



Fairness Score

D5.3 BOOKLET



Funded by
the European Union

Table of Contents

Why Measuring Fairness Is So Difficult	4
What Do We Mean by "Fairness"?	5
From Principles to Outcomes	6
Why a Score? Learning from the Nutri-Score Model	7
What the Fairness Score Is (and Is Not)	8
How the Fairness Score Works	10
How the Fairness Score is calculated	11
What the Music Data Dashboard is (and is not)	12
The Five Dimensions of Fairness	13
From Indicators to a Final Score	14
Parametric Fairness and User Agency	15
Fairness as a Parametric Choice	16
What the Fairness Score Reveals	17
Limits, Risks and Safeguards	18
Why Safeguards are Essential	19
How the Fairness Score Can Be Used	20
Making Fairness Legible	21

INTRODUCTION

Why Measuring Fairness Is So Difficult

Fairness has become one of the most frequently used words in discussions about digital platforms. It appears in debates on algorithms, remuneration, visibility and cultural diversity. Yet despite its central role, fairness often remains difficult to define, and even harder to assess in practice.

In the music sector, this difficulty is particularly evident. Platforms shape how music is discovered, promoted and consumed, but many of the mechanisms behind these processes remain opaque. Creators and users are often asked to trust systems they cannot fully observe or understand. Fair MusE approaches this challenge from a practical perspective. Rather than attempting to provide a universal definition of fairness, the project asks a more operational question: how can fairness be made visible, discussable and comparable in real-world platform environments?

The Fairness Score was developed as a response to this question. It is not intended to resolve all disagreements about what fairness should mean. Instead, it aims to translate complex values into a structured framework that allows different aspects of platform behaviour to be examined side by side.

What Do We Mean by "Fairness"?

Fairness is not a single, universally agreed concept. Different actors in the music ecosystem often associate it with different expectations and experiences.

For creators, fairness may relate to transparency, remuneration or visibility. For users, it may concern choice, diversity or meaningful control over recommendations. For policymakers, fairness is often linked to accountability, market balance and compliance with legal objectives.

These perspectives are not necessarily incompatible, but they do not always align. This plurality makes fairness difficult to operationalise. Any

attempt to measure or assess it must acknowledge that it involves trade-offs between values that cannot always be maximised at the same time.

For this reason, Fair MusE treats fairness as a multi-dimensional concept. Rather than asking whether a platform is fair or unfair in absolute terms, the project focuses on identifying distinct dimensions along which fairness can be discussed, compared and evaluated.

This approach does not eliminate disagreement. Instead, it creates a shared language that allows different viewpoints to be articulated more clearly.

Fairness is not one thing.
It is a combination of perspectives.

From Principles to Outcomes

Many discussions about platform fairness concentrate on internal processes. They examine how algorithms are designed, how decisions are taken, or whether rules are applied consistently.

While these aspects are important, they often remain inaccessible to external observers. Platforms rarely disclose detailed information about their internal systems, and even when they do, technical transparency does not necessarily translate into understandable outcomes.

Fair MusE therefore adopts a different perspective. Instead of focusing exclusively on internal processes, the Fairness Score emphasises outcomes. It asks how platform behaviour affects visibility, choice, diversity and accountability as experienced by users and creators.

This outcome-oriented approach reflects a simple principle: fairness ultimately matters because of its effects. What users and creators encounter in practice is often more relevant than how a system is engineered internally.

By focusing on observable consequences rather than hidden mechanisms, the Fairness Score seeks to bridge the gap between abstract principles and lived experience. It provides a way to discuss fairness in terms that are meaningful beyond technical or institutional settings.

Why a Score?

Learning from the Nutri-Score Model

One of the main challenges in addressing platform fairness is communication. Fairness involves complex trade-offs between economic, cultural and social values, yet these debates often take place far from everyday user experience.

For this reason, Fair MusE drew inspiration from existing public-interest tools designed to make complex information more understandable. One of the most influential examples is the Nutri-Score system used in several European countries to communicate nutritional quality through a simple letter-and-colour scale.

Nutri-Score does not claim to describe food exhaustively, nor to replace expert knowledge. Instead, it translates multiple indicators into a synthetic signal that supports awareness and comparison. Its value lies not in precision alone, but in its ability to make invisible attributes legible to non-specialists.

D5.3 adopts this logic deliberately. The Fairness Score follows a similar principle: it simplifies without denying complexity, and it communicates patterns rather than absolute truths.

Research on algorithmic transparency and public-interest labelling shows that interpretability matters as much as access to information. When users and creators are confronted with large volumes of data or purely technical disclosures, transparency can become ineffective. A score provides an entry point, allowing different dimensions to be seen together and discussed openly.

At the same time, the deliverable is explicit about the risks of scoring systems. As highlighted in the literature on indicators and governance, metrics can oversimplify reality, encourage gaming, or create a false sense of objectivity. These risks are not ignored. They are addressed through methodological safeguards that are integral to the Fairness Score design.

What the Fairness Score Is (and Is Not)

The Fairness Score was developed as a structured interpretative tool. Its purpose is to support reflection and comparison, not to deliver binding judgements.

The Score is designed to:

- make different dimensions of platform fairness visible at the same time;
- translate qualitative assessments into a transparent scoring logic;
- allow comparison across platforms or over time;
- support informed discussion among users, creators and policymakers;

These functions reflect the core objective of D5.3: transforming abstract fairness principles into something that can be examined, questioned and debated.

At the same time, the deliverable clearly defines what the Fairness Score does not aim to do. It is not a certification system, a legal assessment, or a compliance mechanism. It does not replace regulatory enforcement, nor does it claim to measure fairness exhaustively.

Importantly, the Score does not produce a single, unquestionable truth. Its results depend on clearly stated assumptions, indicator choices and weighting decisions. Making these assumptions visible is part of the methodology.

This distinction is crucial. By presenting fairness as something that can be structured without being absolutised, the Fairness Score avoids turning complex social values into rigid benchmarks. Instead, it functions as a shared interpretative framework.

**The Fairness Score informs
debate. It does not settle it.**

How the Fairness Score Works

The Fairness Score is built on the idea that fairness cannot be captured through a single indicator. Instead, it emerges from the interaction of multiple dimensions that reflect different expectations placed on digital platforms.


D5.3 adopts a structured evaluation model inspired by research on multi-criteria decision-making and algorithmic accountability. Rather than asking whether a platform is fair in general, the model assesses how it performs across a set of predefined dimensions, each linked to observable practices or outcomes.

Each dimension is evaluated through specific indicators. These indicators translate qualitative characteristics, such as transparency or user agency, into structured assessments that can be compared across platforms or scenarios.

The deliverable emphasises that this translation is not automatic. Indicators require interpretative judgement, supported by available evidence. For this reason, the Fairness Score does not rely exclusively on quantitative data. It combines qualitative evaluation with structured scoring in order to remain meaningful even where precise measurements are unavailable.

Once indicators are assessed, their results are aggregated through a transparent scoring logic. This aggregation allows different dimensions to be considered together, while still preserving visibility of individual scores. The intention is not to hide complexity, but to organise it.

Throughout this process, the methodology remains explicit about its assumptions. Indicator definitions, scoring thresholds and aggregation choices are all documented. This openness is essential to avoid presenting the Score as a neutral or purely objective measure.



FOCUS 1

How the Fairness Score is calculated

The Fairness Score is built on a structured evaluation framework composed of **16 individual indicators**, grouped into five fairness dimensions.

Each indicator can receive one of three values:

- **-1** when practices are misaligned with fairness principles;
- **0** when practices are neutral or partially aligned
- **+1** when practices meet high fairness standards

The total score therefore ranges from **-16 to +16**. To make results easy to understand, the numerical score is translated into a **five-letter scale**, inspired by public-interest scoring systems such as Nutri-Score:

- **A – Fully aligned**
- **B – Substantially aligned**
- **C – Partially aligned**
- **D – Minimally aligned**
- **E – Unaligned**

This translation follows a **non-linear scaling logic**, designed to encourage platforms to improve practices across multiple dimensions rather than optimising only a few indicators.



What the Music Data Dashboard is (and is not)

The Fairness Score is structured around a set of core dimensions identified through literature review, stakeholder discussions and conceptual analysis within the Fair MusE project.

Each dimension captures a different aspect of platform behaviour that affects how fairness is experienced by users and creators. Together, they provide a multi-faceted view that reflects the complexity of digital music environments.

D5.3 explicitly acknowledges that remuneration cannot be assessed directly without access to confidential financial data. The focus therefore remains on enabling conditions rather than monetary outcomes.

USER AGENCY

User agency refers to the extent to which individuals can meaningfully influence their platform experience. This includes control over recommendations, the ability to adjust parameters, and the possibility to explore content beyond default settings.

This dimension reflects a key innovation of the Fair MusE approach: fairness is not only about protecting users, but also about empowering them.

FOCUS 2

The Five Dimensions of Fairness

The Fairness Score evaluates music streaming services across five complementary dimensions:

Data practices	Examines how user data and metadata are structured, accessible and interoperable, including GDPR exports and metadata standards.
Power balance	Assesses whether users and creators have meaningful control over recommendations, data use and algorithmic choices.
Remuneration	Evaluates transparency in revenue distribution, reinvestment practices and copyright protection mechanisms.
Representation	Measures diversity outcomes within playlists and recommendations, including gender balance, local content and exposure variety.
Transparency	Considers whether platforms explain how recommendations work and disclose key information about catalogue composition. Together, these dimensions reflect the structural conditions that shape fairness in the digital music ecosystem.

From Indicators to a Final Score

Once indicators are assessed, their results are aggregated to produce partial scores for each fairness dimension. These partial scores are then combined to form an overall Fairness Score.

D5.3 adopts a letter-based scale, typically ranging from **A to E**, inspired by communication tools such as the Nutri-Score. The purpose of this scale is not to rank platforms competitively, but to provide an immediately understandable overview of performance across fairness dimensions.

The aggregation process is designed to be transparent. Weighting choices are made explicit, and the methodology avoids hidden calculations. This openness is essential to maintain trust and to allow informed debate about how different aspects of fairness should be prioritised.

The deliverable highlights that different weighting configurations can lead to different outcomes. This is not a flaw, but a feature. By making weighting visible, the Fairness Score invites discussion about values rather than masking them behind technical complexity.

Numerical aggregation therefore serves a communicative function. It summarises multidimensional assessments while preserving the possibility of looking beneath the surface at individual dimensions and indicators.



Parametric Fairness and User Agency

One of the distinctive features of the Fairness Score is its emphasis on **parametric fairness**. Rather than treating fairness as a fixed standard defined once and for all, the model recognises that different users and creators may prioritise different values.

In traditional platform governance models, fairness is usually embedded within system design and applied uniformly. Users have little visibility into how decisions are made, and even less ability to influence them. Fair MusE challenges this approach by introducing the idea that fairness can be explored through adjustable parameters.

Parametric fairness allows individuals to see how different priorities affect outcomes. For example, increasing the weight assigned to diversity may produce a different overall score than prioritising transparency or user control. This does not mean that all configurations are equally valid, but it makes underlying value choices visible.

D5.3 draws on research in human-centred AI and participatory governance, which emphasises that meaningful agency requires more than consent. It requires the ability to understand and influence system behaviour.

By enabling users to engage with fairness dimensions rather than simply accepting predefined rankings, the Fairness Score shifts the focus from passive evaluation to active interpretation. Fairness becomes something that can be discussed, tested and questioned, rather than imposed invisibly.

This approach also responds to concerns raised in the literature about algorithmic paternalism. Instead of assuming that platforms or regulators alone should decide what fairness means, the Fairness Score opens space for plural perspectives.

Fairness as a Parametric Choice

A central innovation of the Fairness Score is its **parametric approach**.

Rather than imposing a single definition of fairness, the system is designed to allow users to **adjust how recommendations are generated**, according to their own values and preferences.

In a possible interface, users could choose, for example:

- whether to activate or deactivate algorithmic recommendations;
- how much diversity or novelty they want in playlists;
- the balance between local and international content;
- the role of popularity versus emerging artists;
- whether AI-generated music should be included.

This approach reinforces **user agency**, recognising that fairness is not a fixed rule but a negotiated outcome shaped by individual choices.

What the Fairness Score Reveals

The Fairness Score demonstrates that platform fairness can be examined in a structured and transparent way, even in complex digital environments.

By breaking fairness down into dimensions and indicators, the model reveals where strengths and weaknesses lie. It shows that platforms may perform well in some areas while falling short in others, and that fairness cannot be reduced to a single attribute.

The Score also highlights the importance of visibility. Where information is accessible, assessment becomes possible. Where information is missing, uncertainty persists. In this sense, the Fairness Score acts as a diagnostic tool that points directly to gaps in transparency.

Another key insight concerns comparability. When fairness dimensions are assessed using consistent criteria, it becomes possible to compare practices across platforms or over time. This does not produce absolute judgements, but it supports evidence-informed discussion.

Overall, the Score makes fairness legible. It transforms abstract concerns into structured questions that can be examined publicly.

Fairness becomes meaningful when it can be discussed, compared and questioned.

Limits, Risks and Safeguards

D5.3 is explicit in acknowledging that scoring systems carry inherent risks. Simplification can lead to misunderstanding, misinterpretation or strategic manipulation.

One risk is reductionism. A single letter score may conceal important differences between dimensions. For this reason, the Fairness Score is always presented together with its underlying breakdown, ensuring that the overall result can be contextualised.

Another concern is fair-washing. Platforms could selectively highlight favourable scores while downplaying weaknesses. To counter this risk,

the methodology emphasises transparency of indicators, weighting and assumptions.

The deliverable also recognises the potential influence of Goodhart's law: when a measure becomes a target, it may distort behaviour. This reinforces the importance of treating the Fairness Score as an interpretative tool rather than a compliance benchmark.

Safeguards are therefore built into the design. These include openness about methodology, avoidance of hidden calculations, and explicit recognition of uncertainty where data is lacking.

The Fairness Score is a guide, not a verdict.

Why Safeguards are Essential

The Fairness Score explicitly recognises the risks associated with scoring systems.

A single letter score may oversimplify complex realities, encourage strategic behaviour, or be selectively used for communication purposes.

To address these risks, the methodology includes several safeguards:

- the score is always presented together with its underlying breakdown;
- assumptions and criteria are publicly documented;
- uncertainty is acknowledged when data is missing;
- the Score is framed as an interpretative tool, not a compliance label.

In this way, the Fairness Score supports informed discussion rather than replacing judgment.

How the Fairness Score Can Be Used

The Fairness Score is intended to support multiple forms of engagement across the music ecosystem.

For users and creators, it can provide an accessible way to understand how platforms differ in their approach to transparency, agency and accountability.

For policymakers and regulators, it offers a structured lens through which platform practices can be examined, compared and discussed, complementing legal and regulatory tools.

For researchers and civil society organisations, the Score provides a flexible framework that can be adapted, refined or extended as new evidence becomes available.

For streaming services, the Score improves their recommendations systems and helps them to be more transparent regarding users and creators.

Importantly, **the Fairness Score is not designed to replace regulation or enforcement**. Its value lies in supporting informed debate and highlighting where deeper scrutiny may be needed.

CONCLUSIONS

Making Fairness Legible

Fairness in digital music platforms is not a fixed property. It is shaped by design choices, governance arrangements and information flows.

The Fairness Score developed within Fair MusE does not claim to resolve these challenges. Instead, it provides a way to make fairness visible, discussable and comparable.

By translating complex values into a structured framework, the Score bridges the gap between abstract principles and lived experience. It shows that fairness can be examined without pretending to eliminate disagreement.

Ultimately, the Fairness Score contributes to a broader objective of the Fair MusE project: strengthening transparency, accountability and trust in the digital music ecosystem. Making fairness legible is a necessary step towards making it actionable.



Fair MusE research and industry partners





Find us on



fairmuse.eu



Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.