

The Next Generation of Spotify Podcasts

The hard part isn't finding you a podcast.


**It's giving you a reason to press play
when there are millions to choose from.**

A listener discovery problem with creator-growth consequences — what Spotify could change, and what I'd test first.

With supporting research from Edison Research, Sounds Profitable, Spotify Research and the choice-overload literature.

Mel Khasbagan

Founder & Strategist, Verité



● WHERE THIS STARTED

Why do I open YouTube when I already know exactly what I want to hear?

Spotify is where I listen to podcasts. But when I arrive with a specific need — I have had a terrible day and I want help letting go — I leave and search YouTube instead. I wanted to understand why.

That question turned into a few weekends of reading. What follows is what the research says about why this happens, what I think Spotify should change, and what I would test first. Page 16 is the part I still don't know.

THE SAME PROBLEM, FROM THE CREATOR SIDE

Spotify can put an unfamiliar creator in front of the right listener. But if that listener doesn't have enough information to take the risk, distribution hasn't become growth.

THE COST OF BEING WRONG

- A song

20 seconds

- An episode

45 minutes

A recommendation paradigm where mistakes are cheap becomes a much harder consumer problem when the wrong guess costs three quarters of an hour — and the listener knows that before pressing play.

More content and better prediction than ever, and choosing is still hard.

The problem isn't a shortage of things to listen to, and it isn't a weak model. It's what happens to a person facing both of those at once.

ABUNDANCE

Not the constraint

There is more podcast content available than any listener could sample in a lifetime. Whatever is stopping this decision, it isn't a shortage of things to listen to.

PREDICTION

World-class

Working out what a listener is statistically likely to play. Spotify has world-class prediction infrastructure and it keeps getting more sophisticated.

DECISION-MAKING

Open

Helping a person actually commit. Prediction produces a ranked list; it does not give the listener what they need to decide whether an unfamiliar hour is worth taking.

The scarce resource in podcasting isn't content — it's confidence. For listeners, confidence that this unfamiliar episode is worth their time. For creators, confidence that Spotify can connect their work with the people most likely to value it. Everything that follows is about closing that gap.

● THE CENTRAL FINDING

Podcast decisions amplify the conditions associated with choice overload.

A meta-analysis of 99 observations (N=7,202) found choice overload isn't caused by option count alone. It appears when four conditions combine.

1

Choice set complexity

Dozens of superficially similar episodes on the same subject.

2

Decision task difficulty

Little pre-consumption evidence about quality or fit.

3

Preference uncertainty

The listener often doesn't know what they want until they hear it.

4

Effort minimisation

The listener is trying to spend as little effort as possible deciding.

Podcast discovery tends to combine all four, and the cost of being wrong is far higher. The literature isn't about podcasting, so this is an argument by analogy rather than a finding. Effort minimisation is the condition I'd expect to matter most in practice — it may weigh heaviest in low-energy listening moments, which is precisely when reducing decision effort is worth the most.

● THE PARADOX

My hypothesis: uncertainty sends people back to who they already know.

When evaluating the unfamiliar costs more effort than a listener wants to spend, familiarity becomes the default. That is what I want to test, not what I claim to have shown.

WHAT CHOOSING A STRANGER COSTS

- 1 Read titles from people I've never heard of
- 2 Judge whether they're credible without much to go on
- 3 Guess whether the episode is really about what I need
- 4 Commit 45 minutes to find out



So rather than absorb that cost, I open a creator I already know.

IF THAT HOLDS

Spotify's discovery surfaces work hardest to put unfamiliar creators in front of listeners.


But if the effort of evaluating an unfamiliar creator is highest when a listener has least energy to spend, the familiar option wins by default rather than on merit.

The effort Spotify spends surfacing new creators can end up reinforcing the listening habits a listener already has.

What's missing isn't more recommendations — it's a reason to trust one. “Listeners who value the same things you value in someone you already follow stayed with this person.” That is borrowed trust.

Discovery is a funnel, and Spotify is weaker at the top of it.

The stage where Spotify is weakest is not the stage where it has the most advantage.



Podcast loyalty is relational, and people pay for it.

Two independent sources point the same way: attachment in podcasting forms around people, not subjects.

12,000+

listener reviews of long-running US podcasts analysed in a 2025 study. It found the host and content played a pivotal role in developing and maintaining parasocial rapport characterised by intimacy and trust.

21%

of weekly podcast consumers gave money to support a podcast last year.

Of those donors,

47%

said they gave to support a host or other individual — not the show, and not the subject.

WHICH MAKES “SIMILAR TOPIC” A WEAK RECOMMENDATION PRIMITIVE. STRONGER ONES ALREADY EXIST IN THE DATA:

A guest you have listened to before

Recommended by a host you follow

Three shows you follow have interviewed this person

Listeners who finish your two most-played shows also finish this

Borrow trust from relationships the listener already values, instead of asking the algorithm to manufacture it from scratch. That is what the model in the next section is for.

Borrowed trust doesn't have to mean view counts.

Every platform lends a listener someone else's judgement. The question is whose judgement, and how well it maps onto the person deciding.

CROWD PROOF

“2.1M views”

What YouTube lends you. Instantly legible, and enough to act on in a second.

But it counts how many people clicked, not whether this is for you — and it tends to compound toward creators who are already large. Spotify has chosen not to surface it, and this deck doesn't argue that it should.

SOCIAL PROOF

“A friend sent it to me”

Word of mouth remains one of the most cited routes into a new show.

But Spotify has very little social graph to draw on, and it doesn't reach the moment someone is actually deciding.

COHORT PROOF

“Listeners like you stayed”

Personalised rather than popular — the judgement of people who value the same thing you value in a host you already follow.


It works the same for a creator with 400 listeners as for one with four million, because the cohort is defined by the reason, not by the size.

The asymmetry worth noticing: a YouTube view counts a click. Spotify can count someone finishing forty-five minutes and returning three weeks later unprompted. That is a deeper behavioural signal — and none of it is currently lent back to the listener making the decision.

The reason line is the product.

TODAY

Episodes you might like



PROPOSED

Interface mockups, not screenshots.

Here's why this is for you

-  Because you save founder interviews with specific examples more than general career advice. 22 min.
-  Listeners who return to Emma Chamberlain also become repeat listeners of this host. 31 min.
-  Because you asked for early-stage perspectives — this is specifically about finding your first customers. 18 min.

And the listener can correct it:

Right reason

Liked the creator, not the topic

More practical, less motivational

Start behavioural, graduate later: these are deliberately behavioural rather than psychological — “because you finish these” is observable and checkable, where “because you value candour” can read as being analysed. A reason that proves right makes the next unfamiliar recommendation cheaper to accept, and a wrong one is visible enough to be corrected. Over time the trust being borrowed shifts from the creator to Spotify itself.

This is not equally true of every podcast.

Saying so up front is the honest version, and it makes the narrower claim considerably stronger.

WHERE THE ARGUMENT IS STRONGEST

Unfamiliar, personality-led, long-form conversational podcasts — where a listener is choosing among people they don't already know.

This is the case where topic tells you least, the time commitment is largest, and the listener has no basis for trusting the host.

WHERE IT IS WEAKER, AND WHY

Daily news subject and habit carry the return, not the host

Sports event-driven — they came because the game happened

Serialised narrative story continuation is the reason to come back

Shows already followed that is continuity, not discovery

For these, better ranking and better continuity are the right investments — not this.

So this is a layer, not a replacement. It sits on top of ranking, for one specific decision: whether to risk an unfamiliar hour on a person you don't know yet. Everywhere else, the existing system should keep doing what it does.


● THE EMERGING-CREATOR PROBLEM

Style features exist on episode one. Behavioural confidence doesn't.

The same signal that helps a listener risk an unfamiliar hour is what an unknown creator needs in order to be found at all.

Illustrative shape, not measured data.

how confident the model is about this creator



WHY THIS INVERTS THE PROBLEM

A recommender built on behavioural volume can't say much about a creator until that creator is already large. That is the loop every emerging podcast is stuck in: you need an audience to be recommended, and recommendation to build one.

A transcript exists the moment the first episode does. Match those attributes against listeners already known to value them, and a creator with 400 listeners can be pointed at a plausible first thousand.

This is also what resolves the paradox on page 5: borrowed trust is what makes an unknown creator a reasonable risk for a tired listener.

72% of creators name discoverability and audience growth as their single biggest challenge. That changes Spotify's creator value proposition: you don't need an audience here before Spotify can begin helping you find one.

Onboard creators with evidence of opportunity, not a list of tools.

Spotify for Podcasters already gives a new creator publishing, analytics and monetisation tools. What it can't yet tell them is whether there is an audience here for them specifically — which is the question they actually arrived with.

1

WHO ARE MY PEOPLE?

The listeners already showing durable attachment to creators with similar attributes — an audience estimate before an audience exists.

2

WHY WOULD THEY STAY?

The characteristics associated with repeat listening for that group: specificity, candour, format, pacing, length.

3

WHAT SHOULD I DO NEXT?

One concrete action drawn from that, rather than another dashboard the creator has to interpret on their own.

The first Spotify creator experience could answer the question every new creator actually has: “is there an audience for me here?” That is a different onboarding promise from a tour of the tools — and it is only possible because the listener-side model produces the answer as a by-product.

● THE DOWNSTREAM CONSEQUENCE

If it works for listeners, it becomes a creator-growth engine.

The listener product creates the signal. The creator product turns that signal into a reason to invest in Spotify rather than simply distribute to it.

Who stays, and why

Your loyal core values specificity; the segment that bounces came for the topic. Two audiences inside one download number.

What produces return, not reach

Which episodes bring listeners in, and which ones make them come back — currently indistinguishable.

What they were looking for

The queries and in-episode questions that led people to you, clustered.

Who hasn't been reached

Listeners who value what your core values and have never been served an episode.

Why it matters strategically: creator investment is currently bought with advances and guarantees, which are expensive and easy to outbid. Intelligence is a by-product of work already done for listeners, and it improves as the catalogue grows.

Detailed creator surfaces, the dashboard mockup and the emerging-creator case are in the appendix. None of it should be built before the listener hypothesis is answered.

● THE BUSINESS CASE I'D WANT TO BUILD

If the hypothesis holds, these are the mechanisms I'd size internally.

I can't size any of this from outside. What I can do is say which questions I'd answer, in what order, and with what data.

● **Unfamiliar creator → voluntary return**

Does better decision confidence actually produce more successful discovery, or only more clicks?



● **Successful discovery → followed show**

Does that exploration convert into recurring consumption, or does it stay a one-off?



● **Recurring consumption → retention and monetisation**

How much economic value does one incremental durable creator relationship produce?



● **Creator understanding → creator investment**

A creator who can see Spotify finding and explaining their audience has a reason to invest here, not merely distribute here. How much does that change behaviour?

Note: each arrow is a question, not a claim. The fourth is the creator-side consequence — detail in the appendix, and not worth building before the first is answered.

Four steps, cheapest first, and I'd stop at any of them.

Desirability before feasibility, and both sides of the market before either. Whether the reason line works is a different risk from whether Spotify can automate it.

01

Validate both sides

2–3 weeks

Targeted interviews, listeners and creators. Do listeners actually experience the decision friction on page 5? Do emerging creators recognise the audience-discovery problem on page 11? If either side is weak, narrow the thesis before building anything.

02

Prototype the explanation

3–4 weeks

Wizard-of-Oz, reasons written by hand. Three cells: no explanation, a generic one, a relationship-based one. B is the cell that matters — without it you learn that explaining helps, not that borrowed trust is what did it.

03

Test the creator value proposition

3–4 weeks

Before building a dashboard, prototype the onboarding moment. Does showing a creator who their likely audience is, and why those listeners might stay, change their understanding of Spotify and their intention to invest here?

04

Run it online for real

~1 quarter

Only once both prototypes show an effect worth automating. This is where the technical question finally gets asked, with real model output in front of real listeners.

North star — durable discovery: did I turn an unfamiliar creator into someone the listener voluntarily came back to? Measured as unprompted return at thirty days, with save/follow, completion, second-episode play and seven-day return read as leading indicators along the way.

Four ways this falls apart.

Each one is answerable early, which is the main reason the plan is ordered the way it is.

Style might not be readable from words

I'm assuming a transcript can tell you a host is candid. Tone, pacing, chemistry and humour may not survive being written down at all.

Needs offline validation against human labels before any experiment runs.

People might not want to be told why

Being shown that Spotify knows you value candour could feel useful, or it could feel like being read. I don't know which, and I don't think anyone does until they ask.

Step 1 answers this.

Topic might matter more than I think

For daily news, sports, and serialised shows, people plainly do return to a subject rather than a person. If that's most of listening, this is a smaller idea than I've described.

Steps 1 and 2 answer this.

An explanation can be wrong out loud

A bad recommendation is quiet. A bad reason is visible, and it makes the system look like it misunderstands you. That means only showing a reason when confidence is high.

A design constraint, not a test.

And the real one: this competes for the same personalisation and creator-analytics people as the generation work. Step 1 is six weeks of interviews, which is small enough to run before that trade has to be made.


Spotify built the Taste Graph.

The next opportunity is the Relationship Graph.

Twenty years of learning what you like. Podcasting is where you can learn who you value, why, and who might matter to you next.

For listeners: a platform they trust to decide what — and who — is worth their time.


For creators: a platform that can show them who their people might be, why those listeners stay, and what to do next to grow that relationship.



Where I'd start: a few weeks of listener interviews before anything is built, then stage one — scoped to two quarters using signals that already exist.

Appendix

How the graph would be built · the attribute space · the search surface · the podcast home ·
the metric · episode merchandising · listener agency · AI · roadmap tension · the creator dashboard




● WHAT THE GRAPH ACTUALLY IS

One listener, then a million of them.

The object is simple: listeners and creators as nodes, and an edge that carries the reason for the attachment rather than just the fact of it.

ONE LISTENER



Spotify already knows the three follows. What isn't legible — to the listener, and in the recommendation — is the green line: why each one holds.

ACROSS MILLIONS OF THEM

The same reason recurs across many listeners, so it stops being one person's taste and becomes a pattern the model can act on.



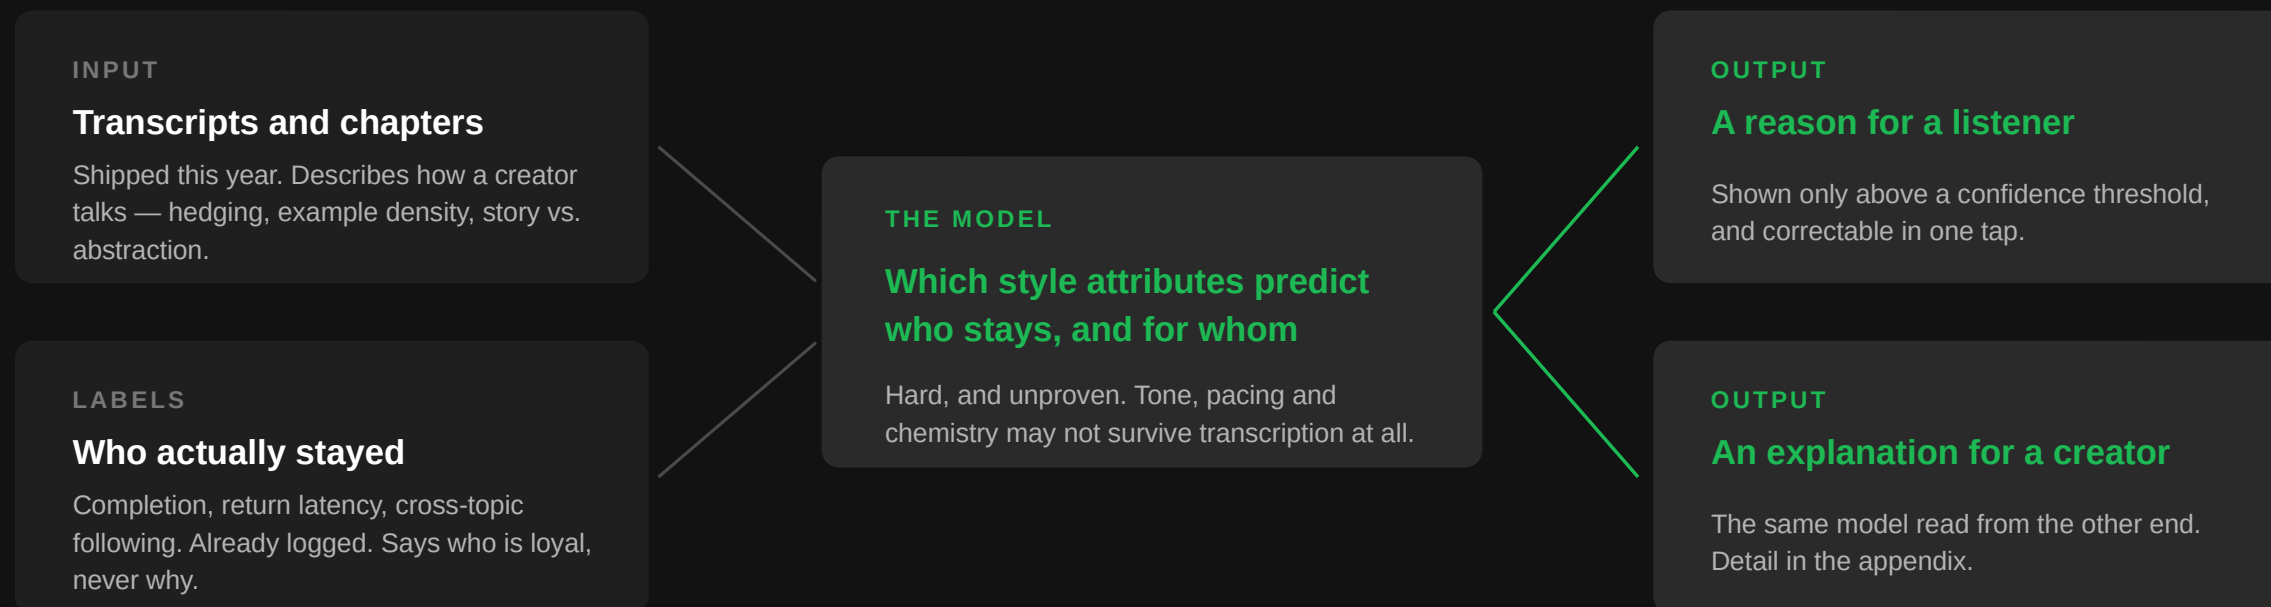
Listeners who value candour and introspection in someone they already follow disproportionately stayed with this creator too.

And that is the recommendation you can actually read: "Because you value candid, introspective hosts — listeners who return to Emma Chamberlain also become repeat listeners of this one." This isn't a personality classifier. It doesn't have to describe a creator accurately; it has to identify which observable attributes predict durable attachment, for which listeners.

● HOW IT WOULD ACTUALLY BE BUILT

Two inputs Spotify already has, and one model.


Inferring “candid” from clicks would be a stretch. Deriving candidate attributes from what was said, then testing which ones predict who stayed, is at least a tractable direction — not a solved one.



Why now: the left-hand box did not exist two years ago. Transcripts, automatic chapters and in-episode questions all shipped recently, and nothing currently reads them for style. The labels have been sitting there the whole time — what was missing was the features to pair them with.

Creators cluster by how they talk, not what they cover.

Two axes drawn from transcript features. Illustrative placement — the point is the shape, not the coordinates.



WHAT THIS SHOWS

Same category, opposite ends

How I Built This and another founder interview show share a genre and share almost nothing else. Topic put them in the same row of the library.

Different category, adjacent

A comedy confessional sits beside a personality-led interview show. Topic similarity alone wouldn't explain why they belong together.

Affinity is a neighbourhood

What a listener returns to is a region of this space, not a row in a catalogue. That region is what the graph stores.

What it looks like to search the way the need arrives.

Searching the way the need actually arrives, answered by the graph.

Interface mockup, not a screenshot.

something calming after work

Anything Goes · Emma Chamberlain

Because you return to candid, introspective hosts.

from 6:40 · 22 min

The Long Way Round · ep. 62

Because you finish unhurried, low-stimulation hosts more often than any other kind.

from 14:20 · 31 min

Nightshift · ep. 12

Because listeners who value what you value in the hosts you already follow stayed with this one too.

from the start · 18 min

WHAT'S DIFFERENT HERE

- The query is a sentence about a situation, not a genre.
- Three results, not a wall of forty plausible ones.
- Each says why it is here, so the listener can judge it.
- The answer is a passage with a timestamp, not a 45-minute show.
- The third result is a creator you've never heard of, recommended on borrowed trust rather than popularity.


Two things the podcast home gets wrong.

Both are ordering problems rather than ranking problems, which is why better prediction doesn't fix either.

ONE · IT MERCHANDISES BY INTEREST, NOT BY MOMENT

Knowing someone likes business doesn't tell you what they want at 8:20am on a 25-minute walk.

SHARE CHOOSING AUDIO VS VIDEO, BY ACTIVITY



Audio in green, video in grey. The better question isn't what a listener is interested in — it's what kind of listening moment they are in.

TWO · IT PUTS NOVELTY BEFORE CONTINUITY

Podcast listening is mostly ritual with occasional exploration. For most sessions the first jobs are practical, and exploration is the last of them, not the first.

- Did a show I follow release something?
- Where did I leave off?
- What did I save?
- What's next in my queue?
- Only now: what else might I like?

72% of heavy listeners use one platform every time, 84% prefer ongoing series, and 48% follow only one or two shows. Continuity is the product; novelty is the garnish.

Spotify's own researchers have already named this problem.

From the paper describing the deployed contextual-bandit system on the Spotify homepage.

The current reward formulation emphasises short-term engagement — whether the user streams from Home. The paper identifies multi-objective rewards balancing short-term engagement with long-term satisfaction as a future direction.

WHAT A STREAM CAN MEAN

- A curiosity click that went nowhere
- An accidental start
- “I played it because you wouldn't stop showing me this”
- Someone else using the device

WHAT WOULD ACTUALLY INDICATE TRUST

- Completion or meaningful dwell after the click
- Save or follow after exposure
- Intentional return to the show
- Whether it increased confidence in Spotify's judgement

A recommendation can generate a stream and still damage the listener's perception of Home. That is the gap between clicking and trusting — and it is why the metric I propose on page 15 is return, not click-through.

The episode is under-merchandised for the size of the ask.

Long-form content asks for an hour of attention with remarkably little pre-consumption evidence that the hour is worth it.

WHAT HAPPENED WHEN SPOTIFY ADDED CHAPTERS

+88%

chapter-initiated plays in the first month of a limited rollout

Adding chapter titles to episode descriptions also improved sparse-retrieval R@50 by 24% offline — navigation and findability moved together.

WHAT CLIPS ACTUALLY DO

81% / 84%

say clips at least sometimes lead them to an episode / to becoming a regular listener

Clips aren't a short-form feed. They are risk reduction — they expose the host's style and a representative payoff before the full commitment is asked for.

WHAT AN EPISODE DECISION SURFACE COULD CARRY

A concise “what you'll get”

not a marketing blurb

One high-signal preview moment

chosen for this listener

Chapters and the most-discussed idea

already generated

Why it is relevant to you

the reason line

Expected time, and whether it fits

your current session

Guest and relationship context

who you already trust

Personalisation needs agency, not just accuracy.

Spotify's own research describes the problem, and the transparency literature describes the constraint on solving it.

SPOTIFY RESEARCH, 2026

Users influence recommendations only through indirect signals — likes, streams, shares, saves — and it is difficult to steer the system in a new direction directly.

This matters more for podcasts, because a play is often not a preference: research, curiosity, sleep, controversy, or someone else using the device.

KIZILCEC, CHI 2016

Transparency is not automatically beneficial at every level. Appropriately designed explanations foster trust; too much or poorly calibrated disclosure undermines it.

So the goal is decision-level explanation — enough to judge the recommendation — not exposing the model itself.

THE MODEL OF PERSONALISATION THAT FOLLOWS FROM BOTH

Not “the algorithm knows you.” Rather: the algorithm has a hypothesis about you, and makes that hypothesis legible and correctable.

Not why I'm interested

Don't use this listen to personalise me

Less of this topic

Never this show

Audio only

Listeners want AI around the human, not instead of it.

The receptivity data splits cleanly, and it splits in a way that has direct product consequences.

69%

approve of AI used for research
and for brainstorming

25%

approve of an AI-generated
voice hosting the podcast

62%

see AI as at least somewhat a
threat to podcast credibility

48%

of heavy listeners would be less
likely to continue with a favourite
show that used AI voices

WHERE AI CLEARLY EARNS ITS PLACE

- Chapters, summaries and key moments
- Semantic search inside an episode or a listening history
- Personalised previews and relevance explanations
- Matching content to available time and activity

WHAT THIS IMPLIES FOR THE ROADMAP

Use AI where humans do not create the core value. Preserve the relationship; automate the friction around it.

That is a real constraint on generated audio as a growth strategy, and it is the tension the next page is about.

Two bets announced the same morning that pull in different directions.

THE GENERATION BET

Personal Podcasts

Private, AI-generated audio tailored to each listener.
Söderström's framing: access, then personalization, now generation.

THE LOYALTY BET

Memberships


Creators sell subscriptions to their most dedicated fans.
This pays off in proportion to how strongly listeners attach to specific humans.

The gap: generated audio gives a listener nothing to be loyal to except Spotify itself. That may be fine for engagement — but Memberships are a creator-loyalty business, so the two can't simply substitute for each other. Both need a shared model of attachment.

What a creator opens on Monday morning.

None of this is survey data or a demographic breakdown. It is the same graph, read from the creator's end. *Interface mockup, not a screenshot. Figures illustrative.*

WHY LISTENERS STAY WITH YOU




Derived from what your loyal listeners have in common, and confirmed by the reasons they tapped as right.

WHAT THEY WERE LOOKING FOR WHEN THEY FOUND YOU



WHERE THE LOYALTY WAS MADE

Episode 41 · 12:04–16:20



Converted more first-time listeners into repeat listeners than any other passage you've published.

WHO YOU HAVEN'T REACHED

180,400

listeners who value candour and self-disclosure in a host they already follow, and have never been served one of your episodes. The attribute that would reach them is the one you use least in your titles.