles to capture harms that are systemic rather than comparative. To bridge this gap, DQUL acts as a "bridge of visibility", transforming abstract statistical disadvantage into legally cognisable evidence of indirect discrimination.[84]

In this respect, a DQUL could serve as a structural transparency mechanism, making visible the distributional consequences of data-driven health maximisation models. By disclosing representational gaps, modelling assumptions, and group-specific reliability levels, such a label would shift the focus from identifying a comparator to scrutinising the model's epistemic foundations. Rather than creating new substantive equality rights, DQUL would enhance the conditions under which indirect discrimination and proportionality assessments can be meaningfully conducted. In doing so, it would reduce the risk that structural disadvantage remains legally invisible simply because it is statistically mediated.[85]

The determination of the details of a seemingly benevolent concept such as health maximisation inevitably raises the question of which model of equality is being adopted, by whom, and on what normative basis. Whatever equality model is selected, it appears to generate divergent outcomes across HTA bodies.[86] In cases of indirect discrimination, relevant legislation seeks to prevent discrimination or, at the very least, mitigate its effects. There must be an objective and reasonable justification for the differential impact.[87] Over time, the Court of Justice's case law has evolved, and the proportionality test has been adopted.[88]

Very clearly, if a group is underrepresented relative to another group, or is deprived of a benefit, this constitutes discrimination. Therefore, within the framework on which HTA decisions are based, a disadvantaged group that generates lower cost-effectiveness and has a higher ICER in threshold evaluations can be considered a case of indirect discrimination. In other words, numerical discrimination is an unintentional form of indirect discrimination and requires objective justification.[89]

---

[83] Court of Justice, Case C-236/09, *Association Belge des Consommateurs Test-Achats and Others*, ECLI:EU:C:2011:100.

[84] J. KWON, H. SQUIRES, T. YOUNG, *Efficiency and equity of community-based falls prevention pathways: a model-based health economic evaluation*, in *Age and Ageing*, 54(8), 2025, afaf212.

[85] *Ibid*.

[86] G.W. TORRANCE, W.H. THOMAS, D.L. SACKETT, *A utility maximization model for evaluation of health care programs*, in *Health Services Research*, 7(2), 1972, 118-133.

[87] M. BELL, *op. cit.*

[88] European Union Agency for Fundamental Rights, Council of Europe, *Handbook on European Non-Discrimination Law*, 2018 ed., Luxembourg, 2018. On proportionality in discrimination law see further M. BAKER, *op. cit*.

[89] M.A. WÓJCIK, *op. cit.*, 93-103.

The central point is that if indirect discrimination can be objectively and legitimately justified, it cannot be considered as constituting discrimination in the first place. According to the Court of Justice, assessing discrimination requires not rigid numerical or statistical criteria, but a substantive approach that takes into account the values and circumstances of different societies: context and impact are considered more important than strict reliance on quantitative data.[90] Court decisions become highly relevant in shaping the dynamics of HTA decision-making.[91]

Returning to the notion of objective and legitimate justification: justification can be both helpful and risky. It allows organisations to pursue legitimate goals because the treatment is not overtly discriminatory. At the same time, if a practice unintentionally harms certain groups and is not clearly linked to a valid and proportionate aim, the organisation may still face claims of indirect discrimination.[92]

For example, as previously discussed, would the prioritisation of young and rare-disease patients on the basis of societal sensitivity criteria be regarded as objective and neutral if justified under provisions on positive discrimination?[93] When positive discrimination is justified by the three proportionality criteria, such measures may be considered reasonable. The test evaluates the stages of legitimacy, effectiveness, and proportionality (or necessity) separately. However, the use of this test as a justificatory framework remains contested, and it is unclear whether it is more advantageous than alternative methods within cost-effectiveness analysis.[94]

Given these complexities, it is necessary to pursue a more refined approach by developing a model explicitly grounded in the principle of non-discrimination. Within this framework, the adoption of "equity weights" or "lenient thresholds" for disadvantaged groups ceases to be a mere efficiency adjustment; rather, it becomes a direct operationalisation of substantive equality. By ensuring that the HTA model accounts for pre-existing structural disadvantages identified through the DQUL, this approach bridges the gap between technical assessment and social justice. It is precisely within this integrated context that an objective and acceptable justification for differential treatment can be grounded.

5. GDPR

Health data processed under the EHDS inevitably intersect with personal data protected by the GDPR, particularly under Article 9. On the one hand, the principle of data minimisation enshrined in Article 5(1)(c) GDPR requires that personal data be adequate, relevant, and limited to what is necessary for the specified purpose. On the other hand, ensuring algorithmic fairness and representativeness in data-driven systems

[90] European Union Agency for Fundamental Rights, *Inequalities and multiple discrimination in access to and quality of healthcare*, Luxembourg, 2013.

[91] K. LIPPERT-RASMUSSEN, *Cost-Effectiveness and the Avoidance of Discrimination in Healthcare: Can We Have Both?*, in *Cambridge Quarterly of Healthcare Ethics*, 32(2), 2023, 202-215.

[92] J.A. LANE, R. INGLEBY, Ind*irect Discrimination, Justification and Proportionality: Are UK Claimants at a Disadvantage?*, in *Industrial Law Journal*, 47(4), 2018, 531-552.

[93] S. Fredman, *Emerging from the Shadows: Substantive Equality and Article 14 of the European Convention on Human Rights*, in *Human Rights Law Review*, 16, 2016, 273.

[94] *Ibid*.

often requires the use of broad, diverse datasets.[95] This structural tension becomes particularly visible in HTA processes that rely on large-scale data analytics.[96]

The GDPR does not provide an explicit definition of algorithmic discrimination, nor does it directly regulate intentional or unintentional forms of discrimination produced by automated systems. Consequently, whether a specific group is deliberately disadvantaged or unintentionally excluded by algorithmic design, the GDPR lacks a comprehensive, explicit anti-discrimination framework tailored to algorithmic decision-making.[97] Nevertheless, Article 9 GDPR establishes strict conditions for the processing of special categories of data, including health data, and significantly limits their use in the context of automated decision-making. Article 22 reinforces this protection by granting data subjects the right not to be subject to decisions based solely on automated processing that produce legal or similarly significant effects.[98]

In this context, the requirement to provide "meaningful information about the logic involved" becomes crucial. It enables the data subject to understand and interpret the consequences of automated decisions taken about them.[99] As established in relation to the NICE HTA model, consistent numerical methodologies may produce systematic forms of numerical discrimination. Because such systems operate in an internally coherent and technically consistent manner, affected individuals – especially patients – may neither understand the structural impact of these decisions nor meaningfully challenge them. One cannot contest what one cannot comprehend. The danger lies precisely in the invisibility of discrimination embedded within consistent algorithmic frameworks.[100]

This risk is closely connected to the use of health data. It cannot be guaranteed that datasets used in HTA processes are entirely free of personal data. Even where secondary-use datasets are formally anonymised, the risk of re-identification persists. For this reason, GDPR compliance remains a structural and ongoing responsibility within HTA processes.[101] The crucial point is that Article 22 GDPR, read in conjunction with Recital 71, grants the data subject specific safeguards in the context of automated decision-making. These safeguards include the right to contest the decision, to obtain meaningful information about its consequences, and to request human intervention.[102]

Accordingly, datasets used to generate evidence in HTA processes must be assessed not only for their potential to contain personal data but also for the risk of re-identification. Where such risks exist, Data

---

[95] M.E. KAMINSKI, G. MALGIERI, *Algorithmic impact assessments under the GDPR: producing multi-layered explanations*, in *International Data Privacy Law*, 11(2), 2021, 125-144.

[96] O. LYNSKEY, *Complete and Effective Data Protection*, in *Current Legal Problems*, 76(1), 2023, 297; N. NI LOIDEAIN, *EU Data Privacy Law and Serious Crime*, Oxford, 2025, 428-430.

[97] *Ibid*.

[98] D. BALDINI, *Article 22 GDPR and Prohibition of Discrimination. An Outdated Provision?*, 2019, available at SSRN: https://ssrn.com/abstract=5399709 (last visited 22/06/2026).

[99] A.D. SELBST, J. POWLES, *Meaningful information and the right to explanation*, in *International Data Privacy Law*, 7(4), 2017, 233-242.

[100] J. GRIMMELMANN, D. WESTREICH, *Incomprehensible Discrimination*, in *California Law Review Online*, 7, 2017, 164, 173.

[101] R.A. ALSHAIKH, K.A. WALSH, S. SPILLANE, P. HARRINGTON, M. O'NEILL, C. TELJEUR, M. RYAN, C.M. O'DRISCOLL, *The Early Experiences of Health Technology Assessment Bodies in the Implementation of the European Union Health Technology Assessment Regulation for High-Risk Medical Devices: A Qualitative Study*, in *Value in Health*, 29(6), 2026, 1034-1044.

[102] R. BECKER, E.S. DOVE, *The EU GDPR and secondary use of health and genetic data for research support purposes*, in *International Data Privacy Law*, 2026, ipag001.

Protection Impact Assessments (DPIAs) become mandatory. This obligation is particularly relevant in light of the EHDS framework, where certain datasets will be made available through the DQUL mechanism. Since DQUL-labelled datasets will be accompanied by metadata and catalogue standards and will be accessible to data users for secondary purposes, their compatibility with Article 9 GDPR must be carefully examined.[103]

Moreover, as DQUL-supported infrastructures will inevitably rely on machine learning and algorithmic systems, GDPR compliance with respect to automated decision-making becomes unavoidable. Although the GDPR imposes the principle of data minimisation, these systems often process data in ways that prioritise health optimisation rather than strict data reduction. This tension underscores the importance of secure processing environments (SPEs), which operate in direct relation to Articles 5 and 9 of the GDPR. An SPE must guarantee adherence to core data protection principles and to the specific conditions governing the processing of health data.[104]

In this regard, the Health Data Access Body (HDAB) bears responsibility for verifying GDPR compliance when evaluating access requests to the SPE. It must ensure that applicants meet data protection requirements and retain the authority to revoke access or notify the competent supervisory authority in cases of non-compliance.[105]

Nevertheless, a structural paradox remains. Achieving high levels of data quality, utility, and accurate labelling (particularly for discrimination monitoring) may require extensive and diverse datasets, which appear to conflict with the principle of data minimisation. One of the central objectives of the EHDS is to enable the secondary use of health data while fully respecting GDPR requirements. However, especially in the context of discrimination monitoring and algorithmic fairness, the demand for comprehensive data may sit uneasily alongside strict data minimisation obligations.

Another major concern from a GDPR perspective is algorithmic transparency. DQUL-based HTA decision-making processes remain normatively indeterminate, particularly in relation to the criteria that other HTA systems traditionally take into account. The integration of automated systems into these processes may generate a range of unintended and unforeseen outcomes.[106] As previously noted, this opacity can prevent data subjects from understanding instances of numerical discrimination and, consequently, from effectively exercising their data protection rights. In HTA contexts, overly narrow datasets may produce biased outputs; yet expansive data collection may conflict with data minimisation principles. A dataset may be legally compliant yet statistically unbalanced. Thus, GDPR compliance is necessary, but not sufficient to prevent structural inequality.[107]

Data minimisation has long been regarded as almost incompatible with algorithmic decision-making, especially when such decision-making requires large amounts of data to be fair and non-discriminatory. The

[103] P. QUINN, E. ELLYNE, C. YAO, *Will the GDPR restrain health data access bodies under the European Health Data Space (EHDS)?*, in *Computer Law & Security Review*, 54, 2024, 105993.
[104] *Ibid*.
[105] T. ZARSKY, *Transparent Predictions*, in *University of Illinois Law Review*, 2013, available at SSRN: https://papers.ssrn.com/sol3/Papers.cfm?abstract_id=2324240 (last visited 22/06/2026).
[106] S. WACHTER, B. MITTELSTADT, L. FLORIDI, *Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation*, in *International Data Privacy Law*, 2017.
[107] G. SARTOR, F. LAGIOIA, *The impact of the General Data Protection Regulation (GDPR) on artificial intelligence*, European Parliament, Brussels, 2020.

question is whether we can interpret the principle of data minimisation, as provided under the GDPR, in the context of AI-based technology appraisal programmes or HTA mechanisms.[108]

For data minimisation to be compatible with the production of non-discriminatory decisions and the prevention of health inequities, the algorithm must be supplied with data that maximises population representativeness. While it is sometimes argued that using demographically structured data is essential for the development of health technologies, it is equally important to test how such data are reflected in the algorithmic model.[109] Data are inherently biased, and this inherent bias can only be mitigated through human intervention.

When automated decision-making processes operate in conjunction with an HTA system structured around DQUL-based technology assessments, the risk of discrimination remains present from a GDPR perspective.[110] This is primarily because the GDPR contains structural gaps regarding a fully articulated right to explanation and the establishment of a genuinely accountable framework for automated decision-making. Creating an accountable system would require full transparency as to how HTA processes function together with DQUL-based standards, how decisions are reached from beginning to end, and in what ways these processes may generate forms of numerical discrimination. It would also require developing a pre-risk identification mechanism capable of detecting and addressing discriminatory patterns before they materialise into concrete decisions.[111]

Such a pre-risk system should be based on the systematic informing and warning of all stakeholders involved, coupled with the effective recognition of the right to object and the possibility of meaningful human intervention. The continuing risk lies in the fact that data subjects may remain unaware that numerical discrimination has occurred in the evidence-based phase of HTA decisions, often unintentionally.

6. AI Act

AI-based decision-making systems have the potential to – and often do – lead to discrimination, posing an inherent risk in almost every context. Barocas and Selbst identify several stages where this discrimination can occur: (i) labelling the training data; (ii) collecting the training data; (iii) feature selection; and (iv) the use of proxies.[112] Each of these stages is directly relevant to DQUL-based HTA: the labelling of health datasets (stage i) is the core function of the DQUL mechanism; the collection of training data (stage ii) is shaped by the representativeness requirements of the EHDS framework; feature selection (stage iii) determines which health characteristics receive statistical weight in cost-effectiveness models; and proxy use (stage iv) is addressed by the indirect discrimination doctrine developed in Section 4. This four-stage

[108] *Ibid*.

[109] R. AGARWAL, M. BJARNADOTTIR, L. RHUE, M. DUGAS, K. CROWLEY, J. CLARK, G. GAO, *Addressing algorithmic bias and the perpetuation of health inequities: An AI bias aware framework*, in *Health Policy and Technology*, 12(1), 2023, 100702.

[110] K. LIPPERT-RASMUSSEN, *Cost-Effectiveness and the Avoidance of Discrimination in Healthcare: Can We Have Both?*, cit.

[111] B. CUSTERS, A.-S. HEIJNE, *The right of access in automated decision-making: The scope of article 15(1)(h) GDPR in theory and practice*, in *Computer Law & Security Review*, 46, 2022, 105727.

[112] F.J. ZUIDERVEEN BORGESIUS, *Discrimination, artificial intelligence and algorithmic decision-making*, Council of Europe, Strasbourg, 2018. https://rm.coe.int/discrimination-artificial-intelligence-and-algorithmic-decision-making/1680925d73.

framework thus provides a structural diagnostic for identifying where numerical discrimination can enter the HTA process.

The nature of discrimination in AI systems differs substantially from the discrimination practised by humans.[113] If we could always trace the root causes of discrimination perpetrated by AI-based systems, we could claim that these systems make perfectly transparent decisions. However, the situation is far more complex, particularly regarding accountability, transparency, and explainability. Perhaps the most significant issue this research seeks to uncover is that individuals cannot effectively challenge a system whose underlying logic they neither understand nor comprehend. For this reason, when AI-based models are used in HTA decision support, they may qualify as high-risk systems under the AI Act (AIA), particularly when they affect access to essential services.[114] Building on this, the fact that the data used is biased does not fully explain how AI leads to discrimination.[115] Article 10 of the AIA requires training data to be relevant, representative, and free of errors to the extent possible. However, representativeness is evaluated relative to the intended purpose, the regulation does not mandate distributive correction mechanisms, and there is no requirement to recalibrate cost-effectiveness thresholds to protect vulnerable groups.[116] Therefore, while the AIA addresses bias at the data level, it does not address the normative architecture of threshold-based valuation. The regulation ensures procedural robustness, not distributive justice.[117] At this stage, the system resorts to the use of proxies – a critical factor often overlooked in general discourse but highlighted in significant studies. As previously defined, proxy discrimination functions as an optimisation strategy or correlation pattern within AI; particularly in threshold-based systems, these systems may seem to uphold fundamental rights on the surface while systematically rendering these rights ineffective. By choosing proxies that correlate with protected attributes, the system circumvents or disregards necessary legal safeguards.[118] When population data that underrepresent certain groups are used in AI model training, the risk of discrimination may be exacerbated, reinforcing the exclusion of specific populations from representation.[119]

Within the HTA decision-making process, where robust evidence is required, access to the datasets used to train AI systems may be requested. It is essential that the datasets employed for technological assessment comply with the requirements of the AI Act.[120] Datasets used for AI training must be relevant to the system's intended purpose, sufficiently representative, and as complete and error-free as possible, with

[113] C. Nardocci, *op. cit.*

[114] R. Di Bidino, S. Daugbjerg, S.C. Papavero, I.H. Haraldsen, A. Cicchetti, D. Sacchini, *Health technology assessment framework for artificial intelligence-based technologies*, in *International Journal of Technology Assessment in Health Care*, 40(1), 2024, e61.

[115] C. Nardocci, *op. cit.*

[116] A. Buscemi *et al.*, *Assessing High-Risk AI Systems under the EU AI Act: From Legal Requirements to Technical Verification*, arXiv preprint arXiv:2512.13907, 2025.

[117] L. Grozdanovski, J. De Cooman, *Individual safeguards in the era of AI: Fair algorithms, fair regulation, fair procedures*, in *Cambridge Forum on AI: Law and Governance*, 1, 2025, e18.

[118] S. Barocas, A.D. Selbst, *Big data disparate impact*, in *California Law Review*, 104, 2016, 671.

[119] S.H.A. Muller *et al.*, *Towards responsible surveillance in preventive health data-AI research*, in *PLOS Digital Health*, 4(12), 2025, e0001146.

[120] R. Di Bidino *et al.*, *op. cit.*

appropriate statistical characteristics. The EHDS will address these deficiencies.[121] The EHDS will provide transparent, reliable, and privacy-preserving access to data for the validation, training, and testing of AI systems, without discrimination. The need for quality data to ensure the protection of fundamental rights in AI is essential for both producing consistent, reliable outputs and maintaining privacy and transparency.[122] At this point, it is important to note that the lack of representativeness in datasets will further complicate the integration of AI in HTA's decision-making processes. Therefore, the importance of DQUL-based HTA will be reinforced, as DQUL meets the requirements under EHDS Art. 78, thereby ensuring that the outputs fully represent the population.[123]

The AI Act identifies certain stand-alone systems as high-risk, as listed in Annex III, which includes inter alia "access and enjoyment of essential private and public services". This category explicitly covers systems that determine eligibility for public assistance benefits and services, as well as those that allocate emergency services, such as medical aid. Consequently, algorithms deployed to assess healthcare benefits would likely fall within this high-risk category and attract a corresponding level of regulatory protection.[124]

While acknowledging bias in AI models has become common, the inability to systematically detect these biases remains a critical challenge. For high-risk AI systems, ensuring data quality is not merely a preference but a regulatory necessity under Article 10 of the AIA. In this context, DQUL-based systems serve as a pivotal component of the data governance process, facilitating access to high-quality datasets and ensuring that AI operates within its intended scope.[125]

One of the fundamental shifts in the industry is the transition from traditional HTA to AI-Enhanced HTA. This evolution moves beyond static algorithmic foundations toward dynamic, representative tools designed to include all demographic groups.[126] By integrating high-quality data through DQUL frameworks, AI significantly amplifies HTA's evidence-generation potential. This synergy optimises cost-effectiveness, a core HTA component, directly improving healthcare outcomes for the elderly and vulnerable populations.[127] Although the AIA does not specifically regulate the decision-making mechanism of HTA, it is no longer difficult to anticipate that AI systems, by their nature, can lead to discrimination. The AIA acknowledges that, due to their inherent nature, AI systems may cause discriminatory outcomes.[128]

---

[121] D. Kalra, *AI Needs High-Quality Health Data at Scale – Will the EHDS Deliver?*, in *Journal of Applied Interdisciplinary Research*, Special Issue, 2025, 116-118.

[122] M. Berglind, O. Kjelkvik, *AI Processing of Special Categories of Personal Data in the EU Healthcare Sector – From a GDPR and EHDS Perspective*, 2025.

[123] Y.S. Tolias, *Unlocking Health Data for AI: Analyzing the European Commission's EHDS Proposal*, in *Journal du Droit de la Santé et de l'Assurance-Maladie (JDSAM)*, 38(3), 2023, 34-38.

[124] D. Golpayegani, H.J. Pandit, D. Lewis, *To be high-risk, or not to be – semantic specifications and implications of the AI Act's high-risk AI applications and harmonised standards*, in *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 2023.

[125] M. Guillen-Aguinaga *et al.*, *Data quality in the age of AI: A review of governance, ethics, and the FAIR principles*, in *Data*, 10(12), 2025, 201.

[126] M.R. Nikolopoulou, V. Nikolova, *HTA75 Utilization of Artificial Intelligence By Health Technology Assessment Agencies to Improve Efficiency of Health Technology Appraisals*, in *Value in Health*, 27(6), 2024, S258.

[127] K. Bowrin, J.B. Briere, P. Levy, A. Millier, E. Clay, M. Toumi, *Cost-effectiveness analyses using real-world data: an overview of the literature*, in *Journal of Medical Economics*, 22(6), 2019, 545 – 553.

[128] L. Grozdanovski, *Non-discrimination law, the GDPR, the AI Act and the – now withdrawn – AI liability directive proposal offering gateways to pre-trial knowledge of algorithmic discrimination*, in *AI and Ethics*, 5(5), 2025, 5039-5062.

This complexity is further compounded by the unclear classification of certain administrative tools. For example, an algorithm designed to identify patients likely to miss appointments may not directly assess eligibility for benefits, appearing merely as an operational tool to optimise scheduling.[129] Yet, such systems can still have significant discriminatory effects if certain groups are systematically flagged. This highlights the importance of considering not only the intended purpose of AI systems but also their potential societal impacts when determining whether AI systems are high-risk under the AI Act.[130]

Article 10 of the AIA specifies that datasets shall, to the extent required by the intended purpose, take into account the characteristics or elements particular to the geographical, contextual, behavioural, or functional setting in which the high-risk AI system is intended to be used. These features are not meant to be considered in a disjunctive manner; rather, the dataset's composition should reflect the weight of all four criteria.[131]

The significance of using actuarial prediction methods to forecast an algorithmic model's future outputs has also been highlighted in the literature. When a dataset, labelled as high-quality under DQUL schemes, interacts with technology appraisal programmes or HTA mechanisms, the use of actuarial prediction becomes essential. This ensures that health technologies reach patients without causing discrimination or inequity. NICE technology appraisal should adopt an actuarial predictive system capable of dynamically assessing evaluations across the full spectrum of health states.[132]

Article 60 of the AIA addresses the testing of high-risk AI systems, requiring providers to produce more forward-looking predictions to evaluate the AI's output. In this context, providers must test AI models not only in simulated trial environments but also in real-world conditions. It is crucial to emphasise that real-world testing must not violate the numerus clausus principle outlined in Article 5 of the AIA.[133]

One of the principles that must be assessed is social scoring. Specifically, algorithmic systems used in the healthcare sector and AI-based health technology appraisal programmes are not intended to implement social scoring.[134] The risk lies in these programmes deviating from their original, socially beneficial purpose, potentially leading to unintended discrimination and inequality. As the model changes, the required dataset, risk factors, and, consequently, the algorithm's ability to adequately define the population may also change.[135]

The issue of social scoring is particularly relevant when datasets, labelled for data quality and utility, are used in health technology appraisal programmes. In some cases, each segment of the population is as-

[129] R. ASIF, *Equality and Inclusivity through Accountability: Regulating Automated Decision-Making in the Public Sector*, in *QMLJ*, 2025, 93.

[130] M. BUYL *et al.*, *Tackling algorithmic disability discrimination in the hiring process: An ethical, legal and technical analysis, in Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 2022.

[131] R. DI BIDINO *et al.*, *op. cit.*

[132] N. TIMMINS, M. RAWLINS, J. APPLEBY, *A terrible beauty: A short history of NICE the National Institute for Health and Care Excellence*, in *F1000Research*, 6, 2017, 915.

[133] P.T. KRAMCSAK, *Can legitimate interest be an appropriate lawful basis for processing Artificial Intelligence training datasets?*, in *Computer Law & Security Review*, 48, 2023, 105765. Note: the numerus clausus principle in Article 5 AIA refers to the exhaustive and closed character of the prohibitions listed therein.

[134] L. COMEL, *Scoring Systems in the General Data Protection Regulation and Artificial Intelligence Act: A Critical Appraisal*, in *MCEL Master's Thesis Series*, 3, 2025.

[135] K. MEDING, *It's complicated. The relationship of algorithmic fairness and non-discrimination provisions for high-risk systems in the EU AI Act*, in *arXiv preprint arXiv:2501.12962*, 2025.

signed a score representing perfect health or poor health. Unfortunately, some populations – especially marginalised groups, vulnerable individuals, or people with disabilities – are often not adequately represented in these datasets, leading to unintentional discrimination.[136] Therefore, it is crucial that the dataset reflects the population's diversity to avoid these issues.

Article 5 of the AIA (Recital 31) highlights that when determining the relative importance of one health technology over another, or evaluating the health status of a population for which the technology is required, social ranking could result in individuals being scored. Consequently, within programmes such as the NICE technology appraisal, individuals who fall below a certain threshold have already been scored. Therefore, if a treatment or resource considered necessary for achieving adequate health does not reach vulnerable, disabled, and marginalised individuals, it constitutes discrimination and results in health inequalities.[137]

Article 27 of the AI Act serves as a significant regulation for scenarios where identifying the intent behind decisions is impossible. Article 27 appears to offer a solution to the problem of objective justification as well. Combined with the DPIA/FRIA approach, the fundamental rights assessment has been adopted as a mandatory requirement; AI systems that pose a risk to fundamental rights must be designed to evaluate, prevent, or mitigate such risks. This assessment constitutes an expert opinion and does not waive the requirement for a DPIA.[138] Within Article 27, paragraphs (c), (e), and (f) are particularly significant for the detection, prevention, or mitigation of discrimination. In any case, it is essential for AI systems to identify the affected group or number of individuals resulting from the system's use for specific risks, establish an effective complaint mechanism, and define the circumstances under which human oversight is triggered.[139] The scope of the FRIA can only be fully determined by the providers after the completion of the conformity assessment procedure under Article 9. The risk factor involved here will determine the effectiveness of the FRIA assessment design.[140]

The legal framework governing the human rights dimensions of AI in healthcare is not limited to the EU regulatory sphere. The Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225), opened for signature on 5 September 2024, constitutes the first binding international treaty on artificial intelligence.[141] The Convention adopts a technol-

---

[136] S. PROPERZI *et al.*, *Exploring methods to assess environmental health inequalities in health impact assessments of local interventions: a systematic review within the JA PreventNCD project*, in *Frontiers in Public Health*, 13, 2025, 1546394.

[137] K. SHRISHAK, S. PÉNICAUD, *European policy guidelines on AI and algorithm-driven discrimination for equality bodies and other national human rights structures*, European Union Agency for Fundamental Rights, 2025.

[138] Artificial Intelligence Act, Article 27(4). See also DONET, *The Fundamental Rights Impact Assessment By Deployers Of Artificial Intelligence Systems In The Regulation (EU) 2024/1689*, 2025, 550.

[139] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (AI Act), OJ L, 12.7.2024, Article 27.

[140] M. MERLER, *Artificial intelligence in law enforcement: a critical analysis of the EU AI Act's framework for protecting public values and fundamental rights*, 2026.

[141] Council of Europe, *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law*, CETS No. 225, opened for signature 5 September 2024, Vilnius. The Convention entered into force upon five ratifications. The European Union signed the Convention on 11 September 2024: see European Commission, The Commission signed the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, https://digital-strategy.ec.europa.eu/en/news/commission-signed-council-europe-framework-convention-artificial-intelligence-and-human-rights (last visited 22/06/2026).

ogy-neutral, human rights-based approach, imposing obligations on Parties – including EU Member States and the EU itself as a signatory – to ensure that activities throughout the lifecycle of AI systems respect human rights, the rule of law, and democratic principles.

Several provisions of the Convention are of particular relevance to the argument advanced in this article. Article 10 requires Parties to adopt measures to identify and address risks to equality and non-discrimination arising from AI systems, expressly covering indirect discrimination effects – precisely the form of disadvantage that DQUL-based HTA systems produce. This obligation is formulated as a continuous lifecycle obligation rather than a one-time compliance check. Article 11 establishes a general framework for risk and impact assessment that, unlike the AI Act's conformity assessment, is explicitly grounded in human rights obligations. Article 14 requires Parties to ensure that persons adversely affected by AI system outputs have access to effective remedies.

These obligations complement the AI Act's FRIA mechanism while operating at a higher normative level. Where the AI Act's Article 27 FRIA focuses on technical documentation of risks to EU Charter rights, the CoE Convention imposes substantive obligations to identify, prevent, and mitigate discrimination risks as a continuous governance responsibility. Applied to DQUL-based HTA systems that involve algorithmic processing, the Convention requires signatory States to conduct human rights impact assessments specifically addressing the risk of structural disadvantage arising from data representational gaps and threshold-based evaluation architectures.

To support the implementation of the Convention's risk management obligations, the Council of Europe's Committee on Artificial Intelligence has developed the HUDERIA (Human Rights Due Diligence for AI Risk Assessment) methodology, which was adopted in November 2024.[142] HUDERIA provides a structured framework for identifying, assessing, and mitigating human rights risks across the AI lifecycle, with particular attention to equality and non-discrimination. Its structured approach to identifying affected groups, assessing differential impacts, and establishing monitoring obligations aligns closely with the institutional proposals advanced in Section 7. DQUL-based HTA systems involving algorithmic processing would, under a HUDERIA-informed assessment, be required to evaluate the representational adequacy of their underlying datasets against the rights of persons with disabilities, elderly populations, and socioeconomically marginalised groups, *i.e.*, precisely the populations identified in this article as vulnerable to numerical discrimination. The CoE Convention and HUDERIA thus situate this article's doctrinal proposals within an emerging multi-level governance architecture linking EU regulatory compliance to binding international human rights obligations.[143]

[142] Council of Europe Committee on Artificial Intelligence (CAI), *HUDERIA – Human Rights Due Diligence for AI Risk Assessment Methodology*, adopted November 2024, available at https://data.consilium.europa.eu/doc/document/ST-16537-2025-INIT/en/pdf (last visited 22/06/2026). HUDERIA is a non-binding but authoritative methodology designed to support Parties to the CoE Framework Convention in conducting systematic human rights due diligence assessments of AI systems throughout their lifecycle.

[143] On the complementary relationship between EU AI governance and CoE obligations, and the significance of the CoE Convention for EU Member States' human rights duties regarding AI, see further the European Parliament resolution of 11 March 2026 on the Council of Europe Framework Convention on AI: European Parliament, PV 10-2026-03-11-ITM-008-03 (last visited 22/06/2026).

7. Conclusion

The EHDS, GDPR, and AI Act collectively establish important safeguards for transparency, data governance, and technical robustness. However, as the preceding analysis has demonstrated, they do not eliminate the risk of numerical discrimination emerging from the interaction between hierarchical data labelling and cost-effectiveness thresholds. Structural disadvantage may arise not from regulatory failure, but from the architectural incompatibility between formally neutral legal frameworks and quantitative decision-making systems that systematically undervalue statistically marginal populations. Numerical precision cannot substitute for substantive equality. The following proposals aim to address this gap at each level of the regulatory architecture examined in this article.

At the outset, the analysis of the DQUL framework in Section 2 revealed that its current architecture functions primarily as a technical scorecard. A dataset may achieve a high-quality label while remaining socially biased – statistically complete yet representationally hollow. The proposal is therefore that the DQUL be redesigned to incorporate an explicit representativeness dimension as a mandatory, not optional, labelling criterion. Representativeness should be assessed not merely relative to the general population but specifically with respect to groups that are structurally underrepresented in health data ecosystems, such as persons with disabilities, rare disease patients, elderly populations, and socioeconomically marginalised groups. A dataset that scores highly on completeness and accuracy but demonstrably lacks representation of these groups should not receive an unrestricted high-quality label. At a minimum, the label should make such gaps visible through a mandatory disclosure requirement, so that data users in HTA contexts are informed of the population scope of the evidence they rely on.

This redesign would require the EHDS implementing acts and the delegated acts specifying DQUL criteria under Article 78 to define representativeness standards operationally, rather than leaving them implied by the general quality indicators. The HDABs should bear responsibility for auditing representativeness at the point of labelling, with revocation authority exercised not only for technical inaccuracy but also for undisclosed representational bias.

Secondly, the analysis of cost-effectiveness threshold logic in Section 3 demonstrated that the QALY framework, when applied uniformly to structurally unequal populations, systematically converts pre-existing disadvantage into lower priority. This is not an HTA malfunction; it is an architectural feature of maximisation-based appraisal. No data governance reform can fully neutralise it without corresponding changes to how HTA bodies interpret and apply their thresholds. The proposal is therefore the mandatory introduction of distributive impact assessments as a procedural requirement within HTA decision-making at the EU level, building on the Joint Clinical Assessment framework established by the HTA Regulation (EU) 2021/2282. A distributive impact assessment would require HTA bodies to produce an explicit account of how a technology's cost-effectiveness profile varies across population subgroups, particularly those likely to generate lower QALY gains due to baseline health disadvantage. This is not equivalent to abandoning cost-effectiveness as a criterion; it is a requirement of transparency about who bears the cost of threshold rigidity.

Thirdly, Section 4 identified that EU anti-discrimination law struggles to capture numerical discrimination because it is organised around relational comparisons between identifiable groups, while the disadvantage under examination is embedded in the architecture of the decision-making model itself. The law is not deficient in ambition but in instruments. The proposal operates at the level of judicial and adminis-

trative interpretation. Where indirect discrimination is alleged in the context of algorithmic HTA systems, the proportionality analysis (specifically, the necessity and suitability stages) should be conducted with explicit reference to DQUL evidence. Where a DQUL reveals material representational gaps in the datasets underlying a cost-effectiveness determination, the use of a rigid QALY threshold as objective justification for differential impact should be treated as legally fragile rather than presumptively valid. This doctrinal argument can be adopted by national courts, the CJEU, and equality bodies within the existing framework, drawing on the proportionality principles established in Bilka-Kaufhaus and confirmed in Seymour-Smith, applied in the healthcare resource allocation context through the positive obligation framework established in Thlimmenos and the statistical evidence standard confirmed in D.H. v. Czech Republic.

Additionally, the concept of structural comparators should be developed in EU equality doctrine. Where disadvantage results not from differential treatment of two identifiable groups but from a threshold design applied to heterogeneous populations, the comparator should be the population as it would be represented in a fully inclusive dataset, rather than a specific advantaged group. This shift in comparator logic would allow equality law to reach the architectural level at which numerical discrimination actually operates.

Furthermore, Section 5 highlights that achieving the representativeness necessary to prevent numerical discrimination requires broad, diverse datasets, whereas the GDPR's data minimisation principle limits data collection to what is strictly necessary. The proposal is that the EDPB issue guidance specifically addressing the application of data minimisation in the context of algorithmic fairness and health equity. Such guidance should recognise that, where a processing purpose includes the prevention of discriminatory outcomes in AI-supported decision-making, the necessity assessment under Article 5(1)(c) GDPR must account for the representativeness requirements of that purpose. The minimisation principle should be read purposively rather than mechanically in this context. Separately, the mandatory DPIA requirement should be explicitly extended to cover datasets used in HTA processes under the EHDS, even where those datasets are formally anonymised.

Moreover, Section 6 showed that Article 10 of the AI Act addresses data quality and bias at the level of training data, but does not reach the normative architecture of the valuation system into which that data feeds. An AI system trained on representative, high-quality datasets may still produce numerically discriminatory outputs if the cost-effectiveness threshold it operationalises is itself structurally biased. The proposal is that, for AI systems classified as high-risk under Annex III that are deployed in HTA or reimbursement decision-support contexts, the conformity assessment obligations should include a mandatory evaluation of threshold design. The fundamental rights impact assessment required under Article 27 of the AI Act is the appropriate vehicle for this purpose and should be interpreted accordingly by market surveillance authorities and notified bodies. This proposal is reinforced by the obligations imposed on EU Member States and the EU itself under the CoE Framework Convention on AI (CETS 225), which requires continuous human rights risk management (including assessment of indirect discrimination risks) throughout the AI lifecycle. The HUDERIA methodology provides a practical implementation framework for these assessments that HTA authorities could integrate into their DQUL and AI-based appraisal procedures.

Finally, the deeper proposal advanced by this article is institutional in nature. The DQUL should be formally recognised as the connective mechanism linking data governance, anti-discrimination law, and AI oversight within the health technology domain. This means three things in practice. First, DQUL findings re-

garding representational gaps should be admissible as evidence in indirect discrimination proceedings before equality bodies and courts. Secondly, DQUL assessments should be a mandatory input into the fundamental rights impact assessments required of high-risk AI systems used in HTA. Thirdly, the distributive impact assessments proposed above should be required to reference DQUL data as part of their evidentiary basis. In this configuration, the DQUL ceases to be a technical label and becomes a legal instrument of structural transparency: the point at which the invisible is made visible, and the formally compliant is made accountable.