



POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

From Signal to Signing

Decision typologies in contemporary A&R and
an operational toolbox

TESI DI LAUREA MAGISTRALE IN
MANAGEMENT ENGINEERING
INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Author: Lorenzo Carissimi

Student ID: 225346

Advisors: Alessandro Biffi, Giulio Arata

Academic Year: 2025-26

Questa tesi è dedicata ai miei genitori, che mi hanno insegnato il valore dello sforzo, della fatica e delle rinunce.

Grazie per i vostri sacrifici, per i quali vi sarò sempre grato.

Abstract

This thesis investigates how contemporary Artists & Repertoire (A&R) work and decide under the conditions of technological change and uncertainty, where discovery and signals of audience response are increasingly mediated by streaming platforms, playlists, recommender systems, and metrics. The study addresses two questions: (RQ1) how the A&R role has evolved and fragmented into more specialized practices; and (RQ2) which empirical dimensions differentiate today's A&R practices and decision logics across contexts.

Methodologically, the thesis adopts an interpretive approach. Evidence is drawn from 10 semi-structured interviews conducted between November 2025 and January 2026 with A&R professionals and adjacent roles across heterogeneous organizational settings. Data were analyzed through a Gioia-informed, iterative coding process that maintains traceability from interview material to first-order concepts, second-order themes, and aggregate dimensions.

Findings are organized into a set of assignable empirical dimensions capturing how A&R manage workload, risk constraints, value orientation, data interpretation, dominant forms of “proof,” time horizons, and decision governance. These dimensions support a configurational typology centered on a single primary axis: Depth vs Volume (D1), defined as the “number of projects per A&R” and the resulting attention structure. Two recurrent positions emerge. A “Depth/Focus” position (P1) operates with a deliberately limited portfolio and concentrates decision-making on defensible signing choices supported by signal triangulation and internal alignment. A “Volume/Parallelism” position (P2) operates with many simultaneous projects and features continuous prioritization and resource reallocation under multi-channel noise and reporting constraints. Across both positions, decision-making appears as managing trade-offs rather than optimizing a single objective, and key tensions shift in meaning and salience between P1 and P2. Building on these results, the thesis translates the typology into practical toolboxes tailored to each position (rituals, one-page templates, decision rules, and metric guardrails), enabling light implementation without bureaucratization.

Overall, the thesis contributes a traceable, data-driven typology of A&R decision-making and a set of operational implications that help future A&Rs make choices more explicit, comparable, and defensible across different organizational conditions.

Keywords: A&R; music industry; streaming platforms; gatekeeping; cultural intermediaries; configurational typology

Abstract in Italiano

Questa tesi indaga come gli A&R (Artists & Repertoire) contemporanei lavorano e prendono decisioni in condizioni di cambiamento tecnologico e incertezza, in cui la scoperta e i segnali di risposta del pubblico sono sempre più mediati da piattaforme di streaming, playlist, sistemi di raccomandazione e metriche. Lo studio affronta due domande: (RQ1) come il ruolo dell'A&R si sia evoluto e frammentato in pratiche più specializzate; e (RQ2) quali dimensioni empiriche distinguano le pratiche e le logiche decisionali dell'A&R di oggi nei diversi contesti.

Dal punto di vista metodologico, la tesi adotta un approccio interpretativo. Le evidenze derivano da 10 interviste semi-strutturate condotte tra novembre 2025 e gennaio 2026 con professionisti A&R e ruoli adiacenti in contesti organizzativi eterogenei. I dati sono stati analizzati tramite un processo di codifica iterativo, Gioia-informed, che mantiene la tracciabilità dal materiale d'intervista ai concetti di primo ordine, ai temi di secondo ordine e alle dimensioni aggregate.

I risultati sono organizzati in un insieme di dimensioni empiriche assegnabili che catturano come gli A&R gestiscono carico di lavoro, vincoli di rischio, orientamento di valore, interpretazione dei dati, forme dominanti di “prova”, orizzonti temporali e governance decisionale. Tali dimensioni supportano una tipologia configurazionale centrata su un unico asse primario: Profondità vs Volume (D1), definito come “numero di progetti per A&R” e la conseguente struttura dell'attenzione. Emergono due posizioni ricorrenti. Una posizione “Profondità/Focalizzazione” (P1) opera con un portafoglio deliberatamente limitato e concentra il processo decisionale su scelte di signing difendibili, supportate dalla triangolazione dei segnali e dall'allineamento interno. Una posizione “Volume/Parallelismo” (P2) opera con molti progetti simultanei e presenta una prioritizzazione continua e una riallocazione delle risorse sotto rumore multi-canale e vincoli di reporting. In entrambe le posizioni, il decision-making appare come gestione di trade-off più che come ottimizzazione di un unico obiettivo, e le tensioni chiave cambiano significato e salienza tra P1 e P2. Sulla base di questi risultati, la tesi traduce la tipologia in toolbox pratiche tarate su ciascuna posizione (rituali, template one-page, regole decisionali e guardrail metrici), consentendo un'implementazione leggera senza burocratizzazione.

Nel complesso, la tesi contribuisce con una tipologia tracciabile e data-driven del decision-making A&R e con un insieme di implicazioni operative che aiutano gli A&R futuri a rendere le scelte più esplicite, comparabili e difendibili in diverse condizioni organizzative.

Parole chiave: A&R; industria musicale; piattaforme di streaming; gatekeeping; intermediari culturali; tipologia configurazionale

Sommario

Abstract	4
Abstract in Italiano	5
Chapter 1 - Introduction	9
1.1 Context and motivation	9
1.2 Research problem and gap: what is missing in the literature	10
1.3 Objective, research questions and expected contributions.....	11
1.4 Research design, data, and logic of the analysis	12
1.5 Chapter roadmap.....	13
Chapter 2 – Context and literature	14
2.1 The A&R in the pre-platform regime.....	14
2.2 The digital transition and the transformation of the A&R role	18
2.3 A&R as cultural intermediary and gatekeeper	23
2.4 Decision-making under uncertainty and trade-offs	28
2.5 Metrics, playlists, and rankings as judgment devices and evaluation regimes.....	31
2.6 Synthesis: sensitizing concepts and what to observe empirically.....	34
Chapter 3 – Methodology.....	35
3.1 Research design and rationale.....	35
3.2 Sampling and interviewee profile	37
3.3 Data collection: protocol, ethics, material management	38
3.4 Data Analysis	39
3.5 Quality and rigor criteria	41
Chapter 4 – Results	43
4.1 — Snapshot of the sample and of contexts as modifiers	43
4.2 — Data structure (Gioia-informed) and themes overview	45
4.3 – Candidate dimensions	50
D1 — DEPTH VS VOLUME.....	53
D2 — SLACK/ECONOMIC RISK	54
D3 — VALUE COMPASS	55
D4 — HOW IT USES DATA.....	56
D5 — MAIN PROOF	57
D6 — Time horizon	58
D7 — Governance (who decides)	59
D8 — Coordination (how many dependencies)	60
D9 — Validation runway and penalties (cliff)	61

M1 — AI (modifier)	62
4.4 — Configurational typology	63
4.5 — Type Cards	66
4.6 — Comparative problems / constraints / trade-offs	71
Chapter 5 – Tool Design.....	77
5.1 Guiding principles for translating empirical evidence into operational implications.....	77
5.2 Toolbox by typology and Toolbox Map	79
5.2a Toolbox P1 (D1=0: Depth/focus)	80
5.2b Toolbox P2 (D1=2: Volume/parallelism)	85
5.3 Lightweight implementation	88
5.4 Closing and bridge to Chapter 6	89
Chapter 6 — AI in A&R work	90
6.1 Why AI is often marginal in the everyday practices of the A&R	90
6.2 What happens today	91
6.3 Future perspectives: plausible shifts and related risks	92
6.4 Minimal guardrails for using AI under platform constraints	93
6.5 Ex-ante adoption criteria	93
Chapter 7 - Conclusions.....	95
7.1 Answer to RQ1 - Evolution of the A&R role between fragmentation and specialization ..	95
7.2 Answer to RQ2 - D1 (Depth vs Volume) and the P1/P2 typology	96
7.3 Contributions.....	97
7.4 Limitations, future research, and conclusion	98
References.....	99
List of Figures	104
List of Tables	104

Chapter 1 - Introduction

1.1 Context and motivation

Over the last fifteen years the music market and the ways music is consumed have changed radically; in this new context the work of A&R (Artists and Repertoire) has had to update accordingly, because it no longer operates in a value chain dominated by majors, physical sales, and radio. The change concerns not only where music is listened to, but how visibility is built: in streaming a large share of listening passes through an enormous catalogue in which the user chooses through tools offered by the platform (search, recommendations, playlists, charts) (Morris & Powers, 2015). In practice, a relevant share of discovery no longer depends only on radio, stores, the press, or word of mouth, but also on how platforms put certain releases and certain artists “in the window”.

This “window” is managed in a hybrid way. On the one hand there is human work (editors/curators who build playlists and select content); on the other hand there are recommendation systems (algorithms that propose tracks and playlists based on users’ behaviours). The combination of the two levels makes visibility more dynamic but also harder to interpret: a spike in listens can depend on an editorial choice, on an automated mechanism, or on external factors (social, live, media), and often these causes overlap (Bonini & Gandini, 2019).

For A&R this scenario has three practical consequences:

First: the weight of tools that help decide under conditions of uncertainty grows. In cultural markets quality is not measurable with a single parameter and no one can know in advance what will work. For this reason, historically, decision-makers have relied on signals and orientation devices (reputation, critics, charts, radio, brands, awards). Karpik (2010) describes these devices as instruments that help make a decision on a unique good by reducing uncertainty through trust and reputation. In the context of platforms, such signals increase in number and in detail (micro-metrics, rankings, playlists) and become more frequent. The advantage is having more signals to orient oneself; the risk is that some become evaluation shortcuts (for example, confusing popularity and quality) and that judgement flattens onto what is more visible and measurable.

Second: data increase, but the ability to understand why something is happening does not increase proportionally. Dashboards highlight recent variations (rises, falls, spikes), and this encourages fast decisions: when to publish, what to push, where to concentrate budget, when to relaunch a track, how to set up pitching. However, the available data are partial and causal attribution remains difficult: seeing an effect is easier than understanding its cause. Moreover, not everyone sees the same data nor has the same analytical capabilities: platforms and large organizations can rely on richer dashboards and dedicated roles, while smaller realities often work on partial indicators; this amplifies differences in the quality (and in the speed) of decisions (Maasø & Hagen, 2020).

Third: the incentives and strategies of the involved actors change. If visibility also passes through mechanisms internal to platforms, artists and organizations are pushed to adapt: taking care of the

first seconds of the track, designing releases and social content in a way compatible with logics of attention and recommendation, optimizing indicators that platforms make salient. Morris (2020) describes this set of adaptations as a process of cultural optimization induced by the infrastructure. The literature on recommenders also helps to understand why: these systems are not neutral; they are often optimized so that the user stays longer and returns more often to the platform (Seaver, 2019).

In sum, when a track's visibility depends at the same time on editorial choices, algorithms, and metrics, for an A&R it becomes harder to understand which signals really matter and why a result is happening; consequently, deciding is more complex and less transparent. This opens a basic question: in what way, concretely, do A&Rs make decisions today and what systematic differences exist between different ways of working? The next section translates this intuition into a precise research problem and makes the gap explicit.

1.2 Research problem and gap: what is missing in the literature

The literature on platforms, playlists, metrics, and recommendation systems is broad; more limited, instead, are the studies that reconstruct up close how these elements enter, step by step, into the decision-making work of A&Rs. Important contributions exist (for example ethnographies, studies on the signing process and on selection criteria), but they often appear dispersed across contexts, phases of the process, and different segments.

The literature offers solid contributions on three levels, which however often remain separate:

1. **Macro framework and industry transformation.** Various works reconstruct the shift to streaming models and the increase in the centrality of platforms (streaming and social) in music discovery, highlighting how playlists and recommendation contribute to organizing attention and consumption (Wikström, 2019; Prey et al., 2020). This literature clarifies well why the environment changes, but it tends to stop before describing how managerial decisions are made in everyday practice.
2. **Evaluation regimes and judgment devices.** Economic sociology studies show that, in cultural markets, quality is uncertain and actors rely on devices (charts, reputation, critics, rankings, playlists, metrics) to make judgment workable (Karpik, 2010). These ideas are crucial to understand the role of playlists and metrics, but they still do not say which combinations of signals are actually used by A&Rs, with which trade-offs and with which routines.
3. **Metrics and datafication in streaming.** Empirical works show how metrics guide micro-decisions and how attention to sudden peaks ("spikes") and to the most salient signals coexists with low transparency on how results are generated (opacity) and with strong differences between actors in terms of available resources and competences (Maasø & Hagen, 2020). Here too, however, the object is often the ecosystem

(platforms/labels/distributors) more than a systematic description of A&R decision logics across different contexts.

From this fragmentation derive three gaps that this thesis intends to address.

1. **Descriptive:** a close-up and comparable empirical reconstruction is missing of how A&Rs work and decide today (mission, workflow, signals/metrics, stakeholders, constraints), beyond generic labels and beyond single cases.
2. **Analytical:** even when descriptions of practices exist, they are rarely organized into a configurational typology, that is, a map of recurring ways of combining dimensions, with clear assignment rules and attention to negative cases. This is relevant because the same signal (e.g., a metric) can have different meanings depending on the mission and constraints of the A&R.
3. **Applied:** the literature rarely translates, in an explicit way, the evidence on decision-making practices into operational tools (templates, rituals, useful vs toxic metrics, guardrails) designed to manage the trade-offs and problems of each configurational typology.

In response, the thesis builds a bridge between these strands: it uses a rigorous qualitative setup to make dimensions and configurations emerge from the data (Gioia, 2012) and to produce operational typologies, which make it possible to describe and compare A&R ways of working and, from there, derive practical implications. Section 1.3 describes in detail the objective, research questions, and expected contributions.

1.3 Objective, research questions and expected contributions

This thesis therefore aims to clarify how A&Rs work and make decisions today, in a context dominated by platforms, metrics, and visibility signals. The objective is not to propose a single best practice, but to build an empirical map of recurring ways of working: starting from interviews, to reconstruct the dimensions that really make a difference in A&R work. On this basis, the dimensions are then read together to understand which combinations recur in practice and to build a configurational typology with explicit assignment rules. Finally, for each type the thesis provides a comparable description (objectives, routines, signals/metrics, stakeholders, trade-offs) and translates the evidence into operational indications designed to manage trade-offs without reducing judgment to a single KPI.

To do this, this thesis adopts two research questions, which guide data collection and analysis.

- **RQ1.** How has the A&R role evolved and how has it been articulated between fragmentation and specialization?

- **RQ2.** Which empirical dimensions distinguish today A&R practices and decision logics across different contexts?

From the empirical dimensions that emerge (RQ2), the thesis then builds a configurational typology (matrix/axis/to be seen) with explicit assignment rules, synthesizes each configuration into comparable type sheets, and highlights recurring constraints and tensions. On this basis, finally, it translates the evidence into operational implications in the form of a toolbox (rituals, templates, useful vs toxic metrics, guardrails).

This thesis aims to contribute, first of all, by making practices comparable that are often described in a fragmented way: it focuses on how A&Rs define success, which routines they adopt, and which signals they use (or avoid), reading these micro-practices within the platform context. On this basis, the analysis does not stop at the description of single cases: it organizes the evidence into a configurational typology that clarifies which combinations of dimensions recur and, above all, helps distinguish what depends on the type of work from what depends on the context (resources, segments, constraints). Finally, the thesis makes explicit what can be done with these results: it translates the evidence into a toolbox by typology (rituals, templates, useful vs toxic metrics, guardrails) designed to manage trade-offs, while keeping traceable where the decision remains inevitably human.

1.4 Research design, data, and logic of the analysis

To answer the research questions, the thesis adopts a qualitative design based on semi-structured interviews with professionals who perform (in different forms) the A&R function, and with the realities that orbit around it.

Sampling follows a maximum variation logic, but within a clear scope: most interviewees directly perform the A&R function, and a minority share includes adjacent roles that interact with A&R and influence its constraints and choices. The planned sample size ($n=9$) is a pragmatic threshold: it is in line with the values reported in studies on saturation and it was monitored during the analysis, extending the sample when new themes emerged (Hennink & Kaiser, 2022)

The interview analysis adopts a Gioia-informed approach: from the interviewees' accounts, first-order concepts (informants' language) are built and then second-order themes, up to aggregate dimensions; the process is iterative and traceable through memos and decision rules (Gioia, 2012; Locke et al., 2022). The emerged dimensions are then read in combination to build a configurational typology and the related assignment rules (Ch. 5).

1.5 Chapter roadmap

Chapter 2 reconstructs the functional evolution of A&R work from the major/retail/radio regime to the platform-centric phase, and introduces a minimal conceptual framework to read the observed practices (A&R as intermediation/gatekeeping, decision-making under uncertainty and trade-offs, evaluation regimes). Chapter 3 describes the research design, sampling, data collection, and analysis. Chapter 4 presents the results: emerged dimensions, configurational typology, type sheets, and recurring trade-offs. Chapter 5 discusses the results, answers the RQs, and translates the results into operational implications (toolbox and guardrails). Chapter 6 discusses criteria and guiding questions to evaluate AI adoption in coherence with the emerged typologies and closes with the final conclusion.

To avoid ambiguities, this thesis adopts some operational definitions, useful to clarify what is observed and what remains outside the analysis.

A&R. By “A&R” I do not refer to a single job title, but to a function: the set of activities and decisions related to (i) identifying repertoire/artists, (ii) evaluating and selecting them up to signing, and (iii) setting the first development and positioning choices in coordination with downstream functions. This reading is consistent with the literature that describes A&R as work of selection and mediation under conditions of uncertainty, in which judgment integrates formal and informal signals (Negus, 1992; Negus, 2002; Zwaan & ter Bogt, 2009).

Boundaries of the A&R process. The definition above describes what the A&R function is; here instead I set an operational boundary useful for the analysis. In this thesis, by “A&R process” I mean the arc that goes from the first signal (scouting/introduction) up to an observable decision outcome (signing proposal/decision not to proceed) and, when one decides to proceed, it also includes the first development/positioning choices directly linked to the A&R function (for example, the setup of the repertoire and of the creative team), so as to be able to compare choices and routines across cases without letting the scope slip onto downstream executive activities. Downstream activities (marketing, distribution, advertising, release operations) are not the object of the thesis as such; they enter the scope only when they become a constraint, stakeholder, or source of signals that changes the way A&R decides.

Organizational scope (label and publishing). The empirical field includes professionals who perform this function in different ways and contexts, from majors (the multinationals that dominate the global music market) to smaller realities (such as independent labels). In practice this means that part of the activities can vary, but what interests the thesis is how criteria, signals, stakeholders, and trade-offs enter decisions.

Chapter 2 – Context and literature

In this chapter, the first two sections (2.1 – 2.2) aim to describe how the profession of the A&R was during the pre-platform regime, and how it has developed until nowadays, in a new context where the music industry is dominated by the platforms, and new players and logic are introduced in the market.

The chapter begins with 2.1, where I describe the ‘classical’ role of the A&R in the pre-platform era, and in 2.2 I analyse the digital transition of the environment and its consequences over the identity of the A&R role. The other sections are dedicated to a minimal conceptual framework, to solidify the theoretical basis of the next chapters: in 2.3 I identify the A&R as a cultural intermediary and a gatekeeper; in 2.4 I describe the decision making process under uncertainty and trade-offs; in 2.5 I describe metrics, playlists, and rankings as judgment devices and evaluation regimes; and finally in 2.6 there is a brief synthesis turning the conceptual framework into a practical guide which will be used during the analysis in the next chapters.

2.1 The A&R in the pre-platform regime

In the pre-platform regime, A&R performs a precise function: first of all, it is a function of selection under uncertainty. It is more correct and useful to identify this professional figure not only as a talent scout who aims solely and exclusively at pure talent, but rather as a job of decision-making and resource allocation under uncertainty. In fact, among very many possible projects, A&R operates in a corporate context in which resources are not unlimited, and is therefore forced to choose where to invest time, budget, and attention, building reasoned and defensible bets. These choices are never simple: in fact, according to the theory of cultural markets (such as music), quality and demand are difficult to estimate in advance and outcomes depend on contingent combinations (as in this case timing, audience reception, media/gatekeepers, internal coordination). In this context, therefore, the project that is judged as “right” can still fail, while an unexpected one can blow up. For this reason, A&R and the label resort to practical heuristics, comparisons with previous cases, and network/reputation signals to make the decision reasonable and defensible, rather than certain (Caves & Richard, 2000; Negus, 1992). This betting logic does not concern only label budgets. It also has an implicit effect of allocating opportunities in a cultural labor market that is structurally oversupplied: if potential projects are many more than the available resources, and careers are often discontinuous, then A&R selection also functions as a mechanism that distributes access to scarce resources (such as studio recording time, budget, team, etc.) (Hesmondhalgh, 2019). In this hyper-competitive context, where the supply of projects far exceeds the capacity to invest and to pay attention, A&R ends up also acting as a gatekeeper, because its choices affect which professional trajectories receive resources and visibility and which remain at the margins.

Within the framework of the cultural industries, A&R works inside a recurring tension: supporting creative processes that require autonomy and time, but at the same time making investment

decisions and coordinating activities oriented toward outcomes of visibility and economic sustainability. However, it often happens that many crucial choices occur before robust market evidence actually exists. In cases like these, where uncertainty reigns supreme, A&R connects different worlds (artists, producers/studios, management, and internal functions) and tries to correctly translate partial signals into a practical judgment on potential, risks, and path.

In this context, the A&R starts scouting work by building over time a portfolio of projects under observation: scouting serves to progressively filter an abundance of proposals and to decide which ones deserve further deepening. To do this, A&R collects signals from multiple channels and returns multiple times to the same names as new information emerges. Scouting, therefore, is not a single moment but a continuous activity of observation and screening, in which A&R progressively decides which projects deserve further deepening and which, instead, should be left behind (Negus, 1992). Scouting therefore does not coincide with the final choice: it is a work of uncertainty reduction that serves to predict with greater precision which projects deserve further time and resources. Negus describes artist acquisition as a sequence of filters and evaluations that passes through multiple channels (demos, live performances, recommendations from managers/producers and independent labels). In this sequence, what is evaluated is not only the quality of a song, but also the compatibility of the project with the roster (that is, the set of artists/projects already in the label's catalogue) and with a plausible initial trajectory.

When the signing decision (or, more generally, the investment decision) gets closer, A&R has to transform observation into a defensible proposal, where it describes the reason for choosing the project and under which conditions. The evaluation often includes both artistic elements (recognizable identity, repertoire or repertoire potential), and feasibility elements (team reliability, ability to produce and perform, long-term resilience), and strategic elements (label positioning, coherence with the catalogue, possibility of building a first credible direction) (Caves & Richard, 2000).

Meanwhile, the development phase of the project can take place both after having closed the actual signing of the recording contract, and in the final phase of selection that happens before the proposal of the project to the label. In this phase, in fact, a good part of A&R work concerns improving the artistic project and aligning it with market logics: from refining the repertoire, to choosing or consolidating key collaborators (producers, writers, arrangers), up to the definition of a coherent identity and of a first narrative of the project. In this phase, by "positioning" I mean the operational translation of artistic identity into a recognizable proposal for the market: which audiences and listening contexts to intercept, with which references (scenes, genres, comparable artists), through which channels and with which distinctive promise. Here development does not necessarily coincide with direct creative control, but more often it is coordination (even if every A&R reasons and works differently): putting competences and resources in relation, facilitating decisions, and creating conditions so that the project reaches the public with a recognizable direction (Hesmondhalgh, 2019). In doing this, "classic" A&R is not an isolated function, but its activities intertwine with marketing, promotion, and an external network of professionals (management, promoters, studios, media): an important share of the work is therefore relational. In this configuration A&R does not control the outcome alone, but it affects which inputs enter the process and how the organization tries to make them circulate. Coherently, before formalizing a

signing, steps of internal validation often emerge. To try to reduce uncertainty and also evaluate other points of view, A&R compares with senior figures and with functions such as marketing, radio/TV promotion, press, and sales. Sometimes an interfunctional project team is also formed and development meetings are held to align musical direction and operational priorities (Negus, 1992).

During this evaluation phase, information is analysed in order to strengthen the power of the decision. However, in the pre-platform regime, the nature of the information is very different from what we are used to today: on the one hand, data on consumption exist but are not very detailed, in fact they arrive late, are extremely aggregated, and therefore say little about who listens, how they listen, and why (unlike the data available nowadays); on the other hand, indirect and contextual signals weigh heavily, such as word of mouth within the professional network, reactions to live shows, and early responses from the press and radio, which however are not objective measures of demand, but clues to be interpreted. Consequently, in this historical context the decision tends to be based on experience, comparisons with previous cases, and interpretation of noisy signals. The classic literature stresses that, to manage *ex ante* risk, labels develop practical rules, comparisons with previous cases, and rely on the reputation of key intermediaries (managers, producers, other A&Rs) (Hesmondhalgh, 2019).

This logic does not make the choice random, but it makes it dependent on the context (genre/scene, label objectives, available resources, career moment, collected signals) and built through successive steps. In practice, A&R does not decide everything in one shot: first it observes, then it deepens, then it activates small tests, and only at the end, if the signals hold, does it propose an investment. The work is articulated in successive micro-decisions up to formulating a path hypothesis (repertoire, team, positioning) sufficiently coherent to justify the label's commitment. Finally, even without granular data, A&R performs a form of reading of the social: it has to anticipate tastes and pleasures of plural publics and translate such hypotheses into scouting and greenlight choices (Wikström, 2019).

Moreover, it is important to analyze how the major/retail/radio regime makes these practices coherent and rational: which actors, institutions, and bottlenecks (channels, gatekeepers, infrastructures) concretely structure A&R work, especially in the access-to-visibility phase. It is therefore within this historical context that by major/retail/radio regime I mean the moment in which the visibility value chain of music had a very precise shape, dominated by majors (or similar realities), physical retail (by now marginal), and national radios. This configuration can be read as a multi-stage preselection system: from internal selection (A&R/label) up to media gatekeepers, who progressively filter an abundance of supply, letting only a minimal fraction of the product reach the public.

Between about 1950 and the end of the 1990s, the recorded music market operated in a relatively stable configuration: visibility (and therefore the probability of scaling) passed through three bottlenecks that reinforced each other. More in detail, these three bottlenecks worked together. On the one hand there were the majors, which were not only big labels: they were multinationals with capital, staff, relationships, and operational capabilities to turn a project into a national or international launch. On the other hand, there was physical retail, that is, stores and the distribution

of physical formats: shelf space was limited, orders were selective, and availability in points of sale decided very concretely what the public could find and buy. Finally, there was broadcast radio, where airtime was scarce: getting into rotation and into playlists meant obtaining continuous repetition that could create familiarity and demand on a large scale, while being left out often meant remaining invisible to the generalist public. On the concrete role of radio as a market lever in the pre-digital regime, Negus observed that it was “to a large extent radio airplay” that “pushed” both retail to stock and consumers to buy records (Negus, (Prey et al 2020). In this setup, commercial success depended to an important extent on coordinated access to these channels: having a major behind you (or other realities with equivalent resources) helped to obtain distribution and sustain campaigns; good distribution made the promotional investment more credible; and radio, in turn, could push both retailers to stock and consumers to buy. For A&R, the consequence was practical: choosing an artist did not mean only judging their potential, but estimating whether and how the project could be activated within these bottlenecks (and with which resources and relationships). From here an operational effect: selection and promotion tended to intertwine, because choosing a project already meant estimating whether and how it could obtain access to the dominant channels and with which resources. Over time, the recording industry tended to consolidate into an oligopoly of large operators, while a larger population of independent labels entered and exited the market. In many reconstructions, this coexistence fed a “major vs indie” logic in which independent labels also functioned as a laboratory of experimentation and incubation, while majors could absorb or acquire projects that had already passed a first market test (a dynamic that becomes particularly visible starting from the 1990s) (Peterson & Berger, 1975; Peterson & Berger, 1996). These asymmetries could also be seen in exposure: in commercial radio in the 1980s, majors produced about 54% of new titles but obtained about 81% of airtime; indies and other operators split the remaining ~19% (Karpik, 2010). For A&R, this means operating inside a constraint: the value of a project is also evaluated in relation to the organization’s capacity to support its access to the channels.

Another important piece of the major/retail/radio regime was the weight of charts and “Top” formats (for example Top 40). When most industry people looked at charts to understand what was working, those charts became in practice a common reference for deciding. The risk, however, is that the chart was used as a shortcut: instead of really evaluating a project (music, identity, possible path), one ended up replacing part of judgment with a simple number, that is, the chart position. This favors coordination by conformity and reinforces winner-take-more dynamics, because what enters high in the charts tends to receive even more exposure and resources. (Karpik, 2010). However, charts and rankings can be read both as signals of demand and as effects of the exposure produced by the regime we were talking about before. In the absence of granular data and measurement tools comparable to today’s, this distinction remained largely an interpretive evaluation.

In sum, in the major/retail/radio regime A&R:

- selects and develops projects under strong constraints of access to external channels
- operates in interdependence with functions and specialists oriented toward activating such channels (radio/TV promotion, press, sales/retail)

- tends to transform uncertainty into “defensible” decisions through comparisons, network reputations, and internal alignment, because consumption signals are less granular and slower.

In this historical moment, visibility depended above all on radio and stores. With the advent of digital, instead, selection gradually shifts onto other nodes: platforms, online catalogues, and recommendation systems. Consequently, the signals on which A&R can rely and the concrete levers through which it can intervene also change.

2.2 The digital transition and the transformation of the A&R role

Following a long period in which music visibility passed mainly through physical retail and broadcast, the major/retail/radio regime saw an important change: with the birth of digital infrastructures and, later, of streaming platforms, music circulation and discovery progressively shifted onto new access and selection nodes. In this new historical period, platforms started (and still do) to dominate the way people approach listening to music. To truly understand the ongoing transition, therefore, it is not enough to say that distribution became digital: it is necessary to reconstruct how the economic model was reorganized and what new gatekeeping architecture was formed.

From the end of the 1990s, the phenomenon of file sharing spread, that is, the direct exchange of music files between users via the Internet, most of the time without authorization. In practice, a person would share on their computer tracks in digital format (for example MP3) and other users could search and download them. Many services worked with peer-to-peer logics: there was not a single store or a single central catalogue, but a network of computers connected to each other. This made copying and distributing music extremely fast and at almost zero cost for users, breaking the economic logic in which the sale of a physical (or digital) copy was the main source of revenue. This led to the weakening of CD and other format sales and pushed the model based on owning the copy into crisis.

To respond to this phenomenon, the recorded music market went through an epochal change, easy to summarize but enormous in its consequences: instead of buying a copy (a CD or an iTunes download for example), people started paying mainly to obtain continuous access to an almost infinite catalogue, through streaming platforms. The market moved from owning music (first on physical formats, then as downloads) to being able to listen to it whenever you want by paying for a service, through a subscription (Hesmondhalgh, 2019).

At the same time, online advertising increased, but a growing share of that spend concentrated on search engines and social media controlled by large companies. This reallocation of advertising spend increased revenues and investment power for large IT companies. With more capital and, above all, with control of digital channels and user data, these firms could also enter adjacent

markets such as online music distribution, up to becoming central actors in streaming and music discovery. Towards the end of the 2010s, the access-based streaming market thus also became a battlefield: alongside platforms born as start-ups (and that survived the initial phase), large tech groups entered with different strategies (Wikström, 2019).

When we talk about streaming platforms, we often imagine them as a neutral environment that has democratized access to music, making it simpler and more available for everyone. However, this narrative forgets to consider an important point. A first practical consequence of consuming music through streaming platforms is that the user does not encounter the catalogue without filters, but the “window” is built by the interface: the home and the “for you” sections foreground some content, while search and categories guide what is easy to find and what remains out of frame. For example, as regards Spotify, recommendations are suggestions generated by combining the user’s listening data and those of their similar profiles, platform rules, and often also editorial choices. Playlists are the most common form of this layouting: collections by mood, activity and/or genre that do preselection and repetition, directly influencing what ends up being listened to more (Maasø & Hagen, 2020).

A music streaming platform therefore is not a simple container that hosts tracks, but a fundamental player in the decision of how those tracks are encountered by listeners, organizing and staging the catalogue through interfaces (home, feed, search), rules (what is suggested, what is pushed, what is hidden), and data systems that observe listening behaviors and translate them into classifications and recommendations. In practice, the platform does not limit itself to transporting music from A to B: it transforms it into a guided experience, where some tracks become visible and repeated (for example through playlists, “for you” suggestions, rankings such as Top50), while others remain in the background, even if technically available. This has a simple consequence: access to the public passes increasingly through decisions and mechanisms defined by the platform (editorial and/or algorithmic) and not only through the choices of labels and artists. In this way, fertile ground has been created to generate strong dependency relationships: artists, labels, publishers, and creators today depend a lot on platform choices, since without visibility inside the platform it is much more difficult to be discovered, build continuity of listening, and convert attention into a career.

This new architecture above all produced a shift in power: a growing share of music circulation and discovery passed through digital distributors (platforms such as Spotify, YouTube, TikTok), therefore outside the direct control of labels and their traditional channels. In operational terms, this also changed the A&R job: part of “making a project happen” (entering listening flows, being found, staying visible over time) depended more and more on rules and choices defined by the platform. Here I take Spotify as an example because it is the most studied platform in the literature and because it makes the mechanism very clear; but the logic is analogous on the other platforms as well, each with its own interfaces and criteria. The consequence is that A&R and labels found themselves having to read and negotiate a new level of gatekeeping, where access to visibility is mediated by socio-technical systems (editorial and algorithmic) that can amplify some content and make other content practically invisible (Wikström, 2019).

Beyond the dynamics already presented, streaming also caused a shift of the bottleneck: with millions of tracks available, the problem is not supply but attention. The growth of supply (tens of

thousands of new tracks per day) makes a strong preselection upstream of listening inevitable, and therefore strengthens the power of those who govern visibility devices, namely the platforms themselves. As streaming matured, differentiation on content tended to saturate: the surviving platforms ended up offering enormous and increasingly similar catalogues, comparable technical quality, and availability on many devices. In conditions like these, competitive advantage shifted from what one owns to how one makes the catalogue used and discovered: search, navigation, organization functionalities, and above all playlists (Wikström, 2019). Within this logic, playlists become the new schedule of the digital music economy: a preselection mechanism that can work in an analogous way to twentieth-century radio, but with finer control over who is exposed to what, when, and in which position (personalization, testing, global scalability) and with a direct link to behavioural data. A strong piece of evidence of curatorial power concerns the dominance of proprietary playlists. As a snapshot of the late 2010s, Spotify had built thousands of its own playlists and tended to promote such assets at the expense of third-party playlists generated by its users: in January 2019, of the 100 most followed playlists, 99 turned out to be all Spotify-owned/curated and 1 was generated by third parties. A study commissioned by the European Commission estimated that playlists curated by Spotify gathered over three quarters of the followers of the top 1,000 playlists; to these was added about a further 9% attributable to Spotify's algorithmic playlists. In this picture, major-label playlists (e.g., Filtr, Digster, Topsify) had much smaller shares (Prey et al 2020). This is not only a symbolic issue: the literature of Prey reports estimates according to which inclusion in large playlists can generate relevant additional revenues per track and that position in the playlist matters a lot, giving maximum impact when the track is placed near the top, requiring less user effort to dig up the track itself. The result is an increase in dependence: to build and maintain a career, artists and labels become more constrained to obtain positions in playlists curated by the platform. From here a useful operational definition arises: for platforms that do not own the rights to the content they distribute, playlists become a key mechanism to exercise curatorial power, that is, the capacity to advance one's own interests and influence those of others through the organization and programming of content.

Moreover, compared to the previous regime, in which information arrived late and in a coarse form, streaming has made much closer and more detailed signals available. Today it is possible to observe a track's trajectory day by day (and sometimes in near real time) and distinguish where and how it is performing: which playlists it enters, in which countries or cities it grows, whether listens increase but saves do not, whether there are spikes linked to a social content or to an editorial push. In practice, data are no longer a final outcome that one discovers afterwards, but a continuous flow that can influence decisions while the campaign is still ongoing: what to push, when, with which message, and on which channels. Qualitative research shows in fact how dashboards and metrics (including spike notifications) enter managers' decision routines, with a growing event-driven logic and a "fresh" data culture. In these analyses it is highlighted that many decisions are made by looking at what is happening in that precise moment (a spike in listens, an entry into a playlist, a trend in a country, a jump in saves), instead of waiting for a balance at the end of the campaign. The point, however, is that these signals often arrive as indicators more than as true explanations: seeing a number go up does not always make it clear why this number is going up. For example, an increase in a track's streams can depend on an editorial playlist, on an automated recommendation, on a push on a social platform, or on a combination of factors, radically changing the correct interpretation of the same data; from the dashboard, however, the causal link is not

always reconstructible with certainty. In other words: data are useful, but they are also opaque in their mechanisms, and this forces managers and A&Rs to work through hypotheses and tests, rather than through certainties (Maasø & Hagen, 2020).

It is for this reason that in a scenario where the same data can lead to different outputs and results, analytical competences became more valuable: knowing how to read the numbers, distinguish them by source (organic vs induced), and connect them to operational decisions (timing, budget, priorities, message) by reacting in a sensible way. At the same time, however, these competences and tools were not distributed equally: those who had teams, time, and know-how could interpret signals better and activate them; those who had less risked chasing only the most visible numbers, without really understanding their origin and without being able to turn them into effective actions. In sum, the increase in connectivity reduces labels' capacity to regulate the informational flow as in the physical-format economy, but it also produces a new lever: more tracks and behaviors become observable.

In the shift from the major/retail/radio regime to the digital ecosystem, A&R work did not change only the tools: the division of labor that makes it possible to discover, develop, and circulate a project changed. Channels increase (streaming, social, UGC), supply grows, and gatekeepers do not disappear but multiply: attention becomes the scarce resource and visibility depends more and more on interfaces, rankings, playlists, and recommendations. In this scenario, A&R tends to specialize (or to split into sub-functions) because properly covering scouting, creative development, positioning, digital activation, and data reading becomes too costly for a single profile. In this framework indeed, there is not a single new way to approach the A&R role, because this figure became more and more specialized and fragmented across practices plausible.

Starting from the 1980s, the cultural industries increasingly emphasized specific segments (age, lifestyle, niches) and forms of targeted marketing; at the same time, however, the mass-market logic and the superstar logic remained central. From this derives a coexistence: micro-scenes and micro-publics proliferate, but the head of the market continues to capture a large share of attention (Hesmondhalgh, 2019). For A&R this changes the practical question. It is no longer only "is this thing strong?", but also: "for whom is it strong, in which listening contexts, and through which channels is it realistic to build continuity?". In other words, positioning becomes an operational job: translating an artistic identity into a recognizable proposal, identifying reference artists and adjacent scenes, clarifying what distinctive promise the project offers, and translating all of this into a first plausible path made of release strategy, content, and collaborations, within a saturated environment.

With digital, the shelf constraint (retail) loses centrality: almost everything is technically available, but the listener can still devote time and attention only to a tiny fraction of the supply. The bottleneck, therefore, is no longer finding space, but obtaining attention. In this situation, to orient oneself within abundance, the tools that narrow choices and make chaos manageable increase (and become central). Here by uncertainty I mean something very concrete: before investing time and resources, no one can know reliably which track or artist will truly become listened to, which will generate revenues, and which instead will remain invisible, even with the same perceived quality. For this reason, listeners and operators (including labels and A&Rs) lean on practical shortcuts,

such as charts, editorial playlists, the platform's automatic suggestions (recommendations), social signals such as trends and UGC, as well as reviews and reputations. These tools therefore function as judgment devices: they reduce the effort of choice and transform a sea of alternatives into a few visible and comparable options, and precisely because they concentrate attention on some options, they often also end up influencing which tracks emerge (Karpik 2010). For A&R, the critical point is that these devices are not only mirrors of demand: they often produce it or amplify it. If a track is put in evidence (for example at the top of an editorial playlist or in autoplay), it can generate many listens even without an "active" choice by the listener; those listens, in turn, become signals that feed new recommendations. In other words, within the digital ecosystem numbers can be at the same time an outcome and a cause of exposure. To decide lucidly, therefore, it is not enough to look at "how much" a number is: it is necessary to understand where it comes from (organic, editorial, algorithm, social) and whether it is creating deeper behaviors (saves, follows, return) beyond passive streaming (Prey et al 2020). As seen above, in the pre-digital phase a label could push a project mainly through channels it controlled or knew well (radio, press, retail). With streaming, instead, a relevant part of how people discover music and how a track continues to circulate over time passes through actors external to labels: streaming platforms, social platforms, and recommendation systems that decide what to make more visible and what to leave in the background. For A&R this meant having to cover a new level of gatekeeping: understanding what the main access points and internal levers of the platform are, and which actions increase the probability of entering and remaining in listening flows, even when the rules are not fully transparent. Concretely, an interface job was added: maintaining operational relationships with DSP teams when possible, preparing materials and pitching for playlists/initiatives, coordinating metadata and releases, and monitoring how the "window" is exposing a track in order to adapt timing and priorities (Maasø & Hagen, 2020). This work became more delicate because streaming also brought datafication, that is, the transformation of listening behaviors into metrics consultable in dashboards (listens, saves, skips, traffic sources). Data do not eliminate uncertainty, but they accelerate decision cycles: many choices become iterative and happen while a campaign is still ongoing (what to push, when, where, with which message). The practical point is that platform data show quickly what is happening, but rarely explain on their own why it is happening and what it is sensible to do next. For this reason, A&R work, beyond the relational component, increasingly requires a specific competence of reading and translating numbers into operational decisions: linking metrics to hypotheses, testing them, and understanding whether a result is only exposure or whether it is building deeper signals. In organizational terms, this further favors specialization and therefore fragmentation of the role: especially in recent years in fact (as we will also see highlighted in the following sections) in many organizations the need for figures, routines, or sub-tasks dedicated to monitoring, analysis, and reporting has grown enormously, because without these capabilities one risks chasing the most visible numbers without converting them into effective choices.

One last factor that pushes fragmentation and specialization concerns a very concrete question: who absorbs creative risk and who bears the costs of experimentation and early development? Who takes on the task, therefore, during the incubation period of a project, in which one has to iterate on repertoire, team, and positioning when market evidence is still limited? Over time, in the cultural industries the same dynamic has often been observed: a share of experimentation and early development tends to shift toward smaller and more flexible actors (indie labels, management,

studios/producers, micro teams), while larger actors such as majors enter later in the value chain, when some signals have already emerged and the investment appears more defensible. In this context, the phenomenon of restructuring is a central key. By restructuring I mean phases in which labels (especially large groups) redesign their setup to respond to economic pressures: cutting and reallocating budgets, rationalizing the roster, greater centralization of decisions, increased control and reporting on projects, and often consolidation operations or changes of ownership that push toward stricter profitability objectives. In these contexts it becomes rational to reduce exposure in the most uncertain phase of a project's life cycle: exploration (tests on repertoire, team, and narrative) tends to be supported by smaller and more flexible actors; larger players enter later, when minimal traction signals emerge (live, professional network, first consumption evidence) that make the investment more defensible (Hesmondhalgh, 2019; Karpik, 2010). For A&R this means two things. First: specialization is not only inside the label; it is distributed along the ecosystem (indies, management, publishers, creator economy, data services). Second: when incubation is outsourced or distributed, what A&R must be able to do at the point where it is located also changes. In some contexts a more "craft" and relational work remains central (scene, network, creative development). In others, instead, the weight of coordination and managerial evaluation grows: reading already-emerged signals, aligning internal functions, and deciding when to intervene and under which conditions (budget, timing, objectives) inside organizations with resources, reporting, and short-term priorities.

Taken together, these evolutions lead to a practical outcome: A&R tends to become less a single figure and more a variable combination of functions (scouting, development, positioning, platform interfacing, data reading), which recomposes itself differently depending on mission and organizational context. This is the premise that the thesis tests empirically: if levers and bottlenecks change, it is plausible that A&R practices and decision logics also change, producing different A&R types rather than a single best practice for this professional figure.

To read these differences in a coherent way, it is useful to adopt a theoretical framework that frames A&R as a cultural intermediary and a decision-maker under conditions of uncertainty, and that treats metrics, playlists, and rankings as judgment devices and evaluation regimes. It is important to underline that, what is presented in the next sections, does not already present the results of my work, but only underlines the evidence of the literature. Indeed, in the next chapters, during the analysis of the interviews, I will test these theoretical evidences, to see what is still valid nowadays, and how the framework has changed from the theory in a market that is developing as fast as ever.

2.3 A&R as cultural intermediary and gatekeeper

Reducing A&R to the sole discovery of talent is misleading for at least three reasons. First, because in the cultural industries talent is not an object measurable once and for all: quality, potential, and fit emerge over time and depend on how a project is produced, narrated, distributed, and supported. Second, because success does not depend only on the work or on the artist, but on social dynamics

of demand that A&R must anticipate, ride, or counter: what works, in fact, is often the result of a relationship between artistic proposal, cultural context, and circulation infrastructures, not the simple sum of individual merits. Third, because A&R makes decisions by coordinating actors (artists, management, producers, label, marketing, radio/PR, platforms) and allocating scarce resources (time, budget, internal attention, release slots), with concrete effects on who gets opportunities and visibility. It is therefore important to underline how A&R also has a mediator role in the space between artist and audience and between artist and organization.

In the music market, there is no direct relationship between production and consumption, but rather a set of figures and practices stand in between. It is in this intermediate space that the concept of **cultural intermediaries** is located. According to Negus (2002), this notion serves to describe a class of professional actors who operate between production and consumption and who, precisely because of their position, contribute to establishing what counts as good and legitimate culture, which styles, scenes, and identities are recognized (or marginalized), and which products acquire visibility and symbolic as well as economic value. In operational terms, Negus argues that cultural intermediaries perform three intertwined functions:

1. **Translation:** they transform an artistic project, often non-standardized and not yet readable, into a language that can be shared and used by other actors (publics, media, platforms, internal departments), building interpretive frames, categories, and narratives (“what it is,” “who it is for,” “in which competitive space it sits”).
2. **Selection:** they filter abundance by deciding what deserves listening, time, and investment, establishing what enters the perimeter of attention and what stays outside.
3. **Legitimation:** they attribute credibility and status to certain projects (and not to others), producing signals of reliability that facilitate access to resources and channels, and thereby contributing to reproducing or reorganizing hierarchies of value in the cultural market.

This definition is coherent with the idea (developed in the literature on cultural industries) that between production and consumption there are zones of mediation in which value is constructed and made comparable: intermediaries’ work is neither neutral nor purely technical, because it affects classifications of taste and criteria of recognition. In the music case, roles such as A&R, marketing, and promotion fall emblematically within this category, but so do more recent roles such as editorial curators and other actors who contribute to making a project positionable and therefore investable.

A second framework coherent with this setup is the production of culture perspective. In short, this perspective argues that what gets created and what manages to circulate does not depend only on the artist, but also on technological, legal, and organizational constraints and enablers, and on how the market is made operational (for example through segmentations, targets, and evaluation practices). A&R therefore is forced to mediate also between these constraints and the artistic proposal, making a project not only interesting on an aesthetic level, but also positionable and therefore investable (Peterson & Anand, 2004).

The coherence between the theory of cultural intermediaries and the dynamics of the music market is emphasized by Negus, who proposes shifting attention from artists to the practices of recording-

industry workers. The author in fact shows how pop music is produced through a network of mediators, such as A&R, marketing, promotion, and commercial functions, who operate in the space between production and consumption and make the encounter between music, the public, and the market possible. In this sense, A&R is a cultural intermediary that is intra-organizational and inter-organizational at the same time. A&R operates in a hybrid ecosystem in which large groups coexist with a multiplicity of small and medium-sized firms, with increasingly complex relationships between formal and informal actors of the cultural market: intermediation therefore happens both through inter-organizational modes and intra-organizational ones. Intra-organizational, because it has to transform aesthetic impressions and fragmented signals (a demo, a live performance, a micro-community, feedback from a producer or manager) into a proposal that is understandable for other departments: what project it is, who it is for, in which competitive space it sits, what narrative can support it, which moves make sense in the short and medium term. Inter-organizational, because mediation also happens between external worlds: local scene and label, artists and management, producers and the network of professionals, promotional and distribution partners. In this sense, intermediation also happens in order to meet the different actors involved in a project's path and to thin out the tensions that inevitably arise between non-coinciding objectives and constraints. A typical tension, recurring in the cultural industries, is the market-art compromise (or, more generally, between creativity/knowledge and commercial objectives): maintaining the singularity and expressive coherence of a project, while at the same time making it readable and sustainable within organizational and market logics. Coherently, the cultural industries literature highlights that managerial activity often consists in making this tension negotiable without reducing it to a simple calculation. A&R operates exactly on this fault line: it has to preserve elements of singularity and difference (what makes a project recognizable) and, at the same time, make them readable as a value proposition in a competitive and segmented context. This explains why many A&R decisions are not formulated only as "it's good / it's not good," but as questions of compatibility (fit): fit with a certain audience segment; fit with a catalogue; fit with a positioning, that is, with the image and the competitive space in which the project can be placed; and fit with a credible development trajectory, without betraying what makes the project recognizable (Hesmondhalgh, 2019). A further implication is that translation work also changes as a function of ownership structure and of the way the organization defines what counts as success and as a sufficiently good profit level. In small labels or independent realities, the sustainability of creative work and the artistic quality of the project can very often count more than the absolute maximization of revenues, as long as the label manages to keep its business sustainable. When firms become large and subject to financial pressures (as in the case of majors, for example because they are publicly listed), acceptable profit tends to be measured against global investment alternatives, and expectations of returns and growth are higher. In practice, when a large label has to justify resources to shareholders or financiers, it tends to raise the bar on success metrics, payback times, and scalability, reducing tolerance for slow or highly experimental paths.

This tension has been read in a radical way as incompatibility between authenticity and commercial logic, but in historical perspectives (for example on rock) art and commerce often turn out to be integrated and co-produced. The case of Prince is often recalled in the literature precisely because it makes the concrete dynamics of this tension observable. At the height of his conflict with Warner Bros (1990s), Prince contested not so much the idea of being successful, but organizational control over his production: timing and quantity of releases, promotional priorities, and, above all, the fact

that rights and the economic value of recordings were structured in a way that limited his artistic autonomy. The friction exploded when the label tried to govern the rhythm and market placement of the output, while he claimed a logic closer to creative continuity and control over the project's identity; the protest became public with the name-symbol and with the word "slave" on his face, as a rhetorical gesture to make contractual asymmetry visible (Morris, 2022). In this sense, the case shows that the market-art compromise is not an abstract debate: it takes shape in clauses, rights, calendars, budgets, and positioning choices that can come into conflict with the artist's creative trajectory. Many majors have historically tried to manage the problem by creating semi-independent internal units (imprints or business units) to which they grant greater creative freedom while still remaining within overall constraints.

The point, therefore, is that translation (understood according to the cultural intermediaries theory) that A&R has to apply concerns not only the outside (public/market), but also the inside (organization). This work emerges with particular clarity when one observes the concrete infrastructure of decisions: word of mouth, referrals, professional reputations, who knows whom. In sociological terms, these dynamics can be read as a network-market: a market coordinated by relationships and trust rather than by fully impersonal signals. Here trust can take different forms, from an "advance of trust" (one bets on a project because it is guaranteed by a credible subject) to trust built over time through repeated interactions and observable performance (Karpik, 2010).

In sum, intermediation is not only personal taste; it is a social and organizational work that transforms incomplete signals and uncertainty into shared categories, defensible reasons, and operational moves. Precisely because it makes a project readable and coordinable, this translation is not neutral: it orients attention and expectations and prepares the ground for subsequent investment choices. This is where the gatekeeping lens becomes indispensable: once the reasons for why a project should be signed or not have been built, another crucial step of the creative development process and product distribution lies in the choice of who gets in, under which conditions, and with which resources.

The gatekeeping lens brings attention to another dimension: A&R as a node that selects and allocates. Defining gatekeeping as "selection + allocation of opportunities and resources" serves precisely to avoid the censorship misunderstanding. By "censorship" people often mean an act of prohibition or suppression (typically with a moral, political, or legal motivation): someone prevents a content from existing or circulating. Here, instead, the point is different and more everyday: A&R works in a context of scarcity, in which listening time, release slots, budgets, internal attention, and access to channels and partners are extremely limited resources. Consequently, the filter here lies in choosing where to put resources and which paths to make possible. Concretely, doing gatekeeping can mean: deciding which projects enter a short list and which do not; which become priorities (and therefore receive team, budget, marketing, pitching) and which remain under observation; which tracks are chosen as the single and which are not; whether to activate a collaboration/producer or not; whether to push a project on certain channels (radio, playlists, sync) and with which timing, and many others. These choices are crucial because they substantially change the probability that that project will obtain visibility and future signals (Caves & Richard, 2000). This logic is also visible in classic preselection models: in a system in which many aspirants compete for few listening and investment slots, A&R is the first critical node inside the

organization. Only a small part of artists manage to get in contact with an A&R, and only a minority of the projects listened to are then proposed to label decision-makers: A&R therefore functions as the first gate between abundance of supply and scarcity of industrial attention.

The economic rationale that makes this selection so central is also linked to a classic concept in the cultural/creative industries literature: the economic properties of “cultural texts” (that is, works/recordings as informational goods), characterized by high fixed costs in producing the “first copy” and low marginal costs (even more so in digital) of reproduction and distribution. Once the break-even point is passed, each additional unit tends to generate very high margins; consequently, hits can compensate for many flops. This incentivizes strategies oriented toward maximizing audience and reinforces the centrality of circulation and marketing as an integral part of value-making projects (Hesmondhalgh, 2019).

However, A&R gatekeeping does not coincide only with the moment of its own choice about which project to start developing and which not. After selection there are further relevant doors: which tracks become priorities, which collaborations are activated, which budgets are unlocked, which narrative is supported, which channels are covered first. Every choice of this type is, in fact, a form of allocation: it changes the probability that a project meets opportunities (professional, promotional, distributive) and therefore that it generates future signals. In the literature, this overall dynamic is described as a matching process between talents and complementary inputs (producers, studio, promotion, access to networks), often governed by contracts and sequential decisions that treat projects as options to be exercised or abandoned over time. Moreover, in many organizations there is also a final formal gate: the signing/greenlight by the label (or by final decision-makers). In this phase, A&R no longer decides only whether a project is interesting, but builds a case that makes the signing defensible: contractual conditions, level of investment, return expectations, acceptable risks, and alignment with portfolio priorities.

The key point from this long analysis is that gatekeeping and intermediation are intertwined: on the one hand translation produces the reasons that make a selective choice defensible; on the other hand selection produces the conditions that make certain development paths possible (or impossible). If an A&R sees potential, they are not simply describing an intrinsic quality: they are often anticipating an investment and coordination path that, if activated, will make that potential more likely.

Framing A&R as both an intermediary and a gatekeeper helps to read the interviews while avoiding some recurring caricatures. On the one hand, the idea that the industry is only determined by capital and that choices are mere impositions from above; on the other, the mirror idea that everything depends on individual creativity and that the industry is a simple neutral container. To these two simplifications a third is often added, more recent: the promise that numbers can replace judgment and turn decisions into automatic optimizations. These two lenses do not deny the three caricatures, but they put them back to scale, showing that each captures a piece but, taken alone, becomes misleading. The intermediation lens scales down both the idea that everything is dictated by capital and, conversely, that everything is the merit of individual creativity, because it puts at the center the everyday work of translation and coordination that concretely makes the encounter between project and public possible. The gatekeeping lens, instead, shows that A&R choices are always

also allocation of scarce resources and access to channels: therefore they produce real distributive effects on who gets opportunities, visibility, and investment.

If A&R is, at the same time, mediation and selection under constraints, then it becomes particularly interesting to analyze how an A&R decides when signals are imperfect and results are not predictable *ex ante*. This leads directly to the theme of Chapter 2.4: uncertainty as a structural condition and trade-offs as the ordinary form of decision action in the creative industries.

2.4 Decision-making under uncertainty and trade-offs

If Chapter 2.3 has shown that A&R operates as intermediary and gatekeeper, the next step is to clarify how decisions happen when the outcome is not predictable *ex ante*.

In the creative industries uncertainty is not noise that can be eliminated with more data or more competence: it is a structural property of the work, linked both to the nature of the product (symbolic goods, evaluated also in terms of meaning) and to market dynamics (scarce attention, fashions, competition and imitation). Within this framework, the A&R decision resembles less an optimizable choice and more an act of commitment toward a path, which must be continuously revised and defended over time (Caves & Richard, 2000; Hesmondhalgh, 2019).

Uncertainty does not mean only risk. In the economic and organizational literature it is useful to distinguish between risk, that is, when there are reasonably estimable probabilities of a damage, and uncertainty, that is, when outcomes are not calculable because there are no stable bases to assign reliable probabilities. In many A&R decisions, the critical variable is not how much a project will return, but whether and how it can build an audience, a recognizable position, and a sustainable trajectory, in an environment in which the context changes more and more quickly (genres, platforms, languages, attention spans, competition). Talking about uncertainty, therefore, serves to avoid an overly financial reading of the problem: it is not only about estimating an ROI, but about making decisions on projects whose quality and demand are partially indeterminate and socially constructed (Knight, 1921).

This condition is often summarized, in the creative industries, by the “nobody knows” principle: before a work is put into circulation and supported by complementary inputs (production, promotion, distribution, access to networks and channels), it is not possible to know reliably what the audience response will be. The point is not that choices are arbitrary, but that even competent decisions must operate with incomplete signals, because economic and symbolic value manifests only after the work has been made visible and tested on the market (Caves & Richard, 2000).

This uncertainty is accompanied by a typical economic structure of cultural markets: the majority of books, records, or films do not reach relevant results, while a significant share of revenues concentrates on a few hits (80/20, 90/10 dynamics, or even more concentrated, depending on the

reference cultural industry). Some examples reported in the literature show extreme orders of magnitude: in 1999, 0.03% of releases would have generated 25% of sales; on iTunes, 0.001% of tracks would have generated 15% of sales (Hesmondhalgh, 2019). The point here is not the precision of the number, but the logic it describes: hits are rare, they drive the portfolio, and they often arrive as surprises, sometimes after repeated rejections, while many “correct” projects fail.

Making the picture even harder is the volatility of demand: fashions and artists can quickly go from “hot” to obsolete, while others break through unexpectedly (Wikström, 2019). Moreover, the sector is particularly vulnerable to technological disruptions coming from consumer electronics and IT (today, in particular, platforms and digital infrastructures), which can change in a short time the rules of discovery, access to publics, and conditions of monetization.

In a context of structural uncertainty, the problem is not only choosing, but building defensible decisions. A&R rarely decides alone: it has to convince other actors (internal and external) that a bet makes sense, that an investment is proportionate, that risks are acceptable, and that the path is coherent with the priorities and constraints of the moment. Consequently, a crucial part of decision work consists in transforming impressions and fragmented signals into a plausible narrative: what did we see? why does it matter? which alternatives are we discarding? which conditions make this choice reasonable? This is an organizational coordination process, not a simple aesthetic judgment (Caves & Richard, 2000). This need to build defensible decisions intertwines with organizational politics and conflict. The literature on cultural industries describes how functions such as marketing often try to exercise influence diplomatically and through dialogue and negotiation; at the same time, the relationship with artists can also include heated confrontation: arguing as part of maintaining a productive working relationship, especially when objectives, timing, and expectations do not coincide.

Here it is useful to recall two general ideas from organizational theory. The first is bounded rationality: in complex and ambiguous environments, actors cannot consider all alternatives nor calculate all consequences, so they instead build solutions that are sufficiently good with respect to objectives and constraints, using heuristics, practical rules, and comparisons. Applied to A&R, this means that the decision does not arise from a total analysis of the market, but from a set of professional shortcuts (comparables, network signals, reputations, minimum criteria) that make it possible to act without pretending to know what will happen with certainty (Simon, 1955). The second is sensemaking: when situations are equivocal, actors must first build a shared meaning of what is happening, and only then decide how to act. In A&R work this emerges when a project is promising but ambiguous: the question becomes “What is it? How do we read it? In which space do we position it? Which story makes it understandable?”, that is, an interpretive activity that precedes (and makes possible) the choice. In practice, many decisions are not the arrival point of an objective evaluation, but the point at which an organization closes an interpretation and decides to commit to a well-defined path (Weick, 1995).

If uncertainty makes a fully objective evaluation impossible *ex ante*, then decision action often takes the form of trade-offs: choices between objectives in tension, which cannot be maximized simultaneously. In the interviews, this point is central because it makes it possible to read differences in practice not as personal styles, but as different ways of managing recurring tensions

under different constraints. A first typical trade-off concerns exploration vs exploitation: investing in new, risky, or not-yet-codified proposals (exploration) or concentrating resources on more known formulas and on already validated trajectories (exploitation). In A&R work this translates, for example, into the balancing between scouting emerging scenes and consolidating more predictable assets; between long-term bets and optimization of results in the short term; between experimentation and repeatability. This tension is not a defect: it is an ordinary condition of portfolio and learning in competitive markets (March, 1991). A second trade-off concerns artistic singularity vs market readability. By definition, cultural projects also compete on difference: to emerge in crowded markets it is not enough to work technically, but it is necessary to propose a recognizable trait (sound, imagery, identity, reference scene) that allows publics and media to distinguish them from very similar alternatives. In other words, a fundamental question for A&Rs becomes “what makes it different and why should someone choose it (and recognize it) over something else?”. At the same time, to be supported by organizations and circulation infrastructures, projects must be understandable, frameable in a value proposition, and positionable. The intermediation discussed in 3.1 works precisely on this point: A&R tends to protect distinctive elements, but it has to make them interpretable and sellable within a segmented market and within organizations with time and budget constraints (Karpik, 2010; Hesmondhalgh, 2019). A third trade-off, very frequent in operational decisions, is short vs long term. A&R choices can be guided by immediate objectives (a release that works, pitching results, fast signals) or by the progressive construction of a career (repertoire development, identity, fanbase, credibility). The tension is concrete because organizational constraints often push toward the short term, while the construction of symbolic value and of a relationship with the public requires time. (Caves & Richard, 2000; Hesmondhalgh, 2019).

A recurring organizational response to risk and uncertainty is portfolio logic: that is, signing or putting many projects into development, observing the public’s initial reactions, and then concentrating budget/promotion on the few that seem to work. This logic is named in the literature as the “shotgun” logic, described as “throwing mud at the wall to see what sticks” (Wikström, 2019). In this scheme, A&R therefore also acts as a portfolio manager: overproduction and catalogue breadth increase the probability of intercepting enough hits to sustain the flops. In analytic terms, portfolio logic makes visible an additional trade-off: breadth vs intensity of investment. Expanding the funnel (many small bets) increases exposure to unexpected opportunities, but reduces the depth with which each project can be followed. Concentrating resources on a few projects, instead, increases coherence and execution, but exposes more to the risk of choosing “wrong” ex ante.

Other recurring trade-offs are:

- **speed vs accuracy** (deciding early in order not to lose opportunities vs collecting additional evidence)
- **catalogue coherence vs opportunism** (protecting an editorial line vs chasing contingent signals)
- **internal criteria vs external criteria** (building judgment inside the team vs relying on devices and external actors) (Karpik, 2010; Hesmondhalgh, 2019).

Talking about trade-offs is not meant to say what the right choice is, but to compare how A&Rs decide: which tensions they see, what they sacrifice, and with which consequences on decisions and resources. If decisions happen under uncertainty and take the form of trade-offs, it becomes understandable why organizations and professionals resort to tools that reduce ambiguity and coordinate expectations: comparisons with similar cases, reputations, endorsements, but also metrics, rankings, and playlists. These tools do not eliminate uncertainty, but they make it more manageable and communicable. The crucial point, however, is that such devices are not neutral: they influence what is seen as promising and what is considered defensible, and they can produce conformity and feedback loops. Moreover, the cultural industries literature insists on a structural attention to power: investment decisions, visibility decisions and, today, also algorithmic or playlisting ones happen within relations of ownership and control. Questions like “who owns the media/platforms?” and “which interests are privileged?” therefore become part of reading numbers, and not an external theme.

For this reason, the next step (2.5) analyzes metrics, playlists, and rankings as judgment devices and as practical evaluation regimes that enter A&R’s decision routines.

2.5 Metrics, playlists, and rankings as judgment devices and evaluation regimes

The spread of metrics, charts, and playlists should not be read as a simple increase in available information, but as a change in the way the industry coordinates attention and value. When products are singularities (i.e., difficult to compare *ex ante* because quality and demand are uncertain and partially constructed), markets work through devices that orient evaluation: they help to choose what to listen to, what to believe, and what to invest in (Karpik, 2010). In A&R work, these devices take a hybrid form: they include social and professional signals (reputations, referrals, endorsements), but also measured and made-public signals (streams, skips, saves, charts) and external selection infrastructures (editorial and algorithmic playlists).

In the lexicon of markets for singularities, this centrality of devices is expected: when products are multidimensional and quality is uncertain, choice cannot be reduced to a one-dimensional calculation or to perfect information. In these markets, judgment devices become crucial, i.e., socio-technical and symbolic devices such as professional networks and word of mouth, brands and denominations, critics and guides, charts, up to playlists and rankings, because they reduce opacity and fill a cognitive deficit for consumers and operators. Following Karpik, judgment devices provide formatted and credible knowledge that dissipates (in part) the opacity of markets for singularities and makes it possible to judge and choose. In this sense, they contribute to qualifying the product (i.e., establishing, from a certain point of view, which qualities matter and for which reasons) and, indirectly, to selecting/hierarchizing also the publics that rely on them (Healy, 2011; Schenk, 2021). Using them implies a delegation grounded in trust, but it also

introduces discretionary power (of those who manage the device) and therefore possible asymmetries and distortions. By delegating judgment to a device, that device decides (in part) what becomes credible and what deserves attention. For this reason, often it is not only products that compete, but also the channels that orient choice (i.e., what is made visible and reliable). Applied to music: playlists, metrics, and reputations are not neutral numbers; they are channels that orient visibility and therefore demand, influencing which projects become credible and investable.

In contexts of uncertainty, a central function of numbers is cognitive delegation: transforming part of the evaluation into a synthetic output (an indicator, a chart position, a percentage growth) allows multiple actors to coordinate more easily even if they have different knowledge and backgrounds (Karpik, 2010). Another function is comparability: metrics and rankings make very different projects comparable along a few dimensions (attention, listens, conversion), enabling fast priority decisions (what enters a funnel, what becomes a priority, what receives budget, etc.). This does not eliminate the uncertainty described in 3.2, but it makes it manageable because it produces communicable signals and, above all, signals that are usable in internal processes (Maasø & Hagen, 2020). The critical point is that, precisely because they are public and comparable, these indicators do not only describe reality, but influence it. The literature on the reactivity of measures shows that when actors are evaluated through rankings or public indicators, they tend to adapt strategies and behaviors to improve position or the indicator, generating feedback effects (behavior changes because of measurement) (Espeland & Sauder, 2007). In the music industry this means that a metric is not only a measure, but also a lever, because it can orient investment choices, release tactics, promotional budgeting, and even product characteristics (such as choices of format and length, shareable moments, and attention optimization).

To understand why metrics are so pervasive, one clarification is needed: the cause is not only greater data availability, but the real lever is that organizational culture and infrastructures make data-driven evaluation the default, i.e., they incorporate measurement into daily routines through dashboards, reports, meetings, and targets. This process is often discussed with the term *datafication*, i.e., the tendency to transform aspects of experience and behavior into trackable, aggregable, and monetizable data (Kitchin, 2014). In A&R work, datafication does not replace judgment, but it changes which judgment can be defended: it makes it easier to justify choices that “sit in the numbers” and harder to sustain paths that require time or that produce weak signals at the beginning. This has at least two consequences. The first is that metrics become a form of governance: they order priorities, define what deserves attention, and establish minimum thresholds of evaluation. The second is that they increase the speed of decision cycles: continuous (almost real-time) measurement facilitates fast adjustments, but it can also push toward overly immediate readings, especially when numbers are used as a proxy for quality rather than as situated signals (Beer, 2016).

In the shift to streaming, playlists and rankings also work as pre-selection, orienting what gets listened to even before listening produces data. Here it becomes central to make a distinction, namely between signal and exposure effect. A number can reflect organic interest (signal), but it can also be the outcome of a previous visibility decision (exposure) that causes that number. The ambiguity is structural. An entry into a relevant playlist (editorial or algorithmic) can generate growth in streams and followers; however, that growth, in turn, can be read as evidence of demand

and trigger further internal investments (marketing, pitching, budget). The analytic risk is to confuse cause and effect: interpreting as proof of quality or potential what is, in part, the result of a device that has already allocated attention, leads to wrong and unfounded decisions (Maasø & Hagen, 2020). This does not mean that playlists are false or that data do not matter: it means, however, that the value of numbers depends on the demand from which they are produced and on the circulation context. Research on Spotify highlights how platform-centric promotion and curation are not neutral: the emphasis on the playlist format and the centrality of proprietary playlists make the platform an intermediary with selection power, with consequences for the ecosystem (who gets listens, under which conditions) (Prey et al, 2020). Consequently, an A&R who reads numbers in a robust way tends to ask: “is this trend driven by exposure (playlist/editorial/paid) or by repeated and converted behaviors (saves, returns, retention)?”; and above all “how replicable is it outside the device that generated it?”.

Regarding judgment devices, it is important to underline one thing: these are not simple measurement tools, but they participate in defining what becomes visible and credible. For this reason, judgment devices are not neutral, but there is an additional reason as well: they are part of structures and relations of power. At a general level, the literature on algorithms insists that algorithms are almost never only technical objects, but incorporate institutional choices (what counts as relevant, what gets amplified), the company’s own interests, and produce hierarchies of visibility (Gillespie, 2014). In the case of music platforms, this means that to make a correct evaluation, it should not be based only on measuring listens, but also on evaluating which listening is made possible and which is made probable. From here follows an operational point for reading the interviews: when A&Rs talk about numbers, it is crucial to understand whether they treat them as objective truths or as outputs of a socio-technical system with its own interests, constraints, and levers.

For Chapter 4, this lens serves to turn metrics and playlists from context into an analytic object. In the interviews, the analysis therefore does not look only at which numbers are cited, but above all at:

- which devices are considered reliable and why;
- how the signal/exposure distinction is handled;
- which thresholds and routines transform numbers into decisions.

The other aspects useful to make interviews comparable (mission, trade-offs, time horizons, constraints, etc.) are summarized in the list of sensitizing concepts in 2.6, to avoid repetitions.

In sum: metrics, rankings, and playlists help coordinate decisions in uncertain contexts, but they also produce reactivity, conformity, and feedback cycles. The result is that data are never only a snapshot, but part of the mechanism that generates visibility and makes some choices more defensible than others. For this reason, section 2.6 closes the chapter by translating this framework into a set of observable sensitizing concepts, useful to code empirically how A&Rs talk about signals, devices, and trade-offs without assuming typologies a priori.

2.6 Synthesis: sensitizing concepts and what to observe empirically

This section closes the conceptual framework by turning it into a practical guide for the analysis. The idea is to use some concepts as sensitizing concepts, i.e., initial lenses that help to notice relevant differences and to ask better questions while coding the interviews. In this sense, the concepts are not meant to “prove” a theory, but to orient attention toward what could matter empirically (Blumer, 1954; Bowen, 2006).

Below, the list is deliberately short and observable: for each concept, I indicate what to look for in the empirical material.

1. **Mission and decision logic:** what is treated as the real objective? Which idea of value justifies choices?
2. **Scope and process gates:** where do decisions sit, entry into the funnel, shortlist, prioritization, budget allocation, single choice, release timing, up to greenlight/signing? Who decides what, and under which conditions?
3. **Signal sources and their hierarchy:** which signals are searched for, and in what order?
4. **Judgment devices considered reliable:** which devices are used to make evaluation credible (reputation, endorsements, charts, playlists, internal dashboards), and why are they considered reliable? What is trust based on?
5. **Signal vs exposure effect:** when numbers or playlists are mentioned, does the interviewee try to distinguish between organic demand and demand generated by visibility? How do they infer causality?
6. **Explicit trade-offs and decision criteria:** which tensions actually recur? What gets sacrificed, and how is the compromise made defensible?
7. **Time horizon and decision rhythm:** how much time is given to a project before judging it? Do they work in fast cycles (real-time) or in longer windows? What role do tests, micro-releases, and iterations play?
8. **Risk management and portfolio logic:** do they reason in terms of many small bets and a few that become priorities? Are there thresholds, options to exercise or abandon? How does this change with budget and context?
9. **Internal coordination and productive conflict:** how are divergences between A&R, marketing, finance/management, and other players composed? Which arguments hold in meetings? When is friction described as a normal part of the work?
10. **Constraints and infrastructural power:** to what extent does dependence on platforms and external channels emerge?

In sum, the concepts above do not automatically produce a typology on their own, but they help make different interviews comparable and let the dimensions that distinguish A&R practices and decision logics today emerge from the material. The next chapter details how these concepts enter the coding process.

Chapter 3 – Methodology

This chapter describes how the empirical research was conducted and how the interviews were transformed into evidence useful to answer the research questions. The objective of this chapter is to make the path carried out replicable and defensible: to clarify the design choices, how the sample was built, how the material was collected and managed, and through which analytical steps we arrived at the construction of dimensions and, subsequently, of a typology. In coherence with the typology-first setting, the chapter also makes explicit the measures of traceability and control (audit trail, negative cases, assignment rules) that make it possible to assess the quality of the interpretation.

To avoid overlaps, the chapter does not anticipate the results: the data structure, the emerged dimensions and the typology are presented in Chapter 4.

3.1 Research design and rationale

The research questions of this thesis call into play phenomena that, by their nature, remain largely internal to processes: operational routines, criteria used to evaluate an artist/project, decision sequences, trade-offs managed under conditions of uncertainty and constraints.

In an interpretive qualitative perspective, organizational practices are not reducible to observable events or measurable performances: they incorporate meanings, judgment criteria, and practical rules that are often implicit. For this reason, actors can be treated as knowledgeable agents, that is, competent sources to reconstruct what they do, with which priorities, and which trade-offs they consider in the different phases of the work. Moreover, in the context of streaming and datafication, a relevant part of the information is mediated by opaque infrastructures (dashboards, rankings, algorithmic affordances) and by asymmetries of access and competence: understanding what one does with data and how one interprets it requires observing the situated use of those signals in decision-making processes, not only their existence as numbers (Maasø & Hagen, 2020). For this reason, the research adopts a qualitative design based on semi-structured interviews with professionals who carry out A&R activities and, when useful, with closely adjacent roles that influence or constrain such decisions.

The choice of semi-structured interviews is motivated by a specific methodological tension: combining comparability and openness. By comparability we mean the possibility of comparing cases on a common core of observable aspects: this stable core prevents every interview from becoming a non-comparable story. At the same time, the thesis aims to bring out real differences that often appear in the details: exceptions, edge cases, conflicts between signals, changes in

priorities, compromises between objectives. A fully structured guide would risk flattening these differences, while a guide that is too free would reduce comparability. Semi-structuring makes it possible to maintain continuity of coverage and to use targeted probes to clarify meanings, sequences, and trade-offs when unexpected elements emerge (Kvale & Brinkmann, 2015).

The unit of analysis, in this design, is not the individual as such nor the organization as a macro category, but the decision-making practice: what actors do recurrently and what they treat as good evidence or a good criterion in the different phases of A&R work. This choice is crucial because it reduces two risks that would make the typology fragile. The first is to confuse differences of context (e.g., resources, structure, market) with differences of type: in the typology-first logic, context is treated as a modifier that conditions constraints and operational possibilities, but it does not automatically define a type. The second risk is to build an impressionistic taxonomy, that is, descriptions based on general impressions without empirical anchoring and without verifiable criteria. To avoid it, the objective is to isolate differences that are visible in the work and that can be compared across cases. In other words, the typology must not capture psychological traits of the interviewee, but configurations of practices and criteria that are reconstructible starting from episodes, examples, and systematic comparisons.

From the point of view of analysis, the design is coherent with a Gioia-informed approach: starting from the language of informants and from their categories (first-order) to arrive, in a traceable way, at comparable conceptual themes (second-order) and aggregate dimensions (aggregate dimensions), making visible the passage from data to concepts through a data structure (Gioia, 2012). In parallel, the analysis is conceived as iterative: coding decisions, definitions, and analytical focuses are refined as new questions and ambiguous cases emerge, through memos and audit trail (Locke et al., 2022).

Finally, it is necessary to clarify what kind of inference is realistic to expect from this design. The thesis does not seek statistical generalizations, but transferable concepts and configurations: categories clear enough to be applied and discussed in similar contexts, and rooted enough in the data to remain defensible. In this perspective, quality depends on the alignment between RQ, sample, interview guide and analysis, and on the transparency with which interpretive choices are documented, so that the reader can assess credibility, transferability and traceability, aspects that will be operationalized in section 3.5 (Lincoln & Guba, 1985; Maxwell, 2013).

The next section (3.2) describes how, starting from this design, a maximum-variation sample was built and how the boundaries between type and context were defined already in the data-collection phase.

3.2 Sampling and interviewee profile

In a typology-first research, sampling is an integral part of the analytical strategy: if the objective is to bring out and compare ways of working and decision logics, then the sample must be built in a way that makes plausible differences in processes visible. In other words, sampling must expose the project to sufficient heterogeneity to allow the construction of dimensions and configurations, but also maintain a common focus that makes the cases comparable.

Consistently with this need, the research adopts a purposive sampling logic: participants are selected such that they can provide dense information on the phenomenon and therefore allow reconstructing and comparing decision-making processes (Palinkas, 2015). Within this logic, the main strategy is maximum variation: including different profiles and contexts to observe how routines, criteria, and trade-offs change (or remain stable) as operating conditions vary (Patton, 2015).

The sample includes informants who:

1. carry out or have carried out A&R activities in the operational sense adopted in the thesis;
2. can describe concrete practices (workflow, criteria, examples of decisions, trade-offs, exceptions) and not only general opinions;
3. operate in positions that allow a credible reconstruction of the process (decision perimeter, interactions with other stakeholders, recurring constraints);
4. cover sufficiently different contexts to support the maximum-variation logic.

Alongside “pure” A&R roles, the research also included adjacent roles when they directly affect A&R decisions and constraints. The inclusion of these profiles is not meant to dilute the object of study, but to make some critical steps of the process more observable. In the analysis phase, these roles are treated as contextual information and triangulation, not as substitutes for observing A&R work. Recruitment took place through a combination of targeted outreach (contacts actively sought by the author) and referral (introductions suggested by interviewees or contacts).

Since the objective of the thesis is to distinguish types of practices without confusing them with the conditions in which they take place, for each case some contextual variables were noted as background information (not as type definitions). Concretely, these are a few descriptive notes useful for comparison (e.g., level of process formalization, availability of resources and analytical tools, breadth of the decision perimeter, dependence on external gatekeepers such as platforms or media).

Finally, the final sample size ($N = 10$) was not fixed *ex ante* as a mechanical threshold, but guided by a criterion of informational adequacy: data collection continued as long as additional interviews contributed to generating new substantive codes or to testing emerging boundaries and definitions; when instead they produced mainly confirmations and marginal refinements, the sample was considered adequate. In the literature this criterion is often discussed in terms of saturation, but it should be understood cautiously and tied to the purpose of the analysis. A systematic review of empirical studies indicates that, in relatively homogeneous populations and with well-delimited

objectives, code saturation tends to be reached within a narrow range of interviews (about 9–17) or focus groups (about 4–8); beyond these thresholds informational gains often tend to be decreasing (Hennink & Kaiser, 2022). In the present study the field is circumscribed (the Italian recording market and A&R and related roles), with a controlled heterogeneity intentionally introduced (different organizational setups and operational constraints) to support the typology-first logic. Consistently with indications from the literature, data collection was planned on an order of magnitude of about 8–12 in-depth interviews and adapted in itinere if, in the final phases, new substantive codes or new configurations still emerged that were relevant for the typology.

To allow the reader to assess transferability and limits, the results section will include a summary table of the sample with: main role (A&R vs adjacent roles), type of organization (macro-categories), market/country of reference, years of experience (bands), interview mode (remote/in person), and an anonymous code for citation in the text. The description of the sample will be sufficiently informative to contextualize the cases, but calibrated to preserve anonymity when the role–organization–event combination would make an informant identifiable.

The next section (3.3) details protocol, ethics, and material management, that is, how this sample was actually interviewed and how the corpus (i.e., the complete set of transcripts and notes) was built in an auditable way.

3.3 Data collection: protocol, ethics, material management

This section describes how the empirical corpus was built in an auditable way: how the interview guide was designed and used, how the interviews were conducted, and how consent, anonymity, and data security were managed.

Data collection was based on a semi-structured interview protocol, consistent with the need to maintain comparability across cases without giving up descriptive depth. The guide was designed starting from the research questions and the operational definition of A&R (Ch. 1), while the concepts discussed in the conceptual framework (Ch. 2) were used as lenses to formulate observable questions and to ensure coverage of the relevant aspects, without turning into categories to be confirmed.

Consistently with the Gioia Methodology, the guide was conceived as a living document: a stable core guarantees comparability across cases, but the wording of probes and the order of exploration can be adjusted when gaps, ambiguities, or unexpected concepts emerge. The objective is not to rigidly standardize the interviews so as to always ask the same questions, but to maintain sufficient openness to intercept new concepts, documenting any modifications and the reasons that motivated them through memos and an audit trail, consistently with Gioia and Locke.

The interviews were conducted in the period between October 2025 and January 2026, with a typical duration of 45 minutes, in person or in remote mode. Before the interview, participants were informed of the purpose of the study and the perimeter of data use; during the conversation, I maintained a role of facilitating the narration: I helped the informant reconstruct episodes and decision-making steps (asking for examples, clarifications, “before/after” sequences, and criteria used), avoiding as much as possible leading or evaluative questions that could have pushed toward desirable or overly general answers. When authorized, the meeting was audio-recorded to guarantee accuracy and reduce dependence on selective notes. In parallel, essential notes were collected during and immediately after the interview and, when necessary, contextual elements that emerged and ambiguities to be checked in subsequent interviews were also annotated. Participants provided informed consent to participation and, separately, to recording (when present), with the possibility to interrupt the interview or to request the removal of specific parts. Data management was set up to reduce the risk of re-identification, particularly relevant in small samples and in professional contexts in which role and organization can make single individuals identifiable. These choices are aligned with data protection principles applied to research in the European context (European University Institute, 2019) (European Data Protection Board, 2021) and with established good practices for the anonymization of qualitative data (UK Data Service, 2022). The recordings (when present) were transcribed into text format and access was limited to the subjects involved in the research; the files were stored in protected archives.

The next section (3.4) describes how the corpus built in this way was analyzed with a Gioia-informed and iterative approach, up to the construction of dimensions and the typology.

3.4 Data Analysis

This section describes how the interview corpus was transformed into a set of empirical dimensions and, subsequently, into a configurational typology. Two principles guide the analysis. The first is traceability: each step must remain linked to the textual evidence, avoiding opaque interpretive jumps. The second is coherence with a typology-first setting: types are not decided *ex ante*, but emerge as recurring configurations of observable differences across cases.

The analysis began with a phase of immersion in the corpus: a full reading of the transcripts and interview notes, with preliminary annotations on recurring terms used by informants, described decision sequences, signals/metrics cited as relevant, stakeholders entering decisions, and constraints and trade-offs recalled. In this phase, the objective is to create familiarity with the material and to build a first inventory of observable elements that can become codes. The coding unit was chosen in a pragmatic way: text segments short enough to retain a specific meaning (e.g., a described practice, an explicit criterion, an example of a decision), but long enough to keep the context needed to correctly interpret the statement.

The first coding phase was conducted with a Gioia-informed approach oriented to the first-order: initial codes were formulated as much as possible close to the language of the informants, privileging descriptive labels and, when useful, in vivo formulations (i.e., words/phrases used by the interviewees). The methodological point here is the balance between adherence to the language and categories of the informants and abstraction (translating those descriptions into comparable and more general concepts): in fact, in the first-order phase the researcher acts almost as a qualified chronicler, minimizing the imposition of external categories; abstraction happens explicitly in the subsequent phase (second-order), where the researcher uses their theoretical competence to recognize patterns and build comparable concepts without losing anchoring to the informants' spoken language (Gioia, 2012). Operationally, this phase produced a high number of initial codes, then progressively cleaned up through merging overlapping codes and separating codes that were too broad into more specific sub-codes; coding was accompanied by analytical memos that record decisions and doubts about why a fragment was coded in a certain way, which alternatives were plausible, which conditions seem to make a practice more or less likely, and which questions remain open.

The second phase was oriented to the second-order: first-order codes were grouped and reformulated into more abstract conceptual themes, capable of comparing cases while maintaining the link with the evidence. Here the main translation happens: from local descriptions to concepts that allow systematically comparing practices and decision logics. The passage from first-order to second-order was managed iteratively, going back to the text multiple times to check whether the theme actually captured what the fragments have in common. Iterativity is therefore treated as daily work of theorization, not as mechanical repetition of steps.

Once the main second-order themes were stabilized, they were further synthesized into a smaller number of aggregate dimensions: higher-level dimensions that describe structural differences in ways of working and deciding. The aggregate dimensions are not already narrated results, but the analytical format with which the thesis prepares comparison and the typology.

To make the passage from data to dimensions transparent, a data structure was built that shows the chain: fragments and first-order codes (informants' language) → second-order themes (conceptual interpretation) → aggregate dimensions (synthesis); the data structure will be presented in Chapter 4 together with the results.

The passage from the list of dimensions to a typology requires one more step: it is not enough to say how cases differ, one must also understand which differences tend to appear together and describe a recognizable way of working. In this thesis, therefore, each case is read as a combination of dimensions: when certain characteristics recur together in multiple cases and have a coherent logic, they are treated as a type (i.e., a way of working comparable with the others). Operationally, this passage translates into very concrete work in three moves:

1. **Compiling a “cases × dimensions” matrix:** a table is created in which each row is an interviewee/case and each column is a dimension (and, when needed, a sub-dimension). For each cell, one does not insert judgments, but evidence, such as a cited episode, a key

phrase, or an example of a decision. The function of the matrix is to make comparison systematic (Miles et al., 2014).

2. **Looking for recurring combinations:** by looking at the matrix, one checks whether there are co-occurrences, i.e., when some dimensions tend to present together in multiple cases. In this phase, provisional groups are annotated, united by the same key dimensions.
3. **Giving a name and a boundary to the groups (provisional types):** a group becomes a type only if it satisfies three simple criteria: internal coherence (cases in the group really show the same pattern), distinctiveness (another group differs by combinations, not by details), and a sense-making logic (it is possible to plausibly explain why those traits stay together, i.e., what rationality/practice holds them together) (Kluge, 2000; Doty, 1994).

A practical corollary of this setting is to recognize limited diversity: the space of theoretically possible combinations is much wider than what is observable in the data. In this case, empty cells are not a failure of the analysis, but information that some combinations do not emerge (or are rare) in the observed field; for this reason, the analysis avoids forcing typologies that artificially fill all possible combinations.

To avoid the typology being absorbed by organizational categories (e.g., “major vs indie”), some contextual variables were recorded and explicitly read as modifiers. In practice, during the construction of the case $\times$ dimensions matrix, for each case relevant contextual elements were annotated (resources, formalization, market, role specialization, dependencies on external gatekeepers) and their possible role in facilitating or limiting the expression of certain practices. The analytical rule was the following: context can explain why a certain way of working is more or less practicable, but it cannot by itself define what type of decision logic is in place. In operational terms, a contextual element was considered a modifier when it changes the intensity, tools, or constraints with which a practice is carried out, without changing its underlying logic; on the contrary, it was considered an indication of typological difference only when it is accompanied by a distinct pattern of criteria, sequences, and trade-offs.

In sum, the analysis proceeds from the text (first-order) to conceptual themes (second-order), to aggregate dimensions (data structure), and finally to the construction of recurring configurations that constitute the typology. Since this typology must be usable and defensible, the next step is to make explicit the safeguards of rigor and control: how boundaries between types are established, how hybrid cases and negative cases are handled, and which assignment rules make the classification replicable. These aspects are addressed in section 3.5.

3.5 Quality and rigor criteria

In a typology-first thesis, quality depends on two very simple questions: “what I claim is really supported by the interviews?” and “if another reader had the same rules, would they arrive at a similar classification?”. To answer these questions, the thesis adopts quality criteria typical of

qualitative research (credibility, transferability, dependability, confirmability) and translates them into a few concrete safeguards: traceability of the analytical path, internal consistency checks, systematic use of edge cases, and explicit rules for assignment to types (Lincoln & Guba, 1985):

1. **Traceability:** the passage from data to concepts remains controllable thanks to two main elements: the data structure, which makes the first-order → second-order → aggregate dimensions path visible, and a cases x dimensions matrix, which forces comparison across cases with homogeneous criteria and reduces a posteriori selection of only the convenient evidence.
2. **Internal consistency:** since the analysis is conducted by a single coder, consistency is supported by targeted re-checks on portions of the corpus and by stabilizing the operational definitions before building the typology. **
3. **Edge cases (negative/borderline):** cases that do not fit well into a configuration are not ignored, but are used to understand whether a dimension is too broad, whether boundaries between types are too vague, or whether a real hybridization exists. Operationally, when a case “does not fit”, one goes back to the original fragments and decides whether to refine definitions/indicators or to treat the case as hybrid. This can make the typology less clean, but more defensible.
4. **Assignment rules:** to make the typology applicable, each type is described with explicit rules based on a simple distinction: core traits (those without which the type does not hold) and peripheral traits (frequent but not indispensable), plus exclusion criteria. In practice, assignment follows four steps: (1) compile the case x dimensions matrix with textual evidence; (2) verify whether the core traits are present; (3) use peripheral traits as reinforcement; (4) if there is conflict, apply exclusion criteria and a tie-break rule (priority to traits linked to the decision logic and to the management of trade-offs). If no assignment is convincing, the case is classified as hybrid or not classifiable and used to refine boundaries and definitions (Doty, 1994).

In closing, Chapter 4 presents the outputs of the analysis (data structure, emerged dimensions, and the final typology in the form of type cards), including for each type a synthesis of core/peripheral traits and assignment rules, so as to make the classification usable also in the chapter on operational implications.

Chapter 4 – Results

This chapter collects and organizes the empirical results of the research: it shows what emerges from the interviews and how this evidence is made readable in an orderly way. For the procedural details of the analysis, see Chapter 3, where the method and criteria are already discussed.

The reading order is designed to accompany the reader step by step. It starts by clarifying who makes up the sample and which contexts it covers (4.1), in order to make transparent the empirical base on which the whole chapter rests. From there, the text offers a map of recurring themes (4.2), useful to understand what the interviewees actually talk about and within which boundaries. On this basis, the results are then expressed as observable empirical dimensions (4.3), defined in an operational and comparable way across cases. Only at that point does the chapter move on to the construction of the typology (4.4) and to its description in a compact and comparable form through the type cards (4.5). Finally, the reading brings into focus the trade-offs and constraints that emerge in the different ways of working (4.6) and closes with a synthesis that prepares the operational translation of the following chapter (4.7).

To guide the reading, the chapter relies on some artifacts (figures and tables) that function as a map and as a support for control:

- [Tab. 1 - Sample snapshot]
- [Fig. 1 - Data structure (Gioia method)]
- [Tab. 2 - Themes overview]
- [Tab. 3 - Dimensions' synthesis]
- [Tab. 4 – Cases x dimensions matrix]
- [Fig. 2 – Configuration scheme]
- [Tab. 5 – Type cards' summary]
- [Tab. 6 – Trade-offs]

Before going into the substance of the themes and dimensions, the chapter therefore opens with a snapshot of the sample: the next section clarifies who the interviewees are and which contexts it covers.

4.1 — Snapshot of the sample and of contexts as modifiers

This section describes the empirical base from which the results of this chapter derive, describing which voices were collected, from which organizational contexts, and with which role scope. The

aim is to make transparent who is speaking before presenting themes, dimensions, and patterns that emerged from the data, while preserving anonymity.

Chapter 4 is based on 10 semi-structured interviews conducted between November 2025 and January 2026. The conversations took place both in person and remotely and have a duration between 39 and 59 minutes. The table below [Tab.1] summarizes, for each case, a macro profile of the role and the context, together with a few descriptive notes at low risk of re-identification:

ID	Role	Context	Interview Date	channel	Length (min)	Context notes
I01	A&R	Indie	nov-25	in person	51	Small-medium team, small roster, strong label's brand identity
I02	Editor	Music media	nov-25	in person	39	
I03	Manager	Booking	dec-25	remotely	54	Freelencer with medium resources
I04	Manager	Collecting	dec-25	remotely	41	
I05	Manager	Federation	dec-25	in person	42	Marketing communication and insight manager
I06	Manager and A&R	Indie	dec-25	in person	41	Team of freelance A&R professionals
I07.1	A&R	Major	jan-26	in person	52	
I07.2	Manager	Major	jan-26	in person	59	Data science and innovation
I08	A&R	Indie	jan-26	remotely	56	
I09	A&R	Major-independent	jan-26	remotely	59	

Tab. 1 – Sample snapshot

The sample includes, as a central core, profiles attributable to the A&R function in different label contexts: indie (I01, I06, I09), major (I07), and a hybrid major-independent context (I09). Alongside these core voices, adjacent positions are included that affect the scope and the operational interfaces of contemporary A&R: a music media outlet (I02), a booking agent (I03), a profile linked to collecting/rights (I04), a famous industry federation with responsibilities for communication and insights (I05), and a managerial profile in the data/innovation area in a major context (I08). This composition makes it possible to observe decision-making practices and

constraints not only within the A&R perimeter, but also along some of its most frequent interdependencies (live, rights, media, data/processes).

In the rest of the chapter, context labels (e.g., major, indie, hybrid, and the non-label functions) are treated as modifiers: they serve to qualify constraints, resources, and intensity of work (for example, availability of tools, access to data, division of labour, organizational interfaces), without being assumed as an automatic explanation of the observed behaviours. In other words, context is used to interpret under which conditions certain practices tend to appear or to become more or less salient; it is not used to define the types *a priori*, nor to infer differences in effectiveness or performance between organizations.

For confidentiality reasons, the interviews are identified with codes (I01–I09) and the roles and contexts are reported in aggregated form, avoiding references to companies, cities, or potentially unique combinations. Context notes are kept only when they do not increase the risk of re-identification. The full criteria for data handling and anonymization, together with the ethical rationale, are described in Chapter 3. Code I07 has two sub-codes, namely I07.1 and I07.2, since these are two distinct separate interviews but done in the same company. When I therefore refer to I07, I consider a single context, in which I combined the data taken from the two separate interviews.

Having clarified the sample and the role of context as a modifier, paragraph 4.2 shows how themes and dimensions emerge from the data, through the data structure and the overview of the themes that feed the analytical dimensions of Chapter 5.

4.2 — Data structure (Gioia-informed) and themes overview

This section presents two artifacts that make the path from the data to the analytical categories traceable: Fig. 1 (data structure) and Tab. 2 (themes overview). The goal is twofold: to show, in a defensible way, the leap from data to concepts, up to dimensions, and to provide the reader with a map of the empirical material, clarifying which themes emerged, what they include, and which dimensions they feed.




Fig. 1 – Data structure (Gioia method)

Fig. 1 organizes the results on three levels: first-order examples (which are two different examples of first-order codes associated to the corresponding second-order theme), second-order themes, and aggregate dimensions. The first-order concepts are labels that adhere to the language and practices described by the interviewees, meaning descriptive units close to the data. The second-order themes group these labels into clusters of observed practices and/or decisions, without transforming them yet into general theory. The aggregate dimensions, finally, synthesize the themes into comparison dimensions: they do not already define types in and of themselves, but they make the cases comparable and prepare the steps to the following sections (4.3 and 4.4), where the dimensions are defined operationally (indicators, poles, boundaries) and then used to build the typology. The overall logic is to make visible where each dimension comes from and to reduce the arbitrariness of the interpretive leap. When quotations are presented, they are reported in short and targeted fragments, since the full excerpt remains traceable through the interview ID. When relevant, the refinement of themes and codes occurred iteratively (e.g., renamings, merges, splits, and clarifications of boundaries) to increase internal coherence and adherence to the data.

Tab. 2, instead, provides an operational reading for each second-order theme: it indicates what it includes, which dimensions it feeds, and a reference to an anchor quote with the related ID. The table therefore has a mapping function: it allows the reader to orient themselves among the themes without anticipating the typology, and it prepares section 4.3, where the dimensions are defined and made usable to compare the cases.

Second-order themes	Includes	Recalled dimensions	Quote (ID)
Scouting	discovery channels (social media, networks, referrals), reputation/track record of contacts, local signals (scenes/“movements”), and minimum criteria for initial attention	Scouting and access to signals	the first contact now, I won't deny it, happens on social media (I01)
Creative potential evaluation	aesthetic/artistic judgement (originality, “emotion,” repertoire, live/performance), repertoire/studio work, and the search for distinctive elements within the noise	Evaluation and interpretation	there is the whole qualitative factor which is often described a bit like that thing you feel in the area (I02)
Individual and team evaluation	attitudinal fit, maturation/ego, ability to work with others, clarity of goals and the artist’s “vision,” trust/reliability, and behavioural red flags	Evaluation and interpretation	I see patterns in everything, even in the way people speak, the way they dress, the way they respond (I06)
Metrics and interpretation	which metrics matter and how to read them (active vs passive, source of streams, longitudinal patterns, cross-platform correlations), rejection of a mono-KPI approach, and the translation from metric to action	Evaluation and interpretation	When an artist/project intrigues us, we try to see it live whenever possible: it provides a basis for evaluating it (I02)
Project planning and development	roadmap and path, organic development vs chasing results, repertoire management (having songs ready, timing), fanbase/superfan and products, and project vs single	Development and career time horizon	Even in marketing: we no longer take generic testimonials, but people with a specific reference cluster. (I02)
Workflow and decision governance	how decisions are made (collective processes, unanimity, roles), routines and processes, speed of execution, allocation	Operating model and decision governance	Out of seventy projects, there are many active projects (I07)

	of time/attention, and internal portfolio management (chats/meetings/rhythms)		
Stakeholders and coordination	interfaces and alignment with marketing/promo/commercial teams, managers/suppliers, legal/administration, platforms/distribution, and multi-counterparty coherence	Operating model and decision governance	We have an entire team available to support our artists in all aspects, from start to finish (I01)
Constraints and trade-offs	budget/runway, risk and tolerance for error, roster size, scarcity of time, organisational sustainability, the “art/culture vs KPI” trade-off, selection vs trawling, and tailoring vs scale	Constraints, risks, capacity and trade-offs	Even when they build an emerging project, their work still has to stay within budget (I02)
AI	actual use (often marginal), desired uses (operational copilot, dashboard aggregation, quick reporting), perceived limits (it does not decide), risks/rights (training data, deepfakes), and the regulatory framework as context	AI and rights	You feed him a demo and then he does his own analysis... honestly, I don’t need it (I01)

Table 2 – Themes overview

The aggregate dimensions are presented in Fig. 1 and referred to in Tab. 2 exclusively as mapping labels; their full operational definition (indicators, poles, boundaries, and examples) is developed in section 4.3.

Below, a synthetic reading of the second-order themes is proposed, following the same principle of the table:

1. **T1– SCOUTING**

Includes: discovery channels (social media, networks, referrals), reputation/track record of contacts, local signals (scenes/“movements”), and minimum criteria for initial attention.

Recalled dimensions: Scouting and access to signals.

Anchor quote: *“the first contact now, I won't deny it, happens on social media”* [CIT: I01]

2. **T2 – CREATIVE POTENTIAL EVALUATION**

Includes: aesthetic/artistic judgement (originality, “emotion,” repertoire, live/performance), repertoire/studio work, and the search for distinctive elements within the noise

Recalled dimensions: Evaluation and interpretation

Anchor quote: *“there is the whole qualitative factor which is often described a bit like that thing you feel in the area”* [CIT: I02]

3. **T3 – INDIVIDUAL AND TEAM EVALUATION**

Includes: attitudinal fit, maturation/ego, ability to work with others, clarity of goals and the artist’s “vision,” trust/reliability, and behavioural red flags

Recalled dimensions: Evaluation and interpretation

Anchor quote: *“I see patterns in everything, even in the way people speak, the way they dress, the way they respond”* [CIT: I06]

4. **T4 – METRICS AND INTERPRETATION**

Includes: which metrics matter and how to read them (active vs passive, source of streams, longitudinal patterns, cross-platform correlations), rejection of a mono-KPI approach, and the translation from metric to action

Recalled dimensions: Evaluation and interpretation

Anchor quote: *“When an artist/project intrigues us, we try to see it live whenever possible: it provides a basis for evaluating it”* [CIT: I02]

5. **T5 – PROJECT PLANNING AND DEVELOPMENT**

Includes: roadmap and path, organic development vs chasing results, repertoire management (having songs ready, timing), fanbase/superfan and products, and project vs single

Recalled dimensions: Development and career time horizon

Anchor quote: *“Even in marketing: we no longer take generic testimonials, but people with a specific reference cluster.”* [CIT: I02]

6. **T6 – WORKFLOW AND DECISION GOVERNANCE**

Includes: how decisions are made (collective processes, unanimity, roles), routines and processes, speed of execution, allocation of time/attention, and internal portfolio management (chats/meetings/rhythms)

Recalled dimensions: Operating model and decision governance

Anchor quote: *“Out of seventy projects, there are many active projects”* [CIT: I07]

7. **T7 – STAKEHOLDERS AND COORDINATION**

Includes: interfaces and alignment with marketing/promo/commercial teams, managers/suppliers, legal/administration, platforms/distribution, and multi-counterparty coherence

Recalled dimensions: Operating model and decision governance

Anchor quote: *“We have an entire team available to support our artists in all aspects, from start to finish”* [CIT: I01]

8. **T8 – CONSTRAINTS AND TRADE-OFFS**

Includes: budget/runway, risk and tolerance for error, roster size, scarcity of time, organisational sustainability, the “art/culture vs KPI” trade-off, selection vs trawling, and tailoring vs scale

Recalled dimensions: Constraints, risks, capacity and trade-offs

Anchor quote: *“Even when they build an emerging project, their work still has to stay within budget”* [CIT: I02]

9. T9 – AI

Includes: actual use (often marginal), desired uses (operational copilot, dashboard aggregation, quick reporting), perceived limits (it does not decide), risks/rights (training data, deepfakes), and the regulatory framework as context

Recalled dimensions: AI and rights

Anchor quote: *“You feed him a demo and then he does his own analysis... honestly, I don’t need it”* [CIT: I01]

Once the themes to dimensions map is defined, section 4.3 moves from an organized catalogue to comparison: for each aggregate dimension, it provides an operational definition, empirical indicators, poles and boundaries, so as to compare cases consistently and prepare the typological construction (4.4).

4.3 – Candidate dimensions

This section defines the empirical dimensions that emerged from the analysis as observable traits that vary across cases in A&R practices and decision logics. The aim is twofold: to make cases comparable through shared operational rules, and to transform the themes discussed in 4.2 into dimensions that can be qualitatively assigned to each case through explicit indicators, poles, and boundaries. The dimensions below are not already the types: in 4.3, archetypes or configurations are not built, but rather the comparative building blocks that make it possible to do so in 4.4 are built.

Each dimension is described in a repeatable way with (1) an operational definition, (2) a set of observable indicators, (3) two descriptive extreme poles, (4) “include/not include” boundaries to reduce arbitrariness and avoid overly generic dimensions, (5) minimum evidence (quotes and a mini episode), and (6) notes on opacity and borderline cases.

In terms of quality, I apply a simple cross-cutting rule: dimensions must be assignable (indicators and boundaries) and distinct from each other (boundaries/tests), avoiding shortcuts based on a single signal.

Where metrics, playlists, or platform signals appear, the section indicates the source when it is distinguishable; when the signal to action chain is opaque or partial, the opacity is stated. The distinction between dimensions that are candidates to structure the typology (CORE) and

dimensions that more often act as modifiers/conditions (SECONDARY) is presented as provisional and will be tested in 4.4 based on the discriminating power of the dimensions over recurring configurations and on their ability to separate borderline cases.

Table 3 offers an overall synthesis of the dimensions; the cards for each dimension are presented after the table in numerical order (D1–D9) to make consultation immediate and reduce reading ambiguity. This order does not imply a hierarchy among dimensions: the underlying logic is that each card defines an observable trait with its boundaries and evidence; in 4.4 the dimensions will then be considered in combination, regardless of the sequence in which they have been presented here. A modifier (AI) is recorded separately because it can cut across different cases without constituting a typological axis in itself.

ID	Dimension	Operational question	Pole 0 (extreme A)	Pole 2 (extreme B)	Provisional status
D 1	Depth vs volume	Does it work on a few projects, taking care of them a lot, or on many with diluted attention?	Few projects, high care	Many projects, diluted attention / urgencies	TO BE DECIDED
D 2	Slack / economic risk	Can it absorb mistakes, or does one mistake compromise sustainability/cashflow?	High slack, experimentation	Low slack, tight constraint	SECONDA RY
D 3	Value compass	Does it decide based on identity/cultural value or on performance/KPIs/market?	Cultural/identity-led	KPI/market-led	CORE (candidate)
D 4	How it uses data	Does it triangulate and contextualize, or does it use a decision rule on a dominant signal?	Cautious / multi-signal	Mono-signal/number decision rule	CORE (candidate)
D 5	Main “proof”	What counts more: live/active fanbase or platforms/algorithms/social?	Live/active	Platform-led	CORE (candidate)
D 6	Time horizon	Does it reason in years/paths or on short windows/spikes?	Long-term	Short-term / reactivity	CORE (candidate)
D 7	Governance (who decides)	Is the decision collective, multi-function, or constrained by role/objectives/reporting?	Collective / multi-function	Constrained by role/targets/reporting	CORE (candidate)
D 8	Coordination (dependencies)	How many interfaces/dependencies absorb day-to-day work?	Low dependencies	High dependencies / overload	SECONDA RY
D 9	Runway / cliff	Are checkpoints far apart or close, with penalties/negative action?	Long runway / distant cliff	Short runway / close cliff	SECONDA RY (candidate)
M 1	AI (modifier)	Level of use/expectation towards AI (0–2/NA)	Absent / NA	Decision-support	SECONDA RY (modifier)

Tab.3 – Dimensions’ synthesis

Here below are the cards for each dimension:

D1 — Depth vs Volume

Operational definition. The dimension describes whether the A&R works on a few projects with a high level of care and follow-up (0), or on a large portfolio with inevitably diluted attention and greater standardization/urgencies (2).

Observable indicators. Number of projects “top of mind” at the same time and intensity of follow-up; signs of overload that reduce focus on the core; small roster by choice vs large portfolio; high artist-per-A&R ratio as a proxy of saturation.

Poles. 0 = depth/focus (few artists/projects, high care, greater tailoring). 2 = volume/big portfolio (many artists/projects, diluted attention, standardization/urgencies).

Boundaries (include / not include).

- **Include:** breadth of the managed portfolio; trade-off between number of projects and intensity of care/follow-up.
- **Not include:** how much coordination is required (D8); who decides (D7); short/long time horizon (D6).

Evidence.

- **Exemplar quote:** “*Our roster is very small by choice... we never go beyond 15/16 artists... and of these 3/4/5 are always emerging or debut. A debut artist has no numbers by definition.*” [CIT: I01]
- **Exemplar quote:** “*betting on too many artists and not focusing on fewer artists*” [CIT: I08]
- **Mini episode:** an operational rule of a roster cap to preserve depth and follow-up, including emerging artists who “by definition” do not have consolidated numbers; rationale: protecting quality of care against dilution. [CIT: I01]
- **Supporting evidence (scale/ratio):** “*In multinationals the number of artists is much larger but A&Rs remain relatively few: the A&R/artist ratio is high.*” [CIT: I01]

Notes / opacity / edge cases. Risk of collapse with D8 when overload is used generically: here what matters is the portfolio vs care trade-off, not the number of dependencies/stakeholders. Also, the excerpt “A&R/artist ratio is high” should be read as a high artists-per-A&R ratio (A&Rs relatively few). [CIT: I01]

Status. TO BE DECIDED — CORE if it discriminates recurring configurations beyond context; SECONDARY if it mainly reflects organizational constraints.

Assignment test. If a case describes few projects with close follow-up and high care → towards 0; if it describes many projects with standardization and urgencies → towards 2.

D2 — Slack/economic risk

Operational definition. It measures the ability of the organizational context to absorb mistakes (0) versus a financial constraint such that one mistake compromises sustainability/cashflow and operational decisions (2).

Observable indicators. Language of error tolerance; references to imminent expenses and margin; budget constraint as the dominant criterion; economic sustainability as the “needle on the scale”; projects stopped due to lack of coherent investment; attention to “breaking even” also on emerging artists.

Poles. 0 = high slack (mistakes absorbable, more room for experimentation). 2 = low slack (tight constraint, mistake potentially compromising).

Boundaries (include / not include).

- **Include:** resources/cashflow and ability to absorb gaps/mistakes.
- **Not include:** personal cautiousness; time-related deadlines/penalties (D9).

Evidence.

- **Exemplar quote:** *“It’s not like you make a major fail because you make a mistake and sign... giving too much money to an emerging artist.”* [CIT: I06]
- **Exemplar quote:** *“We have to make some expenses in the next three months and I’m not so convinced we can make them all and still keep a margin.”* [CIT: I06]
- **Mini episode:** “many projects” are described as being stopped because there is no possibility of an economic investment coherent with what management requires; the stop decision is guided by the sustainability constraint (the “needle on the scale”). [CIT: I08]

Notes / opacity / edge cases. D2 can interact with D7 (governance) when financial constraints enter the decision perimeter (e.g., reporting and “staying afloat”), but the measured trait remains the capacity to absorb mistakes, not who decides.

Status. SECONDARY (modifier) — a resource condition that modulates choices and room for experimentation.

Assignment test. If in the case a mistake is described as “absorbable” and does not threaten continuity → towards 0; if a mistake is described as potentially compromising for cashflow/budget → towards 2.

D3 — Value compass

Operational definition. It indicates whether decisions are primarily oriented by the cultural value/identity of the project and repertoire (0), or by performance/KPIs/market assets and monetization (2).

Observable indicators. Emphasis on music/art/culture and identity; refusal to maximize a metric as the objective; emphasis on monetizability and segments (e.g., “superfan”); use of certifications/recognized assets; language of scale and “dragnet fishing” as a volume-oriented logic; choice of formats/products for economic conversion.

Poles. 0 = cultural/identity-led. 2 = KPI/market-led.

Boundaries (include / not include).

- **Include:** what is considered value and what orients priorities and trade-offs.
- **Not include:** style of reading data (D4); preferred main evidence/proof (D5).

Evidence.

- **Exemplar quote:** “*My real objective... to remember that we are talking about music/art/culture... not maximizing the percentage of active streams.*” [CIT: I01]
- **Exemplar quote:** “*Superfan... more monetizable... spends double on buying physical music... and over 60% more... to attend concerts.*” [CIT: I05]
- **Mini episode:** the creation of product variants (e.g., multiple vinyls/covers) is described as linked to the valorization of a high-conversion target (superfan), with an explicitly economic rationale (maximizing purchases/engagement). [CIT: I05]

Notes / opacity / edge cases. The compass can also be stated as a guardrail (“not maximizing a metric”), and it should be kept distinct from D4: here it is not about how to read data, but what really counts in the choice.

Status. CORE (candidate) — because it directly orients priorities and decision trade-offs.

Assignment test. If the case puts identity/cultural value at the center and rejects KPI-maximization logics → towards 0; if it puts monetizability, assets, and high-conversion segments at the center → towards 2.

D4 — How it uses data

Operational definition. It describes the style of data use/interpretation: caution and multi-signal triangulation (0) versus shortcuts and decision rules based on a single dominant signal/number (2).

Observable indicators. Language of correlation and context; rejection of a single metric as “the answer”; citing multiple patterns in parallel; examples of a “toxic metric” when it becomes a target; heuristic rules tied to a “minimum signal in numbers”; statements that a decision based on a single pattern is a mistake.

Poles. 0 = cautious / multi-signal. 2 = decision rule on a dominant signal (mono-signal/number).

Boundaries (include / not include).

- **Include:** method of reading/interpretation (triangulation, decision rules).
- **Not include:** which evidence weighs more (D5); time horizon (D6).

Evidence.

- **Exemplar quote:** “Numbers on social and on platforms have to be interpreted by putting them in correlation... there isn’t an answer based on a single metric.” [CIT: I06]
- **Exemplar quote:** “I mean, if I see an artist who has been putting out music for 15 years and there isn’t a minimum signal in numbers that something is working, it’s clear that I think about it ten times before going on.” [CIT: I01]
- **Mini episode:** skip rate is used to show how a metric, if turned into a single objective, directly deforms creative/production choices (e.g., removing intros) and therefore alters the object of evaluation. [CIT: I01]

Notes / opacity / edge cases. In the material, the “single-metric” pole often emerges as a negative example (“this is a mistake”) rather than as a claimed practice; useful to build boundaries and guardrails. [CIT: I06]

Status. CORE (candidate) — check in 5.4 its independence from D5 (risk of confusion between method and proof).

Assignment test. If the case talks about correlating signals and rejects decision rules on a single number → towards 0; if it describes rigid rules on a single number/pattern → towards 2.

D5 — Main proof

Operational definition. It indicates which evidence carries the most weight in decisions: live and an active/converted fanbase (0) versus platform/algorithm/social signals and digital performance (2).

Observable indicators. Ticketing as a key data point; searching for live proof when a project “sparks curiosity”; statements of live priority vs social numbers; distinction between passive (playlist-driven) and active streams; use of stream source and playlist type; UGC/TikTok creations and coefficients/correlations; platform alerts on spikes (e.g., catalog).

Poles. 0 = live/active. 2 = platform-led.

Boundaries (include / not include).

- **Include:** the dominant evidence that weighs most in judgement.
- **Not include:** style of reading/triangulating data (D4).

Evidence.

- **Exemplar quote:** *“The most important data we have nowadays is linked to ticket sales.”* [CIT: I03]
- **Exemplar quote:** *“If 80–90% of listens comes from programmed audience (editorial/algorithmic playlists)... it’s a passive stream.”* [CIT: I01]
- **Mini episode:** when an artist/project “sparks curiosity”, it is described as an operational practice to try to see them live as a basis for evaluation, reducing the weight of social numbers alone. [CIT: I02]

Notes / opacity / edge cases. When platform signals are used, the source is sometimes distinguishable (“source of streams”, playlist type), but the translation into action can remain partial or opaque (e.g., interpretation of spikes). [CIT: I01]

Status. CORE (candidate) — because it defines which evidence drives decisions across cases.

Assignment test. If the decisive proof is live/ticketing/fan activation → towards 0; if the decisive proof is platform/playlist/social performance → towards 2.

D6 — Time horizon

Operational definition. It captures the temporal mental model: building over years and paths (0) versus short windows, reactivity, and quick post-release evaluations (2).

Observable indicators. Language of “long paths” versus “hit”; references to trust over multiple years; requests for a 5–10 year vision; benchmarking on the first days of performance; tight monitoring of the single; activation of campaigns/actions following events or spikes.

Poles. 0 = long-term (career building as the unit). 2 = short-term (short windows, quick tests, reactivity).

Boundaries (include / not include).

- **Include:** temporal mindset that structures decisions and routines (years vs short windows).
- **Not include:** institutional runway/cliff (D9).

Evidence.

- **Exemplar quote:** *“I’m interested in long paths, I’m not interested in artists who make the hit.”* [CIT: I06]
- **Exemplar quote:** *“Comparing the first seven days of performance of a track (artist X) with a track (artist Y).”* [CIT: I01]
- **Mini episode:** starting from a signal/event, the A&R describes the need to decide quickly “what campaigns do I run” and “how do I activate myself”, typical of a short-window, reactive logic. [CIT: I01]

Notes / opacity / edge cases. A case can have a long horizon on career but still use early-window benchmarking for tactical decisions; for this reason D6 should be assigned to the dominant mental model, not to isolated single practices.

Status. CORE (candidate).

Assignment test. If the narrative repeatedly uses multi-year paths as the decision unit → towards 0; if it relies on short windows and rapid reactions post release/event → towards 2.

D7 — Governance (who decides)

Operational definition. It describes the mechanism through which relevant decisions are made (signing/renewals/investments): a collective, multi-function process and search for consensus (0) versus a decision constrained by role, objectives, and vertical reporting (2).

Observable indicators. Demos listened to “all together”; consensus/unanimity as a rule; role of multiple functions in deciding; objectives to be met; reports requested by superiors; reporting of revenues/costs and “staying afloat”.

Poles. 0 = collective / multi-function. 2 = constrained by role/targets/reporting.

Boundaries (include / not include).

- **Include:** decision rule (consensus vs reporting/targets) and formal constraints.
- **Not include:** coordination workload (D8).

Evidence.

- **Exemplar quote:** “*For us signing is very much a collective process... we listen to demos all together...*” [CIT: I01]
- **Exemplar quote:** “*He also has to answer to his superiors who ask him for certain reports.*” [CIT: I06]
- **Mini episode:** in the renewal of an expiring contract, the organization must decide “the economics to put on the table”, signaling a decision point that combines economic constraints and renewal governance. [CIT: I07]

Notes / opacity / edge cases. Some cases describe governance as an interweaving of artistic choices and sustainability constraints (“revenues/costs have to stay afloat”): useful to distinguish D7 (decision mechanism) from D2 (slack). [CIT: I06]

Status. CORE (candidate).

Assignment test. If decisions go through multi-function listening/comparison and consensus → towards 0; if they are driven by objectives/reports and top-down constraints → towards 2.

D8 — Coordination (how many dependencies)

Operational definition. It measures how much A&R work is absorbed by coordination with multiple stakeholders/interfaces (2) versus a shorter chain and a reduced set of operational interfaces (0).

Observable indicators. Number of active counterparts; presence of dedicated roles on the partner side (e.g., social/image) that increase plans and constraints; frequency of alignments; time absorbed by emails/chats; perception of “a thousand projects/relationships” to keep together; churn of meetings/ideas with low implementation; need for daily priorities to avoid getting stuck.

Poles. 0 = low dependencies (few interfaces, shorter chain). 2 = high dependencies (many functions/actors, coordination central in the work).

Boundaries (include / not include).

- **Include:** coordination load and number of operational/relational dependencies.
- **Not include:** decision mechanism (D7); portfolio breadth (D1) if not accompanied by many interfaces.

Evidence.

- **Exemplar quote:** *“In a day you might never stop answering emails and not do everything else.”* [CIT: I02]
- **Exemplar quote:** *“You’re inserted in a thousand projects and you have a thousand relationships to keep...”* [CIT: I02]
- **Mini episode:** the informant describes a day absorbed by reactive communication (emails) and by the multiplicity of parallel projects; the implicit decision becomes daily prioritization (“today you do this”) to protect core work. [CIT: I02]

Notes / opacity / edge cases. In the corpus there is no explicit and “positive” description of pole 0 (short chain/few counterparts). For this reason the dimension is treated as a threshold: it is very informative to identify cases with high coordination load (pole 2), while it is less suited to finely distinguish between “low” and “medium” when coordination is not salient in the narrative. In 4.4 it should therefore be used cautiously as a modifier.

Status. SECONDARY (modifier) — it tends to modulate time/quality of work more than defining the decision logic on its own.

Assignment test. If the case describes many interfaces and work absorbed by alignments → towards 2; if the case does not show signals of overload/multiple counterparts and the flow is described as lean → towards 0.

D9 — Validation runway and penalties (cliff)

Operational definition. It indicates how much time is available before an evaluation with a penalty/negative action kicks in: long runway and distant cliff (0) versus quick checkpoints with a close cliff (2).

Observable indicators. Short contracts and commitment conditional on performance; close renewals; formulas such as “if it doesn’t work you go home”; risk of being out if one “creates holes”; long times explicitly stated to reach an output (e.g., an album); economic decisions tied to renewals/checkpoints.

Poles. 0 = long runway / distant cliff. 2 = short runway / close cliff.

Boundaries (include / not include).

- **Include:** institutional/contractual deadlines/penalties that set checkpoints.
- **Not include:** the A&R’s temporal mental model (D6).

Evidence.

- **Exemplar quote:** “*even spending a couple of years to try to make an album*” [CIT: I08]
- **Exemplar quote:** “*They give you a contract for three singles, if they don’t go... you go home, if they go you stay there.*” [CIT: I03]
- **Mini episode:** the “contract for three singles” defines a close checkpoint: after a few releases there is an evaluation and, if negative, a penalty kicks in (exit from the perimeter). [CIT: I03]

Notes / opacity / edge cases. The penalty can also be described in organizational terms (“he risks being sent away if he makes mistakes, if he starts creating holes”), useful to identify a close cliff when it is not formalized contractually. [CIT: I06]

Status. SECONDARY (candidate) — it can become CORE if in 4.4 it structures recurring configurations (e.g., decisions conditional on checkpoints).

Assignment test. If the case is regulated by close checkpoints with explicit penalties → towards 2; if judgement/penalty is distant or not formalized → towards 0.

M1 — AI (modifier)

Operational definition. A modifier that records the level of use/expectation towards AI: 0 absent; 1 productivity/organization; 2 decision support;

Observable indicators. Statements of absence; AI as secretary/agenda; AI for reports/summaries; AI as translation from data → actions (“what do I do with it?”); desire for an end-to-end workflow on a single platform; requests for quick reports for specific channels.

Scale. 0 = absent; 1 = productivity; 2 = decision support;

Boundaries (include / not include).

- **Include:** actual use or operational expectation tied to AI in processes.
- **Not include:** an autonomous typological dimension in 4.3 (here it is recorded as a cross-cutting condition).

Evidence.

- **Evidence quote:** “*We... do not use artificial intelligence tools for data analysis or anything else... nor... for... talent scouting.*” [CIT: I08]
- **Evidence quote:** “*What should AI do... take this info... and turn it into actions: ‘what do I do with it?’*” [CIT: I01]
- **Mini episode:** AI is imagined as operational and organizational support (“an AI secretary... who organizes my agenda”) or as the ability to bring delivery/campaigns/insights into a single workflow; the function is to reduce “hassle” and speed up decisions/actions. [CIT: I02]

Notes / opacity / edge cases. In the material, AI often appears as desire/expectation more than as a consolidated practice; for this reason it is treated as a modifier and not as an axis.

Status. SECONDARY (modifier).

Assignment test. If in the case AI is absent or irrelevant → 0/NA; if it is used for organization and productivity → 1; if it is described as support for diagnosis/actions and decision support → 2.

Once the dimensions are set with their boundaries and minimum evidence, the next step is to observe how these dimensions combine across cases: section 4.4 builds the typology as a set of recurring configurations, making assignment rules explicit and handling borderline cases.

4.4 — Configurational typology

This section builds a configurational typology starting from the dimensions that emerged in 4.3. The goal is practical: to organize the cases along a shared logic, make replicable assignment rules explicit, and openly state when a case is not classifiable in a robust way, instead of forcing it.

The perimeter of the typological construction is restricted to cases that directly describe A&R practices and decisions. For this reason, I04 and I05 are excluded because they are out-of-scope with respect to a “typology-first typology of A&R work” (being, respectively, a company and a federation), but they remain useful only to delimit the field. Within the A&R-only cases, the typology distinguishes one main dimension that structures assignment (the type) and a limited set of secondary dimensions that qualify how that type operates and which trade-offs it tends to generate; context (major/indie, budget, country) can help explain variation, but it is not used to define the type.

The core dimension is D1 (Depth vs Volume). D1 describes the trade-off between portfolio breadth and intensity of care/follow-up (tailoring, close monitoring, ability to keep few projects “top of mind”) versus the management of a large portfolio, with more diluted attention, standardization, and urgencies. By “volume” I do not mean the absolute size of the organization, but the proportional load, which means how many projects, on average, each single A&R has to follow (projects-per-A&R ratio). In the corpus, D1 is supported by direct evidence (for instance, a roster cap rule and the choice of a small roster to protect depth and follow-up) [CIT: I01] and by the reference to the risk of “betting on too many artists” reducing focus [CIT: I08].

To represent the typology, a single-axis form is adopted (Fig. 2). The motivation comes from reading the cases×dimensions matrix: after systematically comparing cases on each dimension (Tab. 4), the only truly clear and cross-cutting variation across the realities turns out to be on D1. Concretely, D1 is also the most polarized dimension in the sample (many cases at the 0 and 2 extremes), while on the other dimensions both poles often do not appear, or a very strong prevalence of one pole emerges, making them less suitable to function as a primary axis without forcing. It is plausible that intermediate configurations exist (tailoring and a “middle” workload), but in the material analyzed here a D1≈1 position does not emerge as a sufficiently stable recurring configuration to be formalized as a third typological position; for this reason, the potential intermediate remains a possibility to validate with additional cases.

Case-Id	D1	D2	D3	D4	D5	D6	D7	D8	D9
I01	0	NA	0	0	2	1	0	2	1
I02	2	2	0	NA	0	0	NA	2	NA
I03	NA	NA	0	0	0	0	2	2	2
I06	2	2	0	0	2	0	1	2	2
I07	2	1	1	0	1	1	0	2	1
I08	0	2	0	1	1	0	0	2	0
I09	0	2	1	0	1	0	0	2	1

Tab. 4 – Cases x dimensions matrix

Tab. 4 summarizes the cases×dimensions matrix used in 4.4 and makes the typological construction traceable. In the current A&R sample, three cases sit on the depth/focus pole (D1=0: I01, I08, I09), three cases on the volume/big portfolio pole (D1=2: I02, I06, I07), and one case is not placeable on the axis due to insufficient evidence (D1=NA: I03), and this is coherent with the nature of the reality represented by the interviewee, i.e., a booking agency, therefore not a proper label. Consequently, the operational positions are P1 (Depth/focus) and P2 (Volume/big portfolio), plus an explicit NC category (not classifiable on D1) when the evidence does not allow a defensible assignment. For P1, the anchors include signals of active protection of depth (small roster by choice, operational cap, emphasis on follow-up and tailoring) [CIT: I01] and the reference to not dispersing across too many projects/artists [CIT: I08]. For P2, the anchors include a high number of simultaneous active projects (*“six, seven, eight projects at the same time”* and *“out of seventy projects... the active projects are many”*) [CIT: I07] and the representation of a multiplicity of projects and relationships to maintain (*“a thousand projects... a thousand relationships”*) [CIT: I02], with further passages compatible with a volume logic and with inputs already “in progress” [CIT: I06].

Once the axis is fixed, the typology becomes configurational through a reduced set of modifiers that explain how the position along D1 translates into practices, criteria, and trade-offs. In 4.4 the modifiers made explicit are four: D9 (Validation runway & penalties / cliff), D5 (Main proof), D7 (Governance), and D3 (Value compass). The other dimensions discussed in 4.3 (in particular D2, D4, D6, and D8) remain tracked in the full matrix, but here they are kept in the background because they do not show an equally “clean” variation or they risk interpretive overlaps, and therefore they do not increase the discriminating power of the typology beyond D1 + the four selected modifiers.

Among the modifiers, D9 is treated as central and as a stress test because it introduces a constraint that can change decision pressure even when D1 remains unchanged: a long runway widens room

for maneuver, while a short cliff compresses evaluation into close checkpoints. This is also visible by comparing two examples: the possibility to “spend a couple of years” to build an output [CIT: I08] versus conditions where permanence is tied to a few releases (“*three singles... if they don’t go... you go home*”) [CIT: I03] or to negative consequences associated with mistakes/“holes” [CIT: I06].

The assignment rules make classification defensible and reduce ex post adjustments. First: only A&R cases enter the typology (the others are out-of-scope). Second: a case is in P1 if $D1=0$ or if at least two depth/focus indicators recur (small roster by choice, operational cap, intense follow-up, explicit tailoring); it is not in P1 if the described norm is a large portfolio with urgencies/standardization. Symmetrically, a case is in P2 if $D1=2$ or if at least two volume indicators recur (many simultaneous projects, overload that compresses follow-up, standardization/urgencies); it is not in P2 if rules that structurally protect depth are present (for instance a cap that limits dilution). When evidence on D1 is insufficient ($D1=NA$) or contradictory with low confidence, the case is not forced into P1 or P2. A rule is also defined for hybrid/transitional cases (even if a stable intermediate position on D1 does not emerge): if a case satisfies strong indicators of both P1 and P2 and prevalence cannot be assessed with at least medium confidence, the case is labeled as hybrid/transition; if evidence remains insufficient, it falls back into NC. Finally, when some modifiers are NA (for example D7 and D9 in I02), this does not block assignment on D1 if evidence on D1 is sufficient, but it limits inferences on “how” that position operates.




Fig. 2 – Configuration typology

In sum, the proposed configurational typology distinguishes a structuring difference in A&R work (D1, single axis) from a limited set of modifiers that qualify how that position operates and which trade-offs it makes more likely, with D9 as the central stress test. Starting from this architecture,

the next section translates each empirical position into an applicable description (type cards), making workflow, signals, stakeholders, and trade-offs explicit, and showing how modifiers (starting from D9) alter the operationalization of the same work logic (4.5).

4.5 – Type Cards

In this section, the typology defined in 4.4 is made more readable and directly usable in the following chapters, by translating the two typological positions (P1 and P2) into descriptive cards written with a homogeneous level of detail and therefore comparable. The setup remains unchanged: one single primary axis (D1 - Depth vs Volume) and four modifiers that influence pressures and operating modes without generating new subtypes (D9 - Runway/Cliff; D5 - Main proof; D7 - Governance; D3 - Value compass). In this section, the aim is to present the two cards following a fixed schema and therefore easily comparable: first the portfolio (D1) and the operational profile, then the workflow and decision points, then the signals that “count as proof” (D5), then governance and stakeholders (D7), then the value compass (D3), then markers and recurring risks, and finally the variation induced by D9 (Runway/Cliff) and a brief anchoring to the sample prototypes.

The first typological position (P1, D1=0: Depth/focus) describes an A&R who chooses to follow a limited number of projects, to be able to dedicate time and follow-up to each one, accepting an opportunity cost in terms of coverage. The difference, therefore, does not depend on the size of the organization, but on the workload per single A&R: few projects-per-A&R make a more tailored work feasible. In the data this emerges both as an intentional roster limit and as a declared trade-off between quantity and care: *“our roster is very small by choice... we never go beyond 15/16 artists”* [CIT: I01]; and the same principle is formulated at the personal-portfolio level: *“if I follow few artists I can take better care of the single project; if I follow dozens ‘you lose a bit everything’”* [CIT: I01]. Operationally, the flow includes a scouting phase that can pass through public signals and relationships (for instance, the first contact “happens on social media” [CIT: I01]) and continues with a more in-depth evaluation that integrates listening, project context, and compatibility with the label’s identity. The main decision point is the explicit choice on signing, formulated as a binary decision (“green light or red light” [CIT: I08]) and, in this type, supported by internal discussion: signing is described as a step in which it is *“fundamental that the whole team is convinced”* [CIT: I01], with dialogue across functions and, when necessary, also an economic-administrative check (*“a part of analysis together with administration”* [CIT: I08]). The goal is not to increase the number of decision-makers, but to increase the probability that the choice is actually implementable.

At the level of signals, P1 does not imply absence of data; it implies rather data used as support and interpreted, not as an automatic criterion. The way “proof” is built (D5) is often based on triangulation and attention to signal quality: “active listens” are read as an indication of fanbase (*“if I see many active listens it makes me say: what a fanbase...”* [CIT: I01]) and numbers are

broken down by looking at the origin of streams (*“one of the metrics I look at most is the source of streams”* [CIT: I01]; *“what type of playlists certain streams come from”* [CIT: I01]). In this logic, some metrics are considered problematic when they push towards optimizations that impoverish the content: skip rate is criticized because *“if we compete to maximize it, intros no longer exist”* [CIT: I01].

The value compass (D3) therefore translates into an explicit guardrail: not confusing music/art with the optimization of a number, *“remember that we are talking about music/art/culture... not maximizing the percentage of active streams”* [CIT: I01]. Coherently, signing is described as a combination of proposal and response: *“I don’t sign an artist who gives a big numerical response, but the artistic proposal does not convince us... the ideal situation is numbers + a compelling artistic proposal”* [CIT: I01]. Even when speed is needed (for instance to get *“a very clear picture in the shortest possible time”* [CIT: I09]), the solution is not to reduce the choice to a single number, but to select a few informative signals and keep them manageable (*“metrics... can never be too detailed, because then you would lose a piece”* [CIT: I09]).

The stakeholders that matter in P1 tend to be few but high-intensity: the internal team, because the decision has to be sustainable in execution, and a selective external network that signals opportunities and projects (for instance managers or other operators who pass on coherent contacts) [CIT: I01], plus non-creative functions such as administration when economic sustainability becomes a concrete constraint [CIT: I08]. As for the decision perimeter, P1 tends to keep with the A&R the artistic evaluation and the setting of project direction, but it can outsource execution components (for example communication/PR) to free up time on the development part (*“the communication of the various projects is outsourced”* [CIT: I08]). The cost of depth is explicit: some phases absorb time and reduce the number of sustainable projects (*“the time effort... means that the number of projects... is generally lower”* [CIT: I09]).

The necessary markers to recognize P1 are (i) a deliberately restricted portfolio as a condition for care and (ii) a signing decision point supported by internal discussion and team alignment, so as to make the choice implementable (without reducing it to a single KPI). [CIT: I01][CIT: I08] Among the traits that are often observed (but that alone are not enough to identify the type) there is a use of data that does not stop at totals, but also considers the origin of streams and signals of “active” audience engagement. [CIT: I01][CIT: I08]. Recurring risks include the risk of underscaling (few projects, smaller experimentation surface) and the risk of closing into a “bubble” or into an area that is too uniform (*“we live a bit in a bubble”* [CIT: I09]; *“a catalogue... of the same genre”* [CIT: I09]). A further risk concerns sustainability: some projects are stopped because they require investments that the context cannot sustain (*“projects... stopped... because there was no... possibility of investment”* [CIT: I08]).

With respect to modifier D9 (Runway/Cliff), P1 tends to express itself differently depending on how much time separates the project from a checkpoint with consequences (long runway vs short runway). In P1, this can be read as a “stress test” of the type: when runway is long, the configuration supports development and long-duration work, coherent with the idea of taking time to build (*“even spending a couple of years to try to make an album”* [CIT: I08]) and with the request for a clearer trajectory (*“a path... more defined”* [CIT: I08]). When instead runway is short and economic-

operational pressure increases, the same logic becomes more cautious: either the selection threshold is raised, or the ambition of some choices is reduced to contain imbalance (*“if a project is cool... but it leans... in a disastrous way... we protect ourselves”* [CIT: I08]). In both cases, no new type is generated: the cost of the trade-offs associated with depth simply increases. In terms of context, P1 is described as more feasible when the workflow is “more stretched out” (for instance in an independent). [CIT: I08]

The main prototypes of P1 in the sample are I01, I08, and I09 (assignment rules and dimensions/modifiers are extensively described in the two previous sections).

The second typological position (P2, D1=2: Volume/big portfolio) instead describes an A&R who follows many projects in parallel: here urgency and standardization are not a preference, but an effect of the load. Also in this case “volume” does not mean “big company”, but a high number of projects-per-A&R. In the data, the idea is explicit (*“we have active six, seven, eight projects at the same time”* [CIT: I07]) and is placed in contexts with high overall numbers (*“out of seventy projects... the active projects are many”* [CIT: I07]). The workflow, consequently, takes the form of a continuous selection of priorities: deciding what to do immediately and what to postpone, with many counterparts and channels competing for attention. This noise is described both as overload (*“they’re writing to me in seven chats...”* [CIT: I07]) and as the risk of absorbing the day into reactive activities (*“in a day you might never stop answering emails and not do everything else”* [CIT: I02]). In this type, the decision point is not only deciding whether to sign or not, but also deciding, every day, which projects to prioritize and how many resources to allocate (time, attention, budget). In P2, care does not disappear, but it becomes more difficult to maintain continuous time for listening and project vision, because the day is often occupied by urgencies and requests. For this reason, the need emerges to protect time dedicated to listening (i.e., actually listening to the music and building a judgement), considered what *“really makes the difference”* [CIT: I02], and to have clear direction on priorities (*“I need someone who tells me: today you do this”* [CIT: I02]) and planning (*“to calendarize it... is already work”* [CIT: I02]). The systemic risk is fragmentation of attention and a less uniform follow-up, because load reduces continuity (*“if days lasted twice as long... more calm and concentration”* [CIT: I07]).

At the level of signals (D5), P2 insists on contextualization: total streams are recognized as the first data point to observe, but they are not enough without interpretation (*“total streams are the first... but... the skill is to contextualize”* [CIT: I07]). The active/passive distinction comes back as a basic grid (*“active stream, passive stream”* [CIT: I07]) to avoid automatic readings of the same numbers, and the rule is to correlate multiple metrics: *“numbers... have to be interpreted by putting them in correlation... there isn’t an answer based on a single metric”* [CIT: I06], *“it’s not a single pattern... this is a mistake”* [CIT: I06]). In some cases, while platform signals remain central, “proof” is shifted to offline activation signals, in particular live interest: *“we rather look at whether there is interest live”* [CIT: I02]. At the same time, social can function as a filter: if one relies almost only on what emerges there, one ends up seeing always the same profiles and the same dynamics, reducing exploration of alternative signals and of proposals outside one’s circuit (*“if you follow only that, you always reiterate the same concept”* [CIT: I02]).

Governance (D7) in P2 tends to include more visible reporting and target constraints: the need emerges to “*answer to superiors... with certain reports*” [CIT: I06] and to “*meet certain objectives*” [CIT: I06], with time pressure (“*you have this amount of time to test employees*” [CIT: I06]). In parallel, choices also appear that pass through inter-department comparison and negotiation: “*beyond the commercial department... other departments... we compare*” [CIT: I07], and some decisions take on an explicitly economic valence (for instance renewals and resource allocation: “*expiring contract... deciding the economics*” [CIT: I07]). The stakeholders that weigh in P2 therefore include both superiors (to report to) and company departments with which to coordinate operational and economic choices. [CIT: I06][CIT: I07] In terms of decision perimeter, P2 oscillates between coordination (project/actors management) and negotiation with functions that control budget and timing: direct responsibilities appear (“*you have budget in hand to manage... a senior A&R*” [CIT: I06]) and explicit trade-offs between projects (“*how much do I invest in one thing compared to another*” [CIT: I07]). Part of the pipeline can also be fed by external figures (freelance) who bring projects already followed, and some promotion/strategy components can be carried out in support or in synergy with other organizational actors. [CIT: I06]

In P2, the value compass (D3) is more tied to economic and organizational constraints: it is said clearly that the A&R does not work only on the artistic side (“*we are a company... but it's not like that...*” [CIT: I07]). Consequently, together with governance constraints, choices tend to often consider sustainability and portfolio management. The typical risk is that the core (listening and vision) is eroded by noise and reactivity; a second risk is standardization, i.e., leaning on shortcuts or channels that produce homogenization (for instance the social filter) [CIT: I02][CIT: I07].

The necessary markers to recognize P2 are (i) a large portfolio that imposes parallel management and (ii) a regime in which priorities and allocation (time, attention, budget) become a continuous decision. Frequent markers (but not indispensable) are reporting/targets as a constraint, communicative noise, and the need to deliberately protect time for listening and vision [CIT: I02][CIT: I06][CIT: I07].

With respect to modifier D9 (Runway/Cliff), P2 also tends to change pace and priorities depending on how much time separates the portfolio from a checkpoint with consequences (long runway vs short runway). In P2, variation is evident because volume amplifies the effect of time and contractual constraints: when runway is short, close decisions increase on continuing/stopping/re-allocating, also through contractual devices and cash constraints (short agreements: “*signed... for three tracks*” [CIT: I06]; quarterly financial planning: “*expenses in the next three months... keep a margin*” [CIT: I06]). In these cases, portfolio management tends to become tighter, because every mistake reflects on many lines at the same time. When instead runway is long, P2 can maintain a development horizon and tolerate non-immediate trajectories (“*it takes you more than a year to understand if a project... works*” [CIT: I06]; “*I'm interested in long paths*” [CIT: I06]); the structural cost of diluted care remains, however, and requires priority and core-protection safeguards.

The main prototypes of P2 in the sample are I02, I06, and I07 (assignment rules and dimensions/modifiers are extensively described in the two previous sections).

To close, the following table quickly summarizes what distinguishes P1 and P2 (markers, way of working, signals, governance, and risks).

Quick comparison		
row	P1 — Depth/focus (D1=0)	P2 — Volume/big portfolio (D1=2)
Necessary markers	Restricted portfolio to be able to follow projects in depth; yes/no signing node supported by team alignment	Large portfolio and continuous parallel management; prioritization/allocation as a recurring decision
Mission / guiding criterion (D3)	Value as project/identity and artistic coherence, with anti-optimization safeguards	Value as a balance between artistic and sustainability/portfolio (company, allocation, trade-offs)
Dominant workflow	Pipeline → in-depth evaluation → yes/no decision supported by the team → intense follow-up	Multiple pipeline → continuous selection of priorities → quick decisions → coordination and parallel management
Proofs and signals (D5)	Triangulation; focus on “active” signals and on origin/context of streams; caution towards problematic metrics	Contextualization; active/passive; multi-metric correlation; live interest as proof in some cases
Governance (D7)	Internal discussion and alignment as an implementation condition (shared signing)	Greater exposure to reporting/targets and cross-functional decisions; negotiation on budget and renewals
Decision perimeter	Evaluation and project direction kept tight; execution can be outsourced to free up time	More coordination tasks; responsibility and negotiation on resources (time, budget, timing)
Typical failure modes	Underscaling; network bubble/homogenization; stop due to investment mismatch	Dilution of care; reactivity and noise; standardization/shortcuts; coordination cost
Stress test D9	Short runway → more caution/reduced ambition for sustainability; long runway → deeper development more feasible	Short runway → close stop/re-allocation (contracts/cash); long runway → long paths but care remains diluted

Tab. 5 – Type cards’ summary

From here, the text moves to recurring problems in comparative form: the trade-offs and constraints that cut across both types, and that become explicit when P1 and P2 are compared on the same set of dimensions/modifiers.

4.6 — Comparative problems / constraints / trade-offs

In this section I discuss the problems, constraints, and trade-offs of each type card identified in 4.5; each trade-off developed in the text is anchored to at least one quote. The aim is to make observable and comparable, between P1 and P2, the recurring tensions that break open in A&R work when objectives and constraints come into conflict. What changes between P1 and P2 is not so much the existence of the same tensions, but how much they weigh and in what order they are addressed. For this purpose, the section presents a comparative table (Tab. 6) and a discursive commentary that explains how each trade-off manifests in the data and how it changes between P1 and P2, using quotes as empirical anchors, in order to make the transition to Chapter 5 immediate, where these same trade-offs will be translated into operational tools (toolbox, rituals, and guardrails) by type.

In the empirical material, A&R decisions typically appear as choices under trade-offs: there is no solution that simultaneously maximizes all relevant objectives. Some trade-offs appear in both types, but they change meaning and priority depending on the position on D1 and on the pressure introduced by constraints. Consequently, problems and constraints help explain why different practices can be rational in different contexts, without assuming that one solution is absolutely superior.

The evidence is presented on two levels. First, Tab. 6 synthesizes, in comparative form, trade-offs/constraints, indicators, typical decision, salience by type, and quote anchors. Second, the text makes explicit for each tension (i) how it appears empirically, (ii) what typical choice it produces, (iii) what cost it entails, and (iv) how it varies between P1 and P2. Salience is assessed as follows: a trade-off is considered high when it is expressed as a rule/necessity with a clear cost and supported by multiple quotes (ideally from multiple interviews within the same type); it is medium when it is present but not central or with concentrated evidence; it is low when it remains an isolated mention and is treated only in a final note.

To build the trade-offs I started from the excerpts in 4.5 and from the quotes in which a concrete friction recurs in day-to-day work. At that point, I rewrote these elements as tensions between two poles: on one side what one wants to obtain (e.g., coverage, speed, control, investments), on the other what is inevitably sacrificed (e.g., follow-up, listening, coherence, sustainability). For each trade-off, in the text I:

- clarify how you recognize it in the data, i.e., which episodes or signals make it emerge (triggers/indicators);
- describe what the typical choice is that A&Rs make when they face that tension;
- make explicit what cost that choice entails (what is lost or postponed);
- compare how it changes between P1 and P2, i.e., whether the same trade-off weighs more, weighs less, or is managed with different routines.

Finally, when the data does not show a real tension between two objectives but rather a rigid limit (for example a close checkpoint, a deadline with consequences, a contractual renewal that forces a decision), I treat it as a constraint and use it to show how it increases or reduces the pressure with which P1 and P2 manage the other trade-offs.

In the text, six trade-offs and one cross-cutting constraint are developed (D9 as a stress test). The three trade-offs treated as most central for density and recurrence in the dataset are TO1, TO4, and TO5. The presentation order follows a progressive logic: it starts from the trade-off that coincides with the typological axis (TO1), moves to the operational costs of working in volume (TO2), then to the construction of “proof” under manageability constraints (TO3), then to the value compass versus company constraints (TO4), to investment and protection choices (TO5), to the tension between units of governance (artist/project vs track/hit; TO6), and closes with D9 as a stress test that modifies pressure across multiple trade-offs without generating new subtypes.

Trade-off	How it appears	Typical decision / consequence	Types where it is most salient (P1 / P2)	Implicit guardrail (only if in data; otherwise “—”)	Quote anchors [CIT: lxx]
TO1 — Care/follow-up vs coverage/throughput	roster cap / small roster “by choice”; explicit quantity↔care trade-off; in P2 many simultaneous projects and urgencies	cap the roster / reduce projects to protect tailoring (cost: less coverage) vs parallel management with dilution risk	P1 High / P2 High	roster cap; follow-up protection	[CIT: I01]; [CIT: I07]; [CIT: I08]
TO2 — Listening/vision (core) vs noise/reactivity	“email/chat” and continuous requests; need for priorities (“today you do this”); erosion of A&R time by management activities	block time for listening + daily priorities vs reactive absorption (cost: judgement and follow-up less uniform)	P1 Medium / P2 High	core-time protection; priority routine (if explicit)	[CIT: I02]; [CIT: I07]; [CIT: I08]; [CIT: I09]
TO3 — Qualified signal (proof) vs quick proxies/manageability	source of streams, active/passive; total streams “first but not sufficient”; triangulation vs “toxic metrics” (skip rate); “not too detailed”	select a few informative signals + triangulate (cost: effort) vs shortcuts on proxies (cost: attribution error/overfitting)	P1 High / P2 High	avoid single metric; avoid “metric-as-target”	[CIT: I01]; [CIT: I06]; [CIT: I07]; [CIT: I09]
TO4 — Artistic/cultural	guardrail “music/art/culture”	raise the artistic threshold even	P1 High / P2 High	“do not maximize a	[CIT: I01];

compass vs economic-organizational constraints	and artistic proposal+numbers; “we are a company”; targets/reporting and portfolio management	with numbers vs prioritize sustainability/targets under pressure (cost: identity coherence)		metric” (if used as a rule)	[CIT: I06]; [CIT: I07]
TO5 — Ambition/investment vs sustainability/protection	projects stopped due to non-sustainable investment; renewals/deciding economics; planning on margin and cash	stop/reduce ambition or re-allocate vs invest increasing exposure (cost: portfolio stress / cash)	P1 Medium / P2 High	check with administration; protection rules	[CIT: I08]; [CIT: I07]; [CIT: I06]
TO6 — Artist/project (long) vs track/hit (short)	“long paths” vs hit; track as more measurable unit; limit of hit→career/live conversion	choose “career/project” as unit vs optimize single/short window (cost: short-termism or slowness/inefficiency)	P1 Medium / P2 Medium (modulated by D9)	—	[CIT: I06]; [CIT: I07]; [CIT: I09]; [CIT: I08]
V1/D9 — Long runway vs close cliff (stress test) (constraint)	“couple of years” / long validation times vs short contracts and checkpoints (renewals expiring, economics decisions)	with close cliff: raise threshold/close earlier/re-allocate; with long runway: more patient development	P1 Medium / P2 Medium (cross-cutting stress test)	—	[CIT: I08]; [CIT: I06]; [CIT: I07]

Tab. 6 – Trade-offs

First, the material shows a structural trade-off between the possibility of maintaining continuity of care and follow-up on a limited number of projects and, on the other side, increasing coverage through a broader portfolio (TO1). In P1 cases this choice is narrated as intentional and made operational through explicit limits, for example a roster “very small by choice” and a ceiling beyond which one does not go [CIT: I01]; the same tension also appears as a dispersion risk when one “bets on too many artists” instead of concentrating [CIT: I08]. In P2 cases, instead, coverage is mainly the effect of load and of the number of projects happening at the same time (“six, seven, eight projects”) [CIT: I07]: consequently, the daily decision is not so much setting a cap, but choosing priorities and accepting that care becomes less uniform.

This difference immediately links to a second trade-off, which concerns defending time dedicated to “core” work (listening, judgement, project vision) against noise and reactivity (emails, chats, requests, and coordination) that devour the day (TO2). In P2 cases the problem is described very concretely (“you might never stop answering emails”) [CIT: I02] and as multi-channel pressure (“in seven chats”) [CIT: I07]; the typical response is trying to impose priorities and protect blocks of time (“today you do this”) [CIT: I02], otherwise the cost is fragmented judgement and discontinuous follow-up. In P1 cases the same tension emerges less as communicative overload and more as a constraint that forces you to choose what to leave behind (“*understand what to drop... there isn’t time to do everything*”) [CIT: I09]; in some passages time is also eroded by “management” activities that take space away from A&R work (“*a lot of time... in the whole product manager activity*”) [CIT: I08].

Within these constraints sits a third trade-off: the tension between understanding well what the data are saying (reading good signals and interpreting them) and, on the other side, using a few simple and fast indicators when time is short and projects are many (shortcuts on proxies) (TO3). In P1 cases, proof is qualified by breaking down and reading the quality of the signal, for instance by looking at the “source of streams” [CIT: I01] and by distinguishing active listens as an indication of fanbase [CIT: I01]. Alongside this, a guardrail appears against blind optimization: some metrics, if turned into a target, deform the content being evaluated (the “skip rate” that pushes you to remove intros, reducing the artistic value of some tracks and flattening the offer) [CIT: I01]. In P2 cases total streams are recognized as the first indicator, but competence lies in contextualizing [CIT: I07] and in rejecting mono-signal readings (“*put them in correlation... there isn’t an answer based on a single metric*”) [CIT: I06]. In parallel, a manageability constraint also emerges: too many metrics and too much detail make you lose the picture (“*they can never be too detailed*”) [CIT: I09]. In sum, the tension is high in both types, but in P2 the probability of relying on quick proxies increases because load forces faster decisions; in P1, more often, pressure is discharged upstream through selection and portfolio limitation.

A fourth trade-off makes explicit the tension between an artistic/cultural compass and the economic-organizational constraints that enter the choice (targets, reporting, sustainability) (TO4). In P1 cases the compass is formulated as both a criterion and a limit: remembering that we are talking about music/art/culture, not maximizing a percentage [CIT: I01], and defending the label’s identity in the long run [CIT: I08]. In P2 cases the same choice is described as being “inside” a business logic, where there is not only the artistic side [CIT: I07] and where objectives have to be met [CIT: I06]. The point is not that one of the two poles disappears: what changes is how the balancing is done and how often sustainability and target criteria enter.

From here we arrive at a fifth, very concrete trade-off: pushing ambitious investments to scale a project versus protecting sustainability and the portfolio through stops or reallocations (TO5). In P1 cases this often appears as a stop when “*there was no possibility of economic investment*” [CIT: I08], and, when the decision is close, the need appears to pair it with an economic-administrative check [CIT: I08]. In P2 cases the same tension takes a more continuous portfolio form (“*how much do I invest in one thing compared to another*”) [CIT: I07] and can be aggravated by close planning on expenses and margin [CIT: I06]. When resources do not allow the required investment, the operational choice becomes stopping or reducing ambition (P1) or redistributing budget across

multiple lines (P2); in both cases the cost is giving up potentially valid options in order to avoid imbalances that compromise the overall portfolio (in the sample, on the P1 side, this tension is documented in a particularly explicit way in I08).

A sixth trade-off finally concerns the unit of evaluation and governance: reasoning on career/project (long horizon) versus governing the performance of the single track (short window) (TO6). In P2 cases the preference for “long paths” is stated [CIT: I06], but it coexists with routines that, in practice, shift attention to the track as the more immediately measurable object (*“recognizing... a successful track more than an artist”*) [CIT: I07]. In P1 cases the tension also emerges as a conversion limit: a track can remain a digital hit without translating into touring capacity [CIT: I09], or as friction between “results first” and *“organic development of the path”* [CIT: I08].

A cross-cutting constraint (D9) modifies the pressure with which these trade-offs are experienced: long runway versus close cliff. When there is validation space, it becomes possible to sustain more patient choices (*“spending a couple of years to try to make an album”* [CIT: I08]; *“more than a year to understand if a project works”* [CIT: I06]). When instead the relationship or permanence is tied to close checkpoints, such as short contracts [CIT: I06] or renewals expiring with economic decisions to take [CIT: I07], pressure increases and the management of trade-offs tends to stiffen: P1 raises thresholds and stops earlier to avoid imbalances on the few lines followed in depth [CIT: I08], while P2 intensifies reallocations and stops across multiple projects because pressure propagates on a broad portfolio and makes the use of manageable proxies more frequent when one has to decide quickly [CIT: I06][CIT: I07]. With a long runway, both types can sustain more patient choices; the structural costs of the two setups remain (coverage in P1, dilution in P2).

In the material, other tensions also appear, that are not developed in the main text in order to respect the stop rule and preserve comparability: first, the tension between implementability (team alignment) and speed emerges especially in P1 cases, but it tends to overlap with the governance mechanisms already recalled in the type cards, so here it remains as an empirical reminder; second, the tension between exploration/diversity and bubble/homogeneity (network/social filter) is relevant for implications, but it has more concentrated coverage; it is therefore kept as an input for Chapter 6 guardrails without being treated as a cardinal comparative trade-off in this chapter.

Overall, some trade-offs are cross-cutting (TO3, TO4, TO5, TO6), but they change form and priority between P1 and P2: in P1, time and resource allocation tends to be solved upstream through limits and stops (protection of depth), while in P2 it tends to be solved in the course of work through priorities, reallocations, and manageable shortcuts. Chapter 5 translates this comparative map into a toolbox of rituals, templates, and guardrails coherent with each type and with the pressure introduced by D9.

Next, chapter 5 will not introduce new concepts, but will translate this evidence into practice, proposing for each type a toolbox of light-weight tools (templates, rituals, decision rules, and guardrails) built to manage recurring trade-offs and to hold up under pressure when checkpoints get closer (D9). The operational indications in Chapter 5 should therefore be read as coherent with

the two types observed in the sample and with the declared assignment rules, without forcing extensions to out-of-scope cases.

Chapter 5 – Tool Design

Chapter 5 has an operational objective: to translate the results of Chapter 4 into tools that can be used in everyday practice. In particular, downstream of the configurational typology built on one primary axis (D1 - Depth vs Volume) and four modifiers (D9 - Runway/Cliff; D5 - Main proof; D7 - Governance; D3 - Value compass), the focus here shifts from what distinguishes the two profiles (P1 and P2) to which operational actions to activate in order to manage the constraints and trade-offs that emerged. The setup therefore remains unchanged: P1 and P2 are two operational positions on D1; the modifiers do not generate subtypes but modulate pressures and modes of execution.

The proposed translation does not assume the existence of universal best practices. On the contrary, each tool is presented as a situated response to an observed trade-off in the previous chapter, and becomes defensible because it is anchored in empirical evidence (quotes and sample cases) and coherent with the organizational constraints described by participants. In this logic, the tools do not solve trade-offs, but make them governable, by making costs, applicability conditions, and side risks explicit. The chapter therefore delivers (i) a set of guiding principles to move from empirical evidence to operational implications, (ii) a Toolbox Map that links recurring trade-offs to tools and expected outputs, and (iii) two comparable toolboxes, one for P1 and one for P2, designed with the same architecture and connected to the most salient trade-offs for each position.

5.1 Guiding principles for translating empirical evidence into operational implications

This section clarifies the chapter's setup. The toolbox aims to be an orderly way to move from the problems that emerged in the data to more explicit operational decisions. Each tool answers a concrete question, for example: how to allocate time in the short term, whether a project justifies an investment, or whether it is appropriate to continue or stop. In this step, a single number is not enough, because the same value can increase for different reasons (greater exposure, playlist placements, campaigns). To reduce interpretation errors, the question must be formulated first and then a few signals should be read together (2-3), including, when possible, also the origin of the data. Below are the 7 guiding principles:

1. **PG1 (guiding principle 1) - Tool = answer to a trade-off.** Each toolbox element must explicitly state which trade-off/constraint it addresses (Tab. 6: TO1 – T06; and D9 as a stress test), which typical decision it supports (e.g., load cap, triage, re-allocation, stop/continue), and finally which cost it makes acceptable or more transparent.

2. **PG2 - No single KPI.** In the data, reading signals and metrics appears useful when it avoids shortcuts and does not collapse into a number-as-target: *“there isn’t an answer based on a single metric”* [CIT: I06], and at the same time complexity must remain manageable (“not too detailed”, otherwise “you lose a piece”) [CIT: I09]. The operational rule is therefore twofold: avoid single metrics as an automatic criterion (even when they are “salient” or easy to communicate), and limit triangulation to a few informative signals, chosen based on the type (P1/P2) and on workload.
3. **PG3 - Separate signal from exposure.** The distinction between total streams and the origin/quality of flows is central to avoid attribution errors and dependence on proxies: in P1 cases, practices emerge of reading the “source of streams” and the type of playlists from which listens come [CIT: I01], while in the sample the need to distinguish active vs passive stream as a minimum interpretation grid is also explicit [CIT: I07]. In guardrail terms, this implies both recording the origin of the signals used as proof and maintaining a distinction between more organic signals and signals potentially induced by editorial or infrastructural devices, so as not to confuse visibility with demand.
4. **PG4 - Protect time and attention on the “core” (listening and vision), especially when load generates reactivity.** Trade-off TO2 shows that, particularly in P2, noise and continuous requests erode A&R time and increase the risk of reactive management (*“they’re writing in seven chats...”*, *“in a day you might never stop answering”*) [CIT: I07][CIT: I02]. The rule is design-level: any tool set for P2 must include at least one explicit time-protection mechanism for listening and judgement (*“it’s what really makes the difference”*) [CIT: I02], otherwise the toolbox ends up optimizing only throughput and coordination.
5. **PG5 - Tools proportional to load (D1).** Volume does not coincide with a big company, but with projects-per-A&R, so tools must scale with load and not with a context label. In P1, care is made possible by a deliberately limited roster [CIT: I01] and by strong follow-up protection; in P2, the same need for decision quality passes through triage and continuous re-allocation, because care is structurally diluted. [CIT: I07] The consequence is that many rules are shared (triangulation, anti-proxy, core protection), but cadence, formalization, and the weight of the ritual change.
6. **PG6 - Make explicit the tension between the value compass (D3) and economic-organizational constraints (D7).** In the empirical material, the value compass appears as a guardrail against myopic optimization (*“we are talking about music/art/culture... not maximizing a percentage”*) [CIT: I01], but at the same time targets, reporting, and sustainability constraints emerge (*“we are a company...”*, objectives and reports) [CIT: I07][CIT: I06]. Tools must therefore: (i) avoid flattening the decision onto performance or compliance, (ii) make the constraints that enter the decision transparent (budget, renewals, timing), and (iii) preserve at least minimal space for implementable artistic evaluation (team alignment in P1; negotiation and portfolio logic in P2).
7. **PG7 - Use D9 (Runway/Cliff) as a cross-cutting stress test.** The same P1/P2 configuration behaves differently when the project/portfolio is close to checkpoints with consequences (renewals, cash, deadlines). In P1, pressure can translate into

caution and protection rules (“we protect ourselves”) [CIT: I08]; in P2, short runway amplifies stop/continue and re-allocation decisions, also through contractual devices and financial planning [CIT: I06]. The rule is therefore to embed in each toolbox at least one mechanism that signals when the context is shifting from runway to cliff and that indicates how thresholds and rituals change (without introducing a new type).

These rules are applied in 5.2 through a synthetic map (Toolbox Map) and two symmetrical toolboxes for P1 and P2: same blocks (rituals, templates, useful vs toxic metrics, guardrails), but content calibrated to the most salient trade-offs and constraints in each position.

5.2 Toolbox by typology and Toolbox Map

This section delivers the operational core of the chapter: a tool map (Toolbox Map) and two comparable toolboxes, one for P1 and one for P2.

The tools are grouped into four families, recurring in both toolboxes: rituals (cadences and explicit decision points), operational schemes (minimal artefacts, for example cards or grids, that make choices replicable and comparable), useful vs toxic metrics (selection of manageable signals and reading criteria), and finally guardrails (protection rules against systematic errors). Each element is anchored to at least one trade-off in Tab. 6 and, when available, to a quote that makes the operational problem or the reading criterion explicit.

The Toolbox Map is conceived as a guide that indicates which tools to activate when a trade-off becomes salient.




Fig. 3 – Toolbox Map

The Toolbox Map therefore functions as an operational index: starting from the most salient trade-off, it points to the relevant tools and finally clarifies the expected output. D9 (runway/cliff) functions as a cross-cutting stress test that changes thresholds and cadences, without introducing subtypes. The two toolboxes follow.

5.2a Toolbox P1 (D1=0: Depth/focus)

P1 describes a configuration in which deep care is made possible by a deliberately limited portfolio: a “very small roster by choice” and an explicit quantity vs care trade-off [CIT: I01]. The stated cost is reduced coverage; the typical risk is staying too small (few projects) and missing opportunities and, in some cases, closing into a “bubble” or into a catalog that is too uniform [CIT: I09]. Tools for P1 therefore have two objectives: first, to protect depth without closing into a bubble and without losing timing at critical passages (signing, investments, deadlines), especially when D9 shortens; second, to avoid merely formal decisions: a “yes” to a project should be accompanied by a minimal definition of operational responsibilities (who does what), essential resources, and the level of internal alignment needed for execution. This matters because, in P1 cases, signing is explicit and binary (“green light or red light”) [CIT: I08] and, to avoid getting stuck afterwards, it requires the team to be genuinely aligned (“*it is essential that the whole team is convinced*”) [CIT: I01].

(1) Ritual: defensible signing decision. In P1, the signing choice is a decision point with high irreversibility: depth of care increases the cost of error and makes economic sustainability more relevant. The data show two elements to translate into a ritual: (i) the decision is framed in a clear-cut way [CIT: I08] and (ii) it is supported by internal discussion and, when necessary, by an

economic-administrative check [CIT: I01][CIT: I08]. The operational proposal is to turn this into a short, repeatable ritual, with three steps: (a) team alignment (implementability), (b) triangulated proof (D5), and (c) D9 stress test (runway/cliff) together with sustainability. To make the ritual truly repeatable without weighing it down, the most effective choice is to use a one-page sheet to be filled in during the signing decision, so that the elements that matter for P1 in the material (clear-cut decision, manageable but qualified proof, constraints and time pressure) emerge in the same place.

In the next page is the one-page sheet referenced above (Figure 4), designed to be completed during the signing decision and to make the step repeatable without introducing additional bureaucracy.

Signing Decision Sheet (P1 — Depth/Focus)

Project/Artist: _____	Date: ____ / ____ / ____
A&R: _____	D9 (Runway/Cliff): ☐ Runway ☐ Cliff
Next checkpoint (renewal / budget / deadline): _____	

1) Decision

Decision: ☐ YES ☐ NO

One-sentence reason (max): _____

2) Proposal + observable response

Artistic proposal evaluation: _____

Observable response (numbers/signals): _____

Note: avoid making only "proposal" or only "numbers" the single criterion.

3) Proof - 3 to 5 signals + source

Signal	Value/evidence	Source (where it comes from)	Interpretation
1	_____	_____	_____
2	_____	_____	_____
3	_____	_____	_____
4	_____	_____	_____
5	_____	_____	_____

Note: include source when possible (stream source / playlist type / active signals).

4) Constraints (resources + sustainability)

Minimum resources needed (time / budget / team): _____

Admin/finance check: ☐ Not needed ☐ Done ☐ To do

Main risk: _____

5) D9 (Runway/Cliff) - operational implications

If CLIFF: what changes? (higher threshold / cautious investment / earlier stop)

If RUNWAY: what is feasible? (more development / more testing)

Closure

Who does what (ownership): _____

Next step (within 7 days): _____

Fig. 4 – Toolbox P1

In practice, the sheet includes:

1. a “yes/no” **decision** + the reason in one sentence (consistent with the clear-cut framing of the decision point) [CIT: I08];
2. **proposal** + **response**, keeping together a “convincing proposal” and numbers without turning them into a single criterion [CIT: I01];
3. **proof** (D5) in 3–5 signals, including, when possible, also the origin (e.g., source of streams, playlist type, active-listening signals) and one line of interpretation [CIT: I01];
4. resource and sustainability **constraints**, with a possible administrative/financial check [CIT: I08];
5. **D9 (runway/cliff)**, making explicit whether one is in cliff mode and how thresholds and ambition change (e.g., higher threshold or a more cautious investment) [CIT: I08]. At the end, the sheet also indicates who does what and the next short-term step: two minimal details that avoid approving a project without making it executable [CIT: I01].

(2) Operational scheme: triangulating proof. In P1, “proof” does not coincide with the absence of data, but with data that are interpreted and broken down: looking at the “source of streams”, distinguishing the quality of flows, and reading signals of active engagement [CIT: I01]. In parallel, a guardrail emerges against metrics that, if optimized as targets, deform the content (skip rate and intros) [CIT: I01]. Operationally, the scheme helps avoid two opposite errors: collapsing everything into an easy proxy, or getting lost in excessive detail (“not too detailed”) [CIT: I09].

Below, therefore, is a minimal guiding grid, which does not replace the “proof” table in the sheet (which serves to record value, source, and interpretation of each signal), but helps choose which signals to include as rows in the proofs, so as to stay on a small set of informative elements that are comparable across projects.

Proof element	What I look at (in practice)	Why it matters (error it avoids)
1) “Active” listening signals	Indicators that suggest a real/engaged fanbase (e.g., “active listens”, return signals, attention). [CIT: I01]	Avoids mistaking passive consumption for stable demand. Avoids attribution errors: the same number can be driven by exposure, not demand.
2) Origin of streams	“Source of streams”: where listens come from; distinguish channels and, when relevant, playlist type. [CIT: I01]	
3) Proposal–response coherence	Put together (in one line) the quality of the artistic proposal and the observable response: “numbers + convincing proposal”, without turning it into a single criterion. [CIT: I01]	Avoids unbalanced decisions: signing “only for numbers” or “only for taste”.
4) One metric not to turn into a target	Make explicit an indicator I should <i>not</i> “maximize” if it risks deforming the content (e.g., skip rate → pressure to cut intros). [CIT: I01]	Avoids the metric becoming the goal and shifting creative behavior in an undesired way.

Tab 7 - guide to filling out the “Proof” section

(3) Guardrail: explicit load limit + follow-up rule. If P1 is possible because an operational limit exists (for example a roster “very small by choice”), that limit should also be made explicit as a tool, because it is what makes deep care sustainable over time. In practice, the guardrail can be formulated as an A&R WIP cap, i.e., a maximum number of active projects that require regular follow-up. By active projects I mean the lines that, in order not to lose coverage, require a minimum cadence of contact and progress, not projects that are simply on the roster. The cap is therefore not a universal value: it can be calibrated pragmatically by observing for a short period how many lines actually receive this minimum follow-up; what cannot be followed at that cadence, by definition, should not stay in the WIP (it goes on hold or exits the operational perimeter). When the cap is exceeded, the typical decision is not “do more”, but to choose what to put on hold, stop, postpone, or to outsource part of execution, so as not to lose continuity of coverage on the projects that remain active (as happens when some components are assigned externally to free up time for the development part) [CIT: I08].

(4) Light anti-bubble ritual. In the P1 material, signals of closure risk appear (“*we live a bit in a bubble*”) and catalog uniformity risk [CIT: I09]. Without introducing new dimensions, this can enter as a minimal guardrail: a periodic routine (e.g., monthly) in which, holding load constant, one checks whether incoming scouting always comes from the same channels and whether some signals (e.g., social) are becoming exclusive filters. The point is not to impose diversity by default, but to prevent a convenient proxy from shrinking exploration.

(5) D9 as a stress test: variable thresholds and ambition. Here D9 functions as a stress test of the P1 configuration: the same work logic (few projects followed in depth) changes in a meaningful way depending on how much time there is before a checkpoint with consequences. When runway is long, the data suggest it is realistic to sustain more patient development trajectories, where it makes sense to invest gradually and leave room for work that takes time (“even a couple of years...”) [CIT: I08]. When instead runway shortens and one enters cliff mode, the priority becomes avoiding a situation where a single project, precisely because it is followed in depth, generates an imbalance that is hard to absorb: under these conditions, the need to protect oneself and contain risks emerges [CIT: I08].

5.2b Toolbox P2 (D1=2: Volume/parallelism)

P2 describes a configuration in which a high number of projects-per-A&R produces urgency, communicative noise, and fragmented follow-up: simultaneous projects (“six, seven, eight”) [CIT: I07], multi-channel pressure (“they’re writing in seven chats”) [CIT: I07], and the risk of spending the day answering emails [CIT: I02]. In this setup, the decision is not only signing/no signing, but a continuous process of prioritization and allocation (time, attention, budget). Tools for P2 therefore aim to: (i) make priorities visible and governable, (ii) protect core listening-and-judgement time (which “makes the difference”) [CIT: I02], and (iii) reduce dependence on quick proxies without introducing unmanageable bureaucracy.

(1) Ritual: weekly portfolio triage + stop/continue rules. Trade-off TO1 in P2 is not solved through a roster cap by choice, but through portfolio governance that accepts volume and manages its costs [CIT: I07]. The proposal is a short, recurring ritual (weekly) that produces two outputs: a priority list (top N), and stop/continue/re-allocation decisions when sustainability constraints emerge (TO8). The ritual reduces daily reactivity (even if it cannot be eliminated entirely) and prevents priorities from being determined only by the noisiest channel or the most insistent request. [CIT: I02]

In the next page is the one-page sheet referenced above (Figure 5), designed to make portfolio management operational in P2: set a small number of weekly priorities and take explicit continue/pause/stop/re-allocate decisions, anchoring them to checkpoints and to a minimal proof, so as to reduce the risk that priorities are dictated only by the flow of requests and by channel noise [CIT: I02][CIT: I07].

Portfolio Triage Sheet (P2 — Volume/Parallelism)

A&R: _____	Date: ____ / ____ / ____
WIP (active projects): ____	D9 (Runway/Cliff): ☐ Runway ☐ Cliff
Main checkpoint (renewal / budget / deadline): _____	Next review date: ____ / ____ / ____

1) Protected core time - defend the “core” against noise

Non-negotiable listening/judgement blocks (date + time):

1) _____ 2) _____

Listening queue / focus (what must be heard this week):

2) Weekly priorities - Top-3 with ownership

Project/Line	Why now (urgency / trade-off)	Key decision	Next action

3) Portfolio triage board - stop/continue/re-allocate

Project/Line	Next checkpoint	Decision (C/P/S/R)	Resource move	Proof snapshot (2–3 + source)	Notes
		☐ C ☐ P ☐ S ☐ R			
		☐ C ☐ P ☐ S ☐ R			
		☐ C ☐ P ☐ S ☐ R			
		☐ C ☐ P ☐ S ☐ R			
		☐ C ☐ P ☐ S ☐ R			
		☐ C ☐ P ☐ S ☐ R			
		☐ C ☐ P ☐ S ☐ R			

Legend: C=Continue P=Pause S=Stop R=Re-allocate. Avoid single-metric decisions; include source when possible.

4) D9 (Runway/Cliff) - fast-track rules

If CLIFF: increase triage frequency, tighten stop/continue rules, and prefer reversible moves.

If RUNWAY: schedule check-ins and avoid mistaking exposure spikes for stable demand.

Budget/finance check (if needed): ☐ Not needed ☐ Done ☐ To do

Closure (minimum trace)

One-sentence portfolio narrative (for reporting/alignment):

Fig. 5 – Toolbox P2

In practice, the sheet makes fillable what in P2 would otherwise remain implicit and scattered across many conversations and requests. It contains:

1. a **protected core time block** for listening and judgement (when and on what), because the data make it explicit that this activity is what “really makes the difference”, but it is easily eroded by emails/chats and coordination [CIT: I02];
2. a choice of a few **priorities** (Top-3) with a rationale and next action, to avoid work being determined by the most insistent channel and to translate the need to “know what to do today” into an explicit decision [CIT: I02];
3. a **portfolio-management grid** in which, for each line, one records:
 - the project name
 - the next checkpoint
 - a synthetic management decision (continue/pause/stop/re-allocate),
 - in case of re-allocation, an explicit resource re-allocation choice (time/attention/budget: what increases, what decreases),
 - an intentionally lean proof snapshot (2–3 signals) with source indication, which explains the rationale for the decision. The objective is to hold together two constraints that coexist in the sample: avoiding single-metric decisions and, at the same time, not getting lost in unmanageable detail when projects are many [CIT: I06][CIT: I07];
4. an operational reminder of D9 (runway/cliff) that activates a fast track when the portfolio approaches deadlines, renewals, or close-range economic constraints (more frequent triage, tighter stop/continue rules, more reversible moves) [CIT: I06][CIT: I07].
5. a space for a minimum of responsibility (who does what) and a one-line “narrative” of the portfolio: useful elements when work requires coordination and reporting, because they make the choice communicable and defensible without turning it into an extended report [CIT: I06].

(2) Operational scheme: rapid and manageable triangulation. In P2, the temptation to rely on proxies grows because workload forces speed; however, in the sample the rule to correlate signals and reject single-metric readings is explicit (“put them in correlation...”, “it’s not a single pattern”) [CIT: I06]. The translation is already embedded in the P2 sheet reported above: the “Proof snapshot / mini proof frame” column functions as a template because it forces the recording of a few signals (2–3) in a consistent way across projects, together with the source, without turning the choice into a report.

(3) Guardrail: protected listening blocks. The data include an explicit request for protection and planning (“*to calendar it... is already work*”; “*I need someone to tell me: today do this*”) [CIT: I02]. In P2, the rule is to create a minimal, non-negotiable protection of listening time and

judgement building. This is coherent with the marker that listening is what “really makes the difference” [CIT: I02] and it avoids decision quality being eroded by coordination.

(4) Governance and reporting as an explicit constraint and as an input to triage. In P2, reporting and objective constraints emerge (“answering to superiors... reports”, “reaching objectives”) [CIT: I06], and interdepartmental negotiation on economic and operational aspects [CIT: I07]. The toolbox incorporates this as an implementability requirement: weekly triage must produce communicable outputs (priorities + rationale), because under governance constraints the decision is not only personal but must be defensible and coordinable.

(5) D9 as a stress test. In the P2 material, short runway appears associated with contractual devices and close-range financial planning (short contracts, expenses and margin on short windows) [CIT: I06], as well as economic decisions linked to renewals [CIT: I07]. In these cases, the toolbox makes explicit a decision “fast track”: (i) higher triage frequency, (ii) greater use of stop/continue rules, (iii) timely re-allocation to prevent pressure from propagating across many lines at once.

Overall, the two toolboxes share the same architecture and the same underlying guardrails (anti single-metric, separating signal/exposure, core protection), but they differ in intensity and in the point where the system absorbs and manages operational pressure: in P1 mostly upstream (limits and selection), in P2 in-flight (priorities, re-allocations, anti-reactivity routines).

5.3 Lightweight implementation

This section clarifies how to implement the tools described in the previous section in a “lightweight” way, i.e., without introducing new procedures or requiring major changes in the firm’s working logic. The idea is to integrate the tools into ordinary work, with minimal steps and an essential written trace that makes clear what was decided and why. Below are the practical rules for applying the tools in the two typological positions, P1 and P2.

In P1 (Depth/focus), the rule is to use the signing sheet as a standard step every time a “heavy” and weakly reversible decision is made (in particular: signing a project or committing resources in a meaningful way). The sheet is not meant for reporting: it is meant to concentrate, on one page, the clear-cut decision and the elements that, in the data, make the choice implementable (a manageable but qualified proof, constraints, the checkpoint, and the next step). In this way, the toolbox enters where P1 is most vulnerable: when a deep-care choice can absorb time and resources and becomes hard to correct ex post.

In P2 (Volume/parallelism), the rule is to use the portfolio sheet as a regular routine (typically weekly) to set a small number of priorities and make portfolio-management decisions explicit (continue, pause, stop, re-allocate). Here too, the sheet is not a report: it is a quick way to prevent

priorities from being determined only by channel noise and urgencies, and to record, in a lean way, a mini proof frame (a few signals with source) and, when needed, the re-allocation move.

Finally, when D9 is “short” (close-range checkpoints, renewals, economic constraints), the same rule applies with a tighter cadence and more cautious thresholds: pressure increases, stop/continue and re-allocations grow, and it becomes more likely to slide into shortcuts in reading signals.ù

5.4 Closing and bridge to Chapter 6

With this chapter, the translation of Chapter 4 results into operational tools is complete: guiding principles (5.1), a Toolbox Map (5.2), and two comparable toolboxes for P1 and P2, made usable through one-page sheets and lightweight application rules (5.3).

During the interview setup, the initial aim of the thesis was also to understand whether and how AI tools entered A&R’s decision logics. However, as data collection progressed, it emerged that these tools are often marginal or absent in operational-managerial practices, despite a declared and widespread interest among interviewees. For this reason, I chose to include a short concluding chapter on the topic, which can offer interpretive cues and some future perspectives on how AI might affect A&R work, taking into account constraints, opacity, and infrastructural dependencies that are typical of the platform ecosystem.

Chapter 6 — AI in A&R work

The previous chapter ended with an empirical observation that influenced the structure of this thesis. In the phase of setting up the interviews, an additional objective was to understand whether and how artificial intelligence tools were entering the decision-making logics of A&R. As data collection progressed, however, it emerged that, despite a declared and widespread interest on the part of the interviewees, such tools are often marginal or absent in everyday operational-managerial practices. For this reason, this chapter is deliberately brief and plays a different role compared to the results chapters, namely offering interpretive keys and some future-oriented implications on how AI could affect A&R work, explicitly taking into account constraints, opacity, and infrastructural dependencies that characterize the platform ecosystem.

The chapter therefore combines two ingredients. On the one hand, it takes as a baseline the result that emerged from the interviews (marginal adoption). On the other hand, it integrates external evidence and industry documentation to map where AI can enter the workflow in a realistic way, why it often does not enter, and under what conditions adoption would make sense without treating AI as an automatic, non-verifiable verdict.

6.1 Why AI is often marginal in the everyday practices of the A&R

Starting from the interviews' baseline ("high interest, low use"), the question becomes: why does adoption remain marginal in operational-managerial practices? The evidence discussed in the report suggests that the bottlenecks are not only technical. They are also structural and organizational, and they closely reflect the constraints of the platform ecosystem. In some cases, this marginality is formulated clearly: "*As for data... there still aren't (AI tools for data analysis).*" [CIT: I05].

First, data access and observability are partial by design. A large part of the signals relevant for A&R is mediated by platforms that decide what is visible, with what granularity, and under what conditions. In this context, many promises of AI (for example, robustly predicting long-term success) are structurally difficult: a model can only learn from what the infrastructure allows it to see. Second, adoption is constrained by skills, resources, and asymmetries. In practice: even if the tool exists, time (and capability) is needed to understand what it is measuring, how reliable it is, and in which cases it can be wrong. If these conditions are missing, AI risks staying outside everyday work, or being used "on trust" without checks. Not by chance, when AI is described as a concrete need, it is imagined more as practical help than as a substitute for judgment (for example: "*an AI secretary... who organizes my agenda*") [CIT: I02]. The point here is not that AI is useless, but that its value depends on how much the organization can verify its outputs. If resources are few, the temptation increases to take the score as an answer; if, instead, resources are many, the opposite risk increases: building routines that chase and optimize platform signals too much, ending up confusing visibility and value. Third, AI enters an environment in which signals are not purely organic. Visibility can be induced by editorial choices, playlisting infrastructures, marketing spend,

or platform affordances. Without separating organic demand and induced exposure, AI systems risk formalizing the platform's selection policy as if it were neutral evidence. This explains why interviewees can declare interest in AI but use it little: an in-depth qualitative assessment can appear more reliable than a score that cannot be attributed to stable and interpretable sources. More than a score, the need expressed is often a translation from fragmented data to decision (*"taking this information... and turning it into actions: 'what do I do with it?'"*) [CIT: I01]. Finally, there is a problem of accountability and legitimation. Signing and allocating resources are decisions with high reputational risk. If a tool does not make clear why it produces a recommendation, it becomes difficult to defend the decision internally and, if necessary, externally. When AI is imagined as a direct evaluator of music, the judgment is often skeptical (*"You feed it a demo and it tells you whether to sign... honestly it's not useful to me"*) [CIT: I01]. In practice, this tends to shift the use of AI toward more "defensible" tasks (triage, anomaly detection, synthesis) and away from final judgments.

6.2 What happens today

If we look at what happens today in documented practices and public cases, AI enters above all as a support infrastructure: it accelerates initial screening and makes analysis of scattered signals more manageable, while "heavy" decisions (signing and major investments) remain anchored to an in-depth qualitative assessment of the music, the person, the team, the context, and execution capability. This picture is consistent both with industry accounts on the use of data/ML tools to support A&R, and with works that treat signing as a multi-criteria decision (not reducible to a single score) (Knibbe, 2020; Choicharoon et al., 2024).

In practical terms, four uses recur with more evidence:

1. **Screening and watchlists in scouting** (internal tools or data platforms). Here AI is used to build shortlists of artists/projects to deepen (i.e., a restricted list of priority candidates), to intercept early growth signals, and to organize large volumes of artists/releases. An explicit case is Soldatone in Warner Music Group: in corporate communication it is described as proprietary technology that combines streaming/social/touring data with machine learning methods to identify early signals and support operational insights (Soldatone, 2018). Trade press writes that these tools would help manage more scouting and evaluate more artists (Soldatone, 2020).
2. **Multi-signal coherence and sanity checks** (reducing false positives). A recurring use of AI in the A&R area is to check whether traction is coherent across multiple sources: streams, follower growth, engagement, playlist adds, trend over time. In practice, AI (or scoring systems) is used to surface anomalies and inconsistencies and to speed up the first filter, before moving to qualitative assessment (Knibbe, 2020).
3. **Routines linked to events and windows** (monitoring + timing). AI often helps decide when to act on a project already under observation (for example: monitoring performance in the first 24-72 hours, choosing release windows, coordinating pitching activities, or understanding if and when to intervene on budget and content). In an ecosystem marked by

editorial cycles and seasonality, these tools are used above all to keep pace and prioritize operational attention, consistently with the way metrics and algorithmic affordances structure decision rhythms in streaming (Maasø & Hagen, 2020).

4. **Tagging, organization, and discovery** (especially on the platform side). This point does not concern A&R directly in its way of working, but it concerns it because these AI functions define the visibility context within which A&R itself works. Platforms use AI to classify and organize the catalog, make recommendations, do playlisting (build and update playlists), and do tagging (assign labels such as genre, mood, or activity). For A&R this matters in two ways: it influences how a track is found and in which contexts it enters (search, mood, playlists, recommendations), therefore it changes positioning and pitching levers, and it influences the signals that A&R observes, because part of traction can depend on these classification/curation mechanisms. In the SoundCloud case, for example, the company clarified that tools such as “Musio” are used for discovery and content organization (not to train generative models on users’ music), together with other “non-generative” uses such as anti-fraud and recommendation (Music Business Worldwide, 2025; TechCrunch, 2022).

Across all these points, the most controversial part remains using AI to predict what will be a success or not, and making the decision to sign only based on that. AI can also make predictions or give scores in the short term, but turning that score into a “yes/no” is fragile for two simple reasons: first because an artist’s long-term outcome depends on many things that AI does not see (team, strategy, investments, context, execution), and second because many platform dynamics are opaque, so it is difficult to understand why a number goes up or down and attribute causes in a safe way (Knibbe, 2020).

6.3 Future perspectives: plausible shifts and related risks

Even if AI is today marginal in operational practices, it remains interesting to ask how, in the future, AI technology may significantly reshape A&R work. Three trajectories appear plausible.

First, it is likely that AI will strengthen its role as scalability infrastructure: better triage, better synthesis, better internal documentation. In this trajectory, the main benefit is not prediction, but the reduction of cognitive load and the increase of comparability and traceability of decisions. Second, AI may increase pressure toward two linked dynamics: waiting for more data signals before signing, and ending up basing choices on the same numbers. If models are trained on what becomes visible and measurable, they tend to privilege what already produces data, shifting exploration toward already-validated candidates, penalizing new artists who do not yet have data, and, according to judgment devices theory, increasing the risks and costs of standardization (Knibbe, 2020; Karpik, 2010). Third, a growth of the reliability problem of numeric signals is likely. With the increase of tracks uploaded at scale (including AI-generated ones) and with the possibility of inflating numbers with bots, “artificial” patterns can imitate organic traction. In 2025, Deezer reported that a very high share of streams of AI music on the platform was linked to fraud (up to ~70% in the cited report) and that detection systems are used to block payments on

anomalous activity. For A&R, this strengthens the need for multi-signal validation and the idea that spikes should be treated as hypotheses to verify, not as evidence to trust (The Guardian, 2025).

Overall, these trajectories suggest that AI's impact may concern less an automation of taste and more a reconfiguration of attention allocation and standards of proof in A&R work, under infrastructural constraints that largely remain outside the control of A&R itself.

6.4 Minimal guardrails for using AI under platform constraints

Since the central risks are organizational (not only technical), guardrails must be light enough to be adopted, but strong enough to prevent the most damaging failure modes.

A first guardrail is never to use a single metric. Decisions about singularities should not collapse into a single score, especially in an era where AI can further distort metrics: signals should be read in relation to each other and cross-validated across sources. A second guardrail is to label the source of the signal: as much as possible, separate organic signals from induced visibility (editorial, paid amplification, event-driven spikes). Without this step, AI systems risk automatically incorporating (and therefore later normalizing) the platform's selection policy. A third guardrail is to treat cold start as the rule. If the tool implicitly excludes what lacks data, the organization should define an explicit "low-data" lane with qualitative criteria and checkpoints, instead of a silent exclusion. A fourth guardrail is to make AI failure modes explicit before using outputs: virality bias, short-lived spikes, manipulation and trend-chasing, so as to treat AI output as a signal to verify with checks (other data and qualitative assessment). Finally, guardrails must recognize opacity: even well-designed tools operate on partial observability, especially in a market (that of art) where unpredictability remains naturally high. It follows that, the less reversible a decision is, the more AI should be used to generate questions and checks, not to replace the decision.

6.5 Ex-ante adoption criteria

Finally, this chapter aims to provide current and future A&Rs with criteria to decide when and how the adoption of an AI technology is sensible. The identified criteria are the following:

1. **Specify the task and the minimum output:** triage, validation, synthesis, timing, or catalog organization. A tool is useful if it produces an actionable and defensible output (shortlist, alerts, red flags, comparative synthesis), not if it simply adds a score.
2. **Test data and observability:** which sources are used, with what coverage, and with what ability to distinguish organic signals from induced signals, making explicit what happens in case of cold start.
3. **Require plurality and interpretability:** if the system returns a single, non-debatable score, it behaves like a single-criterion judgment device, an anti-pattern in this domain

according to Karpik. Outputs should show which signals drive the recommendation and in which direction.

4. **Define controls and accountability:** who decides what to do with the output and where the reason (even minimal) is recorded to avoid AI becoming an ex post justification.
5. **Assess coherence with typology and constraints.** In high-risk contexts with low tolerance for error, AI should mainly reduce false positives and increase defensibility; in contexts with greater slack it can support broader exploration; in data-rich contexts, the risk is present-bias and dependence on the platform; in data-poor contexts, the risk is adopting opaque outputs without verification capability.
6. Finally, one last criterion, increasingly relevant when using general-purpose models for synthesis and textual analysis, concerns **vendor transparency and compliance**. In the EU framework, the AI Act introduces obligations for providers of general-purpose models, including the publication of a summary of training contents according to a Commission template. For an A&R team this translates into minimal vendor due diligence (data provenance, copyright policy, usage constraints, logging and traceability), not as an abstract compliance exercise, but to avoid downstream reputational or legal risks (European Union, 2025).

Consistently with the result that emerged from the interviews (AI still marginal in operational practices), this chapter frames AI not as a new core of A&R work, but as an emerging set of assistive tools that can reallocate attention and structure preliminary evaluation, under infrastructural constraints and opacity. The key managerial implication is not adopting AI in general, but deciding when it makes sense to use it and with which rules, choosing in advance what it must do, what data it can use, how to control typical errors, and who takes responsibility for the final decision. Under these conditions, AI can support scalability and traceability without becoming an additional constraint or a substitute for judgment. The next chapter presents the conclusions of the study, synthesizing the main results and the implications that emerged.

Chapter 7 - Conclusions

This chapter closes the thesis without introducing new results: it synthesizes and interprets, coherently, what has already been presented in Chapters 2-6. The aim is to provide explicit answers to RQ1 and RQ2, using exclusively the evidence already discussed, clarify what the thesis adds (contributions) and what cannot be supported in a defensible way (limitations), and deliver a closing that is consistent with the typology-first approach, in which the logical chain remains: dimensions → axis D1 → positions P1/P2 → trade-offs → toolbox. In Chapters 4-6, the thesis worked on a corpus consisting of 10 interviews and made the chain from data to dimensions traceable; on this basis it defined a single-axis typology (D1) with two empirically stable positions (P1 and P2), an intermediate area (D1≈1) that is plausible but did not emerge as a stable configuration, and a set of modifiers (in particular D9/D5/D7/D3) that modulate decision pressure without generating sub-types. This architecture is the interpretive constraint for everything that follows.

The chapter proceeds as follows: 8.1 closes RQ1 (the evolution of the A&R role between fragmentation and specialization); 8.2 closes RQ2 (empirical dimensions and configurational typology); 8.3 makes contributions explicit (2 theoretical + 2 managerial); 8.4 states limitations, suggests future directions, and closes with a final synthesis.

7.1 Answer to RQ1 - Evolution of the A&R role between fragmentation and specialization

RQ1 asked how the A&R role has evolved and how it has been articulated between fragmentation and specialization.

The first element of the answer concerns the shift in bottlenecks. In the reconstructed trajectory, the growth of supply abundance and the centrality of platforms shift part of selection and visibility onto evaluation and coordination devices (playlists, rankings, metrics), with relevant effects on how signals and attention are produced and distributed. In this context, A&R does not operate only in a market of “unique” projects characterized by quality uncertainty: it also operates within an ecosystem in which some signals become more salient and more actionable (but not for this reason more causal or more true), and in which attention allocation tends to be mediated by external devices that reduce uncertainty but introduce risks of delegation, conformity, and standardization. The second element concerns the operational consequences of this transformation, namely a reconfiguration of constraints and competences. Platform metrics and affordances feed feedback loops in which action, interpretation, and outcome can reinforce each other; in this scenario, the pressure to react to fresh signals (spikes, alerts, dashboards) coexists with the opacity of causal links and with asymmetries in analytical capacity between actors. This produces new required competences (reading signals, managing noise, the ability to avoid overfitting to proxies) and new constraints (dependence on external devices, more compressed timelines, more fragmented

attention). The third element concerns the articulation between fragmentation and specialization. Within the discussed change framework, the thesis reads fragmentation not as simple dispersion, but as a condition that pushes toward practices that are more specialized and sensible in context: in some cases, what matters most is working on a few projects, curating them deeply and building solid proof before committing resources; in others, what matters most is managing many projects in parallel and continuously deciding where to put time and attention, even when there is a lot of noise. This articulation is not deduced a priori, but is traced back to the empirical configurations observed and formalized in the answer to RQ2.

In synthesis, RQ1 is closed with a central idea: in the contemporary ecosystem, A&R operates under a growing pressure of selection and allocation in conditions of abundance, in which judgment devices and metrics become coordination infrastructure; this does not eliminate the interpretive component, but shifts the trade-offs and incentivizes practical specializations that will be made comparable through dimensions and typology (RQ2). It follows that, in the sample, the evolution of the role manifests above all as a differentiation of work regimes and decision priorities (more “depth” or more “volume”) rather than as a linear replacement of judgment with KPIs.

7.2 Answer to RQ2 - D1 (Depth vs Volume) and the P1/P2 typology

In the sample, the answer to RQ2 (“which empirical dimensions distinguish today A&R practices and decision logics across different contexts, and which configurational typology follows?”) is built by making cases comparable through explicit rules, without using context (major/indie, budget, country) as a shortcut to define types. Comparability stems from the traceable chain data→dimensions (Fig.1 / Tab.2) and leads to a single-axis configurational typology: D1 (Depth vs Volume). D1 does not say whether an organization is big or small, but how much load falls, on average, on a single A&R figure, i.e., how many projects they can actually keep “top of mind” and follow with continuity. In the analyzed material, this is the clearest and most polarized difference; for this reason, two stable positions emerge, P1 (D1=0) and P2 (D1=2), while when the evidence is not sufficient the classification is not forced (NC). An intermediate point (D1≈1) is plausible, but here it does not appear as a recurring configuration solid enough to formalize. The two decision profiles differ as follows:

- in P1, work tends to concentrate on a few projects, because care and follow-up have a cost and require time. As a consequence, signing often emerges as a “yes/no” choice that, to be sustainable, is accompanied by internal alignment and by a search for more qualified proof;
- in P2, instead, the load makes parallel management of many projects normal: the recurring decision is not only “do I sign/do I not sign”, but above all what to put at the top of the list today, where to allocate time, attention, and budget, and how to defend space for listening and vision amid noise and reactivity.

The typology remains configurational because the same position on D1 can work differently under different pressures: for this reason, a few modifiers enter (D9/D5/D7/D3), with D9 (runway/cliff)

as a stress test. In particular, in the sample P1 tends to manage “upstream” the care vs coverage trade-off (TO1) by limiting and selecting the portfolio; P2 more often lives the trade-off between listening/vision and noise/reactivity (TO2) as structural. Within these constraints, TO3 also becomes central: the compromise between building truly qualified proof and using manageable proxies when time is tight. In this framework, metrics and judgment devices do not eliminate quality uncertainty: they make it more manageable, but they also introduce risks (single-criterion, standardization, confusion between signal and exposure) that motivate dedicated guardrails and routines.

7.3 Contributions

The thesis therefore offers four main contributions (two theoretical and two managerial), all internal to the outputs already presented.

Theoretical contributions:

1. **A typology-first configurational typology of A&R practices:** the thesis proposes a parsimonious single-axis typology (D1) with explicit assignment rules and with context treated as a modifier (intermediate not forced; NC when the evidence does not hold). In this sense, the typology works as a transferable conceptual grid to describe A&R work regimes beyond the observed cases, leaving open empirical verification on larger samples.
2. **Trade-offs as the comparative unit:** the thesis makes cases comparable through observable and assignable trade-offs (TO1–TO6) and uses D9 as a stress test to read how decision pressure changes. This explanatory frame is generalizable in a theoretical sense because it allows interpreting decision practices under uncertainty and constraints as solutions to recurring tensions, rather than as simple individual preferences.

Managerial contributions:

3. **Translation into a typology-consistent toolbox:** the thesis translates the differences that emerged between P1 and P2 into a set of operational tools consistent with each work regime and, in fact, leaves two documents (a toolbox for P1 and a toolbox for P2) that an A&R can use as a reference if they see themselves in that profile. The Toolbox Map and the symmetric toolboxes for P1/P2 organize tools into families (rituals, templates, useful vs toxic metrics, guardrails) and show how to activate them based on the dominant trade-off and pressure (in particular D9), avoiding the idea of a single “best practice” valid for all contexts.
4. **Light implementation and defensibility of the decision:** the thesis also proposes a “low-friction” adoption mode, designed to enter ordinary work without introducing new bureaucracy. The principle is to leave an essential trace of what was decided and why, and to use a few targeted routines to protect signal quality and accountability (reducing the effect “salient signal = decision”). In practice, the setup aims to hit typical vulnerabilities: in P1, avoiding over-investments or poorly

reversible choices without checkpoints; in P2, reducing reactivity and noise that erode listening and vision, making the allocation of time and attention more stable.

7.4 Limitations, future research, and conclusion

The main limitations of the research are the following. The first is that the evidence derives from 10 interviews: the research therefore does not claim to describe all contexts, but offers a grid that can be applied and tested on other cases; in the meantime, here it remains anchored to the corpus of 10 interviews and the conclusions are formulated within the sample, and not presented as valid for the entire population of A&Rs. The second limitation concerns the evidence and the coverage of the typology. For rigor, when on D1 the proof is not clear enough or appears to be in conflict, I do not force the assignment to P1 or P2 and I use the NC category. This is a prudential choice (therefore a point of rigor), but it also implies that the typology, in the corpus, does not cover all cases with the same informational strength. The third limitation is that, when for some modifiers information is missing (NA), I can say less about how that profile works, even if assignment on D1 is possible. Finally, the last limitation concerns the method: Gioia-informed and iterative analysis makes the data→concepts path more transparent, but it does not eliminate the limits of a non-longitudinal study and of coding not done by multiple people.

Consistently with these limitations, three developments are a priority for subsequent research: (1) extend and stress-test the typology on larger samples, with attention to the intermediate area $D1 \approx 1$; (2) introduce longitudinal designs centered on D9 (runway/cliff) as a stress test; (3) validate the toolbox in real contexts to understand whether it reduces proxy-driven decisions and protects core time without introducing bureaucracy.

In conclusion, the thesis argues that the recent evolution of the music ecosystem makes the A&R work more exposed to judgment devices and salient metrics, intensifying recurring trade-offs. In the sample, the most discriminating difference between the observed cases is the position on D1 (Depth vs Volume), which stabilizes two configurations (P1 and P2) and makes choices and vulnerabilities readable through a trade-off map, with D9 as a stress test. The operational consequence is not a single best practice, but a typology-consistent toolbox: different tools to manage different pressures, with a minimal and defensible goal, namely preserving signal quality and decision accountability in a context in which what is measurable and salient would otherwise risk replacing what is qualified and relevant.

References

- Aguiar, L., & Waldfogel, J. (2018). *Platforms, promotion, and product discovery: Evidence from Spotify playlists* (JRC Working Papers on Digital Economy No. 2018-04). European Commission, Joint Research Centre. <https://ideas.repec.org/p/ipt/decwpa/2018-04.html>
- Bakare, L. (2025, July 14). *An AI-generated band got 1m plays on Spotify. Now music insiders say listeners should be warned. The Guardian*. <https://www.theguardian.com/music/2025/jul/14/an-ai-generated-band-got-1m-plays-on-spotify-now-music-insiders-say-listeners-should-be-warned>
- Beer, D. (2016). *Metric power*. Palgrave Macmillan. <https://doi.org/10.1057/978-1-137-55649-3>
- Blumer, H. (1954). What is wrong with social theory? *American Sociological Review*, 19(1), 3–10. <https://doi.org/10.2307/2088165>
- Bonini, T., & Gandini, A. (2019). “First week is editorial, second week is algorithmic”: Platform gatekeepers and the platformization of music curation. *Social Media + Society*, 5(4), 1–11. <https://doi.org/10.1177/2056305119880006>
- Bowen, G. A. (2006). Grounded theory and sensitizing concepts. *International Journal of Qualitative Methods*, 5(3), 12–23. <https://doi.org/10.1177/160940690600500304>
- Caves, R. E. (2000). *Creative industries: Contracts between art and commerce*. Harvard University Press.
- Chartmetric. (2018, April 23). *How Warner Music uses AI and data to spot talent in a world of streaming*. <https://chartmetric.com/blog/how-warner-music-uses-ai-and-data-to-spot-talent-in-a-world-of-streaming/>
- Choicharoon, A., Hodgett, R., Summers, B., & Siraj, S. (2024). Hit or miss: A decision support system framework for signing new musical talent. *European Journal of Operational Research*, 312(1), 324–337. <https://doi.org/10.1016/j.ejor.2023.06.014>
- Dalugdug, M. (2025, May 12). *SoundCloud addresses AI clause in terms of service, says it “has never used artist content to train AI models”*. *Music Business Worldwide*. <https://www.musicbusinessworldwide.com/soundcloud-addresses-ai-clause-in-terms-of-service-says-it-has-never-used-artist-content-to-train-ai-models/>
- Doty, D. H., & Glick, W. H. (1994). Typologies as a unique form of theory building: Toward improved understanding and modeling. *Academy of Management Review*, 19(2), 230–251. <https://doi.org/10.5465/amr.1994.9410210748>

Espeland, W. N., & Sauder, M. (2007). Rankings and reactivity: How public measures recreate social worlds. *American Journal of Sociology*, 113(1), 1–40. <https://doi.org/10.1086/517897>

European Commission. (2025, July 24). *Explanatory notice and template for general-purpose AI model providers to summarise the data used to train their model*. Shaping Europe’s digital future. <https://digital-strategy.ec.europa.eu/en/library/explanatory-notice-and-template-general-purpose-ai-model-providers-summarise-data-used-train-their>

European Data Protection Board. (2021). *Guidelines 07/2020 on the concepts of controller and processor in the GDPR*(Version 2.0, adopted 7 July 2021). https://www.edpb.europa.eu/system/files/2023-10/EDPB_guidelines_202007_controllerprocessor_final_en.pdf

European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

European University Institute. (2019). *Guide on good data protection practice in research* (3rd ed.). <https://www.eui.eu/documents/servicesadmin/deanofstudies/researchethics/guide-data-protection-research.pdf>

Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 167–194). MIT Press. <https://doi.org/10.7551/mitpress/9780262525374.003.0009>

Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), 15–31. <https://doi.org/10.1177/1094428112452151>

Healy, K. (2011). Judgment and distinction. *Socio-Economic Review*, 9(4), 787–800. <https://doi.org/10.1093/ser/mwr010>

Hennink, M., & Kaiser, B. N. (2022). Sample sizes for saturation in qualitative research: A systematic review of empirical tests. *Social Science & Medicine*, 292, 114523. <https://doi.org/10.1016/j.socscimed.2021.114523>

Hesmondhalgh, D. (2019). *The cultural industries* (4th ed.). SAGE.

Karpik, L. (2010). *Valuing the unique: The economics of singularities* (N. Scott, Trans.). Princeton University Press.

Kitchin, R. (2014). *The data revolution: Big data, open data, data infrastructures and their consequences*. SAGE.

- Kluge, S. (2000). Empirically grounded construction of types and typologies in qualitative social research. *Forum Qualitative Sozialforschung / Forum: Qualitative Social Research*, 1(1), Article 14. <https://doi.org/10.17169/fqs-1.1.1124>
- Knibbe, J. (2020, November 10). “A&R has always been about data”: A deep dive into the role of data in A&R with Chaz Jenkins, CCO at Chartmetric. *Music Tomorrow*. <https://www.music-tomorrow.com/blog/ar-has-always-been-about-data-a-deep-dive-into-the-role-of-data-in-ar-with-chaz-jenkins-cco-at-chartmetric>
- Knight, F. H. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.
- Kvale, S., & Brinkmann, S. (2015). *InterViews: Learning the craft of qualitative research interviewing* (3rd ed.). SAGE.
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. SAGE.
- Locke, K., Feldman, M. S., & Golden-Biddle, K. (2022). Coding practices and iterativity: Beyond templates for analyzing qualitative data. *Organizational Research Methods*, 25(2), 262–284. <https://doi.org/10.1177/1094428120948600>
- Maasø, A., & Hagen, A. N. (2020). Metrics and decision-making in music streaming. *Popular Communication*, 18(1), 18–31. <https://doi.org/10.1080/15405702.2019.1701675>
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87. <https://doi.org/10.1287/orsc.2.1.71>
- Maxwell, J. A. (2013). *Qualitative research design: An interactive approach* (3rd ed.). SAGE.
- Miles, M. B., Huberman, A. M., & Saldaña, J. (2014). *Qualitative data analysis: A methods sourcebook* (3rd ed.). SAGE.
- Morris, J. W. (2020). Music platforms and the optimization of culture. *Social Media + Society*, 6(3), 1–10. <https://doi.org/10.1177/2056305120940690>
- Morris, J. W., & Powers, D. (2015). Control, curation and musical experience in streaming music services. *Creative Industries Journal*, 8(2), 106–122. <https://doi.org/10.1080/17510694.2015.1090222>
- Negus, K. (1992). *Producing pop: Culture and conflict in the popular music industry*. Edward Arnold.
- Negus, K. (2002). The work of cultural intermediaries and the enduring distance between production and consumption. *Cultural Studies*, 16(4), 501–515. <https://doi.org/10.1080/09502380210139089>

- Palinkas, L. A., Horwitz, S. M., Green, C. A., Wisdom, J. P., Duan, N., & Hoagwood, K. (2015). Purposeful sampling for qualitative data collection and analysis in mixed method implementation research. *Administration and Policy in Mental Health and Mental Health Services Research*, 42(5), 533–544. <https://doi.org/10.1007/s10488-013-0528-y>
- Patton, M. Q. (2015). *Qualitative research & evaluation methods* (4th ed.). SAGE.
- Peterson, R. A., & Anand, N. (2004). The production of culture perspective. *Annual Review of Sociology*, 30, 311–334. <https://doi.org/10.1146/annurev.soc.30.012703.110557>
- Peterson, R. A., & Berger, D. G. (1975). Cycles in symbol production: The case of popular music. *American Sociological Review*, 40(2), 158–173. <https://doi.org/10.2307/2094343>
- Peterson, R. A., & Berger, D. G. (1996). Measuring industry concentration, diversity, and innovation in popular music. *American Sociological Review*, 61(1), 175–178. <https://doi.org/10.2307/2096413>
- Prey, R. (2020). Music streaming playlists and curatorial power. *Social Media + Society*, 6(2), 2056305120933291. <https://doi.org/10.1177/2056305120933291>
- Prey, R., Esteve Del Valle, M., & Zwerwer, L. (2022). Platform pop: disentangling Spotify's intermediary role in the music industry. *Information, Communication & Society*, 25(1), 74–92. <https://doi.org/10.1080/1369118X.2020.1761859>
- Schenk, P. (2021). Can judgment devices explain the demand for fine wine? How quality assessments affect the demand for a credence good. *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 73, 59–85. <https://doi.org/10.1007/s11577-021-00794-4>
- Seaver, N. (2019). Captivating algorithms: Recommender systems as traps. *Journal of Material Culture*, 24(4), 421–436. <https://doi.org/10.1177/1359183518820366>
- Silberling, A. (2022, May 3). *SoundCloud acquires Musiio, an AI music curator, to improve discovery*. TechCrunch. <https://techcrunch.com/2022/05/03/soundcloud-acquires-musiio-an-ai-music-curator-to-improve-discovery/>
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118. <https://doi.org/10.2307/1884852>
- UK Data Service. (2022). *How to anonymise qualitative and quantitative data* (A. Vlad & M. Haaker). <https://ukdataservice.ac.uk/app/uploads/howtoanonymisequalandquantdata2022-04-21.pdf>

Warner Music Group. (2018, March 28). *Warner Music Group acquires Sodatone*. Warner Music Group. <https://www.wmg.com/news/warner-music-group-acquires-sodatone-33396>

Weick, K. E. (1995). *Sensemaking in organizations*. SAGE.

Wikström, P. (2019). *The music industry: Music in the cloud* (3rd ed.). Polity Press.

Zwaan, K., & ter Bogt, T. F. M. (2009). Research note: Breaking into the popular record industry: An insider's view on the career entry of pop musicians. *European Journal of Communication*, 24(1), 89–101. <https://doi.org/10.1177/0267323108098948>

List of Figures

Figure 1: Data Structure (Gioia method).....	46
Figure 2: Configuration scheme.....	65
Figure 3: Toolbox Map.....	80
Figure 4: Toolbox P1.....	82
Figure 5: Toolbox P2.....	86

List of Tables

Table 1: Sample snapshot.....	44
Table 2: Themes overview.....	47
Table 3: Dimensions' synthesis.....	52
Table 4: Cases x dimensions matrix.....	64
Table 5: Type cards' summary.....	70
Table 6: Trade-offs.....	72
Table 7: Guide to filling out the “Proof” section.....	84