TIRESIAS AI · WHITEPAPER
April 2026

Mapping the Deep Structure of Human Preference

Why Psychographic Modeling Outperforms Behavioral Data in Entertainment Investment — and What That Means for the Decisions That Matter Most

The entertainment industry runs on taste. Hits and misses are determined by whether a property's emotional and cognitive offer matches what specific audiences want at a specific moment. The challenge: taste is not directly observable. You can only observe its outcomes — box office receipts, stream counts, cancellation decisions — after the money is spent. This paper sets out the scientific case for a different approach: measuring the deep structure of human preference directly, before capital is committed, using the personality psychology that governs why people like what they like.

1. The Taste Problem

Entertainment investment is a fundamentally different problem from most capital allocation decisions. When a pharmaceutical company funds a drug trial, the target is biochemically defined. When an infrastructure fund finances a toll road, demand is modelled from traffic patterns and population data. Both involve uncertainty, but the uncertainty is about external factors — market growth, regulatory risk, execution quality.

Entertainment investment has a different uncertainty at its core: *taste*. Success depends on whether a large enough group of people will find a specific emotional and cognitive experience sufficiently compelling to pay for it — not once, in a test environment, but repeatedly, voluntarily, and competitively against every other option available to them that week.

Taste is not random. It has structure. People's preferences cluster in predictable patterns that persist across time, across domains, and across formats. A person who seeks out morally complex narratives in literature will tend to seek them in television. A person who values emotional warmth and resolution in film tends to value it in the books they read. These patterns are not coincidences of demographic overlap. They are expressions of stable underlying personality traits — traits that are measurable, mappable, and directly predictive of content preference.

The scientific literature on this relationship is now extensive. Yet the entertainment industry's primary measurement tools — box office history, genre categorisation, demographic surveys, focus groups, collaborative filtering — treat taste as an emergent aggregate rather than a mappable structure. They measure taste's outcomes rather than its causes. This produces a systematic vulnerability: the tools work adequately when the historical data is abundant and the product is similar to previous products. They fail precisely when the stakes are highest.

The further you get from behavioral data, the wider our advantage.

2. Where Current Methods Fail

2.1 The Behavioral Data Trap

Collaborative filtering — the recommendation engine architecture underlying most streaming platforms — works by finding users who behave similarly and assuming their future preferences will also converge. It is effective within a stable, well-populated catalog. Its failure mode is structural: it requires behavioral history to make predictions. When that history doesn't exist — new IP, new audience, new platform — it has nothing to work from.

This is not a data quality problem. It is a category error. Behavioral data tells you what people have done. It cannot tell you what they will do when the product is new, the platform is new, or the audience profile has shifted. For the decisions that require the most capital commitment — greenlighting an original series, acquiring a catalog, licensing a gaming IP for screen adaptation — behavioral history is structurally inadequate.

2.2 The Library Merger Problem

When two large content catalogs merge, the recommendation and audience modelling systems face a cold-start problem at scale. The engagement history of Catalog A's subscribers tells you nothing reliable about their response to Catalog B's content, because the behavioral signals were generated in a different product context, against a different competitive set, for a different price point.

The merger that looked like a data asset turns out to be a data liability. Two separate behavioural models, built on incompatible signal sources, cannot simply be averaged. The resulting recommendation system produces lower-quality outputs for both catalogs until new engagement data accumulates — which takes months, during which subscriber churn risk is elevated.

Psychographic models do not have this problem. Personality traits are stable across contexts. A subscriber's preference structure — their appetite for moral complexity, their threshold for emotional intensity, their response to ensemble-driven storytelling — does not change when the catalog around them changes. A personality-grounded model transfers across catalog mergers without degradation.

2.3 The Gaming-to-Screen Translation Failure

The adaptation of gaming IP to film and television has become the entertainment industry's most expensive recurring problem. Properties with hundreds of millions of engaged players have produced a consistent string of theatrical underperformers. The conventional explanation — "gaming audiences don't translate to theatrical audiences" — is accurate but unhelpful. It describes the symptom without diagnosing the cause.

The cause is a systematic confusion between *platform audience* and *IP affinity*. The behavioral data available for a gaming property — completion rates, in-game purchase patterns, character selection, time-in-game — describes the gaming audience's relationship with *the game,* not with the IP's underlying narrative and emotional features. These are different things.

A game succeeds through agency, mastery, and reward loops. A film succeeds through catharsis, character arc, and emotional resolution. The audience that responds to an IP's story in game form may respond to the same IP's story in film form — but not through the same emotional pathway, and not in the same demographic slice. Mapping gaming behavioral data onto a theatrical audience prediction is a category error that produces consistent mispricing.

Psychographic modelling addresses this directly. By scoring the IP's underlying content features independently of format, and then mapping those features against personality-driven audience segments for the target format, it becomes possible to identify which portion of a gaming property's audience has the personality profile that predicts theatrical engagement — and, equally importantly, what that audience needs the adaptation to deliver.

2.4 The Demographic Proxy Problem

Age and gender remain the primary segmentation axes for most entertainment audience research. They persist because they are measurable, legally defensible, and historically predictive at the population level. They fail at the IP level because the variance within demographic groups vastly exceeds the variance between them for most content decisions.

A 28-to-35 female demographic contains: people who prefer psychologically complex prestige drama; people who prefer light romantic comedy; people who prefer crime procedurals; people who prefer supernatural fantasy. Their demographic membership is identical. Their taste profiles are not. A marketing campaign designed for "the 28-to-35 female audience" that doesn't account for this internal variance is, at best, finding the centre of a distribution that doesn't describe anyone precisely.

2.5 The Tracking Era Is Ending

The critique of behavioral data tools is not only methodological. It is structural — and it is now being enforced by regulatory and platform-level decisions that the industry has spent a decade trying to avoid.

GDPR arrived in 2016. The Cambridge Analytica scandal broke in 2018, triggering legislative attention across every major market. CCPA took effect in 2020. Apple's App Tracking Transparency framework launched in 2021, cutting 55% of mobile traffic tracking and costing Meta alone an estimated $10 billion in lost annual revenue. In 2024, Google formally admitted defeat on cookie deprecation — not because the privacy concern had been resolved, but because the behavioral tracking infrastructure that the industry built its business model on had become legally and technically untenable.

The platforms that depend on behavioral tracking are already paying the bill in lost attribution, broken targeting, and unpredictable subscriber churn. Digital journalism has been hollowing out for a decade; CPC is higher than ever; new regulatory motions routinely wipe out billions in market cap with a single ruling.

This is not a temporary headwind. It is a structural shift. The behavioral data infrastructure — tracking pixels, third-party cookies, identity graphs, cross-site retargeting — is being dismantled by regulatory pressure, platform policy, and changing consumer expectations simultaneously. The question for every company that runs its audience intelligence on behavioral data is not whether this affects them. It is what they replace it with.

Psychographic modelling is, structurally, the answer. Personality traits are inferred from consented first-party signals. They are stable — unlike behavioral patterns, they do not degrade when cookies are blocked or device IDs are masked. They are portable across platforms and catalogs. And critically: as PII regulation tightens, psychographic prediction becomes more competitive, not less. The privacy tailwind that is destroying behavioral data businesses is a direct advantage for trait-science infrastructure.

> ***"Stop studying ripples. Study the rock."***

3. The Scientific Foundation

The case for psychographic modelling rests on a body of evidence in personality psychology that is now sufficiently large, sufficiently replicated, and sufficiently well-validated to constitute a reliable scientific foundation for commercial application. The research is not recent. Since the 1930s, lexical and factor-analytic research has independently re-derived the same personality dimensions across languages, cultures, and decades — arriving at the same five stable trait dimensions from entirely separate methodological starting points. This convergence across 90 years of research from independent scientific traditions is what makes the Big Five framework suitable for commercial prediction at scale. Three bodies of work are most directly relevant.

3.1 Personality Predicts Content Preference

Rentfrow and Gosling's 2003 study in the *Journal of Personality and Social Psychology* established the foundational empirical basis for personality-to-content-preference mapping.[1] Across a series of studies with thousands of participants, they demonstrated that Big Five personality traits reliably predict music preferences — with predictive accuracy for behavioural outcomes ranging from 41% to 66% of a colleague's independent behavioral observation. Openness to Experience predicted preference for reflective, complex musical genres. Extraversion predicted preference for energetic, rhythmically-driven music. Agreeableness correlated with preference for upbeat, emotionally warm content.

The two decades of research that followed have substantially extended and quantified this finding. Schäfer and Mehlhorn's 2017 meta-analysis of personality and musical style preferences[2] documented that Openness is the dominant predictor of preference for classical, jazz, and sophisticated genres (r = .13), while Neuroticism drives preference for sad and melancholic music (r = .22) — a distinct emotional regulation pathway. Across all traits and all musical genres, the average effect size in personality-music research is r = .058. That figure requires context: in population-level prediction applied across millions of users, an effect of that magnitude produces commercially significant audience segmentation.

Nave, Minxha, Greenberg, and colleagues extended this work in a 2018 study published in *Psychological Science,* demonstrating that musical preferences predict personality not only from self-report but from active listening behaviour and Facebook Likes — linking the passive signal of what someone engages with to the psychographic profile that explains why.[3] The relationship between personality and content preference is bidirectional and robust at scale.

The predictive relationship extends well beyond music. Openness predicts preference for arthouse and unconventional film, literary fiction, philosophy, and immersive game genres (RPGs, indie titles) with consistently replicated positive effects. The specific content features that drive Openness-aligned preference — narrative complexity, thematic novelty, moral ambiguity, symbolic depth — are exactly the features that distinguish high-value IP with underserved audiences from mass-market content with saturated audiences. This is the commercial application the entertainment industry has been missing.

[1]Rentfrow, P. J., & Gosling, S. D. (2003). The do re mi's of everyday life: The structure and personality correlates of music preferences. Journal of Personality and Social Psychology, 84(6), 1236–1256.
[2]Schäfer, T., & Mehlhorn, C. (2017). Can personality traits predict musical style preferences? A meta-analysis. Personality and Individual Differences, 116, 265–273.
[3]Nave, G., Minxha, J., Greenberg, D. M., Kosinski, M., Stillwell, D., & Rentfrow, J. (2018). Musical preferences predict personality: Evidence from active listening and Facebook Likes. Psychological Science, 29(7), 1145–1158.

3.2 Personality Is Inferable from Behaviour at Scale

Kosinski, Stillwell, and Graepel's 2013 study in the *Proceedings of the National Academy of Sciences* demonstrated that digital behavioral traces can reliably infer Big Five personality traits at scale.[4] Using a dataset of over 58,000 US Facebook users, they showed that personality dimensions could be predicted from digital behavior with meaningful accuracy — establishing the bidirectional relationship that makes psychographic modelling commercially viable: behavior predicts personality, and personality predicts preference. The myPersonality project that underpinned this research eventually grew to include data from over six million participants, making it the largest personality dataset ever assembled from a single platform.

The inference relationship has since been replicated across platforms and behavioral signal types. Liu and Campbell's 2017 meta-analysis of Big Five traits and social media, drawing on 113 independent samples (N=53,913), confirmed that Extraversion predicts social media use frequency (r = .11–.40 across platforms) and Neuroticism is the strongest predictor of problematic social media use.[5] The consistent finding across studies: social behavior generates reliable personality signal. Every interaction — what someone follows, what they share, what genre they engage with, how long they linger — is a data point in a psychographic inference model.

This finding has a direct implication for entertainment audience intelligence. The behavioral data that streaming platforms and social media platforms already collect is not just a record of past preferences — it is a latent personality signal. Properly interpreted, it provides the psychographic profile needed to predict future preferences, including preferences for content the user has never encountered. The inference direction matters as much as the prediction direction: both are now empirically established at sufficient scale to support commercial deployment.

3.3 The Replication Question

A common concern about psychological research applied to commercial contexts is replicability. The so-called replication crisis has rightly introduced scepticism about effect sizes in social psychology. The personality-behaviour literature, however, has fared substantially better than many other areas of psychology under replication scrutiny — and the reason is structural.

The Big Five framework replicates across 50+ countries and cultures. Its structure has been independently re-derived by lexical and factor-analytic research from entirely separate starting points since the 1930s. Multiple large-scale meta-analyses confirm the pattern: the music preference evidence synthesises populations exceeding 260,000 participants; the social media meta-analysis covers 113 samples and nearly 54,000 participants; political ideology research spans 73 studies and over 71,000 participants; risk propensity research covers 133 samples and nearly 70,000 participants. These are not single-study findings.

[4] Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. Proceedings of the National Academy of Sciences, 110(15), 5802–5805.
[5] Liu, D., & Campbell, W. K. (2017). The Big Five personality traits, Big Two metatraits and social media: A meta-analysis. Journal of Research in Personality, 70, 229–240. N=53,913; 113 independent samples.

The LOOPR (Large-scale Online Ongoing Personality Research) project, which conducted systematic preregistered replications of personality-outcome associations, found that 87% of tested associations replicated in the expected direction.[6] Barrick and Mount's landmark 1991 meta-analysis of Big Five traits and job performance — finding Conscientiousness predicts job proficiency across all occupational categories (rho = .19–.28) — has been replicated in every major occupational psychology review conducted since.[7] The robustness of the Big Five framework specifically reflects the quality of its underlying measurement: the traits are dimensionally stable, normally distributed, and reliable across instruments and contexts. Effect sizes are generally small to medium (r = .10–.30). Personality accounts for meaningful but partial variance — the appropriate framing is probabilistic population-level prediction, not individual determinism.

87%	41–66%	6M+	50+
Replication Rate	**Predictive Accuracy**	**Participants**	**Countries**
LOOPR personality-outcome associations	Personality → content preference (Rentfrow & Gosling 2003)	myPersonality database · personality-behavior inference (Kosinski et al.)	Cross-cultural Big Five replication · structural invariance confirmed

3.4 The Five Classical Dimensions

The Big Five personality model — Openness to Experience, Conscientiousness, Extraversion, Agreeableness, and Neuroticism — is the most empirically validated framework for measuring stable individual differences in personality. Each dimension has decades of measurement research behind it, cross-cultural validation, and established relationships to a wide range of behavioural outcomes.

For entertainment content prediction, each dimension maps to distinct content preferences:

Dimension	TASE Label	Content Preference Signal
Openness	Intellectual Curiosity	Narrative complexity, thematic novelty, symbolic depth, unconventional structure
Conscientiousness	Discipline & Order	Clear moral frameworks, structured narrative, resolution-oriented storytelling
Extraversion	Social Energy	Ensemble casts, social stakes, high-energy pacing, group-experience content

[6] LOOPR Project: Large-scale Online Ongoing Personality Research. 87% of personality-outcome associations replicated in expected direction across preregistered replication attempts.
[7] Barrick, M. R., & Mount, M. K. (1991). The Big Five personality dimensions and job performance: A meta-analysis. Personnel Psychology, 44(1), 1–26. Conscientiousness: rho = .19–.28 across all occupational categories.

Agreeableness	Warmth	Interpersonal warmth, emotional resolution, low-conflict storytelling, catharsis
Neuroticism	Emotional Sensitivity	High emotional intensity, psychological depth; also a repeller for distressing content

4. The Tiresias Framework

The Tiresias Audience Simulation Engine (TASE) operationalises this scientific foundation as a five-stage simulation pipeline, converting IP feature profiles and personality segment data into affinity scores, audience maps, and commercial recommendations.

Our proprietary technology includes additional, scientifically validated traits beyond the Big Five, allowing us to A) better detect human taste signal, and B) make better predictions than any Big Five or classical model. In our core technology offering lies proprietary machine learning models, custom-trained to map preferences towards entertainment domains specifically, and trained on consented user data.

4.1 Content Feature Scoring

Every IP is scored across content feature dimensions (video, text, music, imagery) on a 0–1 scale across thousands of dimensions, including: narrative complexity, emotional intensity, violence intensity, visual spectacle, warmth, moral complexity, thematic novelty, humour, romance, character depth, world-building, horror tension, and pacing intensity. These dimensions are drawn from the personality-content preference research literature and are designed to map directly onto the Big Five trait dimensions that drive audience response.

This scoring is format-agnostic. The same source material — a novel, a graphic novel, a game — receives the same feature scores regardless of what adaptation format is under consideration. This is the architectural decision that enables cross-domain prediction: by separating the IP's content identity from its format identity, it becomes possible to predict audience response to an adaptation before the adaptation exists.

4.2 Psychographic Segmentation — TriggerMap™

Audience segments are derived via TriggerMap™, Tiresias's proprietary T-Score segmentation system. Our approach allows for overlapping boundaries and non-equal segment sizes — accurately reflecting the actual structure of personality distributions in populations. Segments are defined by their T-Score profiles, their content preferences, and their predicted behavioural patterns: seek-out, word-of-mouth, wait-for-recommendation, or skip.

Segments are not demographic proxies. A "Cultural Seeker" segment — defined by high Openness, moderate Conscientiousness, and prestige-signal responsiveness — may span age brackets from 24 to 58. An "Immersive Escapist" — defined by franchise loyalty patterns and high world-building affinity — may be evenly distributed across gender demographics. The segmentation reflects the actual structure of audience preference, not the convenience of measurable external characteristics.

4.3 Affinity Simulation

For each audience segment, the simulation computes a 0–1 affinity score representing the probability-weighted likelihood of positive reception. The overall affinity score is the population-weighted average across all segments, with a 95% confidence interval derived from segment variance. The confidence interval is a substantive data point, not a disclaimer: a wide interval (e.g., 0.22–0.84) indicates genuine audience polarisation — some segments will respond strongly, others will not — which is a different commercial picture from a property with a narrow interval around the same mean.

4.4 Text-to-Personality Extraction (Augur)

Tiresias's Augur LLM was custom-built for our methods. Augur provides text-to-personality extraction at higher accuracy than ever before, making it ideal for mapping existing audiences to new IPs — new novels, unpublished scripts, graphic novel runs, game narratives. Augur infers Big Five personality-relevant features directly from user-generated text, achieving a Spearman correlation of $r > 0.50$ against validated personality measures. This enables cold-start analysis for genuinely new IP: no behavioral history, no box office comps, no prior adaptation required.

4.5 Benchmark Performance

Performance is measured against two baselines. The primary commercial benchmark is industry-standard collaborative filtering — the recommendation architecture that underlies most major streaming platforms. Tiresias outperforms this baseline by +39% for film, +57% for TV, and +193% for books in top-K precision against held-out preference declarations (N=6,000 U.S. representative sample, proprietary focus group). The music genres benchmark shows +287% lift — reflecting the particular weakness of behavioral models in a format where personality is the primary predictor.

The secondary benchmark is MBTI-based segmentation, still used by a significant portion of the market intelligence industry despite its well-documented reliability problems. Tiresias's agent-swarm simulation architecture outperforms MBTI-based prediction by 30%+ on MiroFish validation datasets — a result driven by the Big Five's superior psychometric properties: the traits are dimensionally stable, normally distributed, and derived from empirical factor analysis rather than typological categories.

+39%	**+57%**	**+193%**	**+30%**
Film Prediction Lift	**TV Prediction Lift**	**Books Prediction Lift**	**vs. MBTI Baseline**
vs. collaborative filtering · N=6,000 U.S. sample	vs. collaborative filtering · N=6,000 U.S. sample	vs. collaborative filtering · N=6,000 U.S. sample	Agent-swarm vs. typological segmentation (MiroFish)

5. Three Problems, One Framework

5.1 New IP with No Behavioral History

When a studio is evaluating rights to a novel published eighteen months ago, or a game from an independent studio with no theatrical precedent, the collaborative filtering and box-office-history approaches return the same answer: insufficient data. There is no prior engagement history to model from.

Psychographic simulation has no such dependency. Augur extracts the IP's personality-relevant features from the text or narrative directly. TASE maps those features against the target audience's personality segments. The affinity score and segment breakdown are available within days, before any audience has encountered the property in its adaptation form.

5.2 Cross-Domain Adaptation

The gaming-to-screen problem is the most commercially consequential version of the cross-domain challenge. The IP's gaming audience and its potential theatrical audience share an affinity for the IP's story — but they are not the same people, in the same proportions, responding to the same features.

TASE resolves this by scoring the IP's content features independently of format, then simulating affinity against the personality segment distribution relevant to the *target format*. A gaming IP's theatrical audience simulation uses the theatrical audience's personality segment model, not the gaming audience's behavioral data. The result is a prediction of who will show up to the film — and what the film needs to deliver to retain them — rather than a description of who plays the game.

This also works in reverse: a film IP being adapted for gaming, or a literary IP being extended into audio drama, can be simulated against the relevant format audience before the cross-domain investment is made.

5.3 Post-Merger Catalog Integration

For M&A teams evaluating catalog acquisitions, and for post-merger integration teams managing combined libraries, psychographic modelling offers a stable substrate for audience mapping that does not degrade when the behavioral context changes.

Because personality traits are stable, a subscriber's psychographic profile — derived from their engagement with Catalog A — is directly applicable to predicting their affinity for Catalog B's content. The profile transfers; the behavioral context does not need to.

This enables two specific applications: (1) pre-acquisition audience mapping — identifying which segments of the acquirer's subscriber base will engage with the target catalog's content, and at what affinity levels; and (2) post-merger programming strategy — identifying content gaps and commissioning priorities that serve the combined audience's psychographic profile rather than the historical programming pattern of either catalog in isolation.

6. The Speed Imperative

The traditional content development timeline — option, develop, test, greenlight, produce, release — is measured in years. Streaming has compressed the release cycle at the back end while increasing the pressure at the front end: more content decisions, faster, with higher consequences for wrong calls.

Traditional audience research tools have not kept pace. A theatrical test screening program takes months. A focus group study for a major IP takes a month. A collaborative filtering engine requires months of engagement data accumulation to produce reliable signal.

You can have them all in a month. Feature scoring, segment simulation, affinity modelling, and commercial recommendation — the full pipeline from IP input to actionable output — is a computational process, not a data accumulation process. It does not wait for audience engagement to accumulate. It derives predictions from the stable psychological architecture of audience preference, which is already known.

This speed advantage is not just a convenience. In rights negotiation contexts — where an IP is available for a limited window, where competing bids are being evaluated, where a library acquisition is time-sensitive — the ability to generate a rigorous audience affinity analysis before the window closes is the difference between an informed decision and a gut call.

De-risk content transactions with audience intelligence built on proprietary psychometric models — before a dollar is spent.

7. What This Changes

Psychographic modelling does not replace creative judgment. A simulation cannot tell a writer's room what to write, or tell a director how to shoot a scene, or tell a marketing team what the campaign line should be. These are human decisions, and the quality of their execution is the primary determinant of whether a well-positioned property succeeds or fails.

What psychographic modelling changes is the quality of the decision framework within which those judgments are made. It provides answers to four questions that conventional methods address poorly:

First: who is the core audience for this property, defined at the level of personality-driven preference rather than demographic proxy? Second: how large is that audience, and how strongly does the property's feature profile match what they want? Third: what are the specific features that drive their affinity, and which features create risk? Fourth: which adaptation decisions — format, sequencing, tone, campaign positioning — are most consequential for audience alignment?

These are the questions that determine whether a content investment is priced correctly, positioned correctly, and developed correctly. They are answerable with psychographic data. They are not answerable with box office history, genre categorisation, or demographic surveys alone.

The entertainment industry makes billion-dollar decisions about what people will like before they have had a chance to like it. The scientific tools to make those decisions with dramatically more precision than current methods allow have existed in the academic literature for decades. The gap has been in operationalising them at the speed and scale that commercial decision-making requires.

That gap is now closed. And the timing is not coincidental.

The regulatory and platform-level dismantling of behavioral tracking infrastructure has created a forced migration. Every company that built its audience intelligence on third-party cookies, cross-device identity graphs, and behavioral retargeting is now in the market for something to replace them with. The replacement has to work without PII. It has to be privacy-compliant by design, not by exception. It has to be stable across platform changes and regulatory motions. And it has to predict preference rather than merely describe historical behavior.

Trait science is that replacement. It was the right answer before the tracking era ended. It is the only durable answer now that the tracking era is ending. The infrastructure window for building the new substrate is open. It will not stay open indefinitely.

About Tiresias AI

Tiresias AI is a psychographic intelligence platform for the entertainment industry. The Tiresias Audience Simulation Engine (TASE) provides audience affinity prediction, segment profiling, and commercial recommendations for IP licensing, greenlight decisions, M&A due diligence, and pre-market audience strategy. Tiresias is a member of the NVIDIA Inception Program. Research and client case studies are published at **www.t-me.ai**

References

1. Rentfrow, P. J., & Gosling, S. D. (2003). The do re mi's of everyday life: The structure and personality correlates of music preferences. Journal of Personality and Social Psychology, 84(6), 1236–1256.
2. Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. Proceedings of the National Academy of Sciences, 110(15), 5802–5805.
3. LOOPR Project. Large-scale Online Ongoing Personality Research: Systematic replication of personality-outcome associations. 87% of tested associations replicated in expected direction.
4. Schäfer, T., & Mehlhorn, C. (2017). Can personality traits predict musical style preferences? A meta-analysis. Personality and Individual Differences, 116, 265–273.
5. Nave, G., Minxha, J., Greenberg, D. M., Kosinski, M., Stillwell, D., & Rentfrow, J. (2018). Musical preferences predict personality: Evidence from active listening and Facebook Likes. Psychological Science, 29(7), 1145–1158.
6. Liu, D., & Campbell, W. K. (2017). The Big Five personality traits, Big Two metatraits and social media: A meta-analysis. Journal of Research in Personality, 70, 229–240.
7. Barrick, M. R., & Mount, M. K. (1991). The Big Five personality dimensions and job performance: A meta-analysis. Personnel Psychology, 44(1), 1–26.
8. Miranda, D., & Blais-Rochette, A. (2020). Neuroticism and emotion regulation through music listening: A meta-analysis. Music Perception, 37(4), 274–288.

