

Voices of Authority?! Trust in Human and AI-Generated Educational Podcasts

By Andreas Hebbel-Seeger*

This paper examines the extent to which awareness of an educational podcast's provenance, whether human-produced or AI-generated, shapes perceived credibility, trust, and content evaluation on the part of listeners. Using a quasi-experimental between-subjects design, 60 participants were randomly assigned to one of two groups; both listened to the same AI-generated audio contribution but were given different information about its origin. The findings show that participants who knew the content was AI-generated reported a significantly lower need to fact-check it and also showed a (non-significant) descriptive tendency to rate the synthetic hosts as more trustworthy and knowledgeable than participants who believed they were listening to a human-produced episode. These results partially contradict the AI disclosure literature, though they can be accounted for by the sample's AI-affine, highly educated profile and by an experience-driven "competence heuristic" toward generative AI. Owing to the limited sample size and low variance in AI-related attitudes, potential moderating effects could not be meaningfully explored. The study brings together podcast research, parasocial relationship theory, and AI trust research, offering a first empirical impetus for the educational-media examination of AI-generated instructional content.

Keywords: Podcast, Trust, AI, Disclosure

Introduction

Listening to serial audio or audiovisual content via subscription platforms such as Apple Podcasts and Spotify is no longer a niche activity. Current research shows that roughly one third of the German population consumes podcasts on a monthly basis; predominantly as background listening while exercising, commuting, or doing household tasks (Gattringer, 2024). Studies with an international scope report comparable figures despite cross-cultural and national differences (Gohil, 2025). A wide array of motivations drives this use, and the supply is correspondingly diverse in terms of genre and format (McClung & Johnson, 2010; Markman, 2012). Within this landscape, podcasts with an explicitly educational or instructional focus occupy the lower end of the demand spectrum (Edison Research, 2024), embedded in varied usage contexts ranging from informal self-directed learning to blended learning arrangements (Carvalho et al., 2009; Engzell et al., 2025; Meden et al., 2024).

In informal settings especially, trust in the scholarly quality and factual accuracy of what is being communicated is a basic prerequisite, yet one that listeners can typically only infer from formal cues. These include information about a host's credentials (academic qualifications, professional titles, or field of employment), the institutional

*Professor for Media Management, Macromedia University, Campus Hamburg, Germany.

affiliation of the podcast (e.g., with a university or a broadly recognized employer), or the reputation of the platform itself. Trust also develops longitudinally: as listeners accumulate experience with a podcast, they build a parasocial relationship with its host (Horton & Wohl, 1956; Vilceanu, 2025).

Research suggests that it is precisely the “human” dimension of a podcast, verbal slips, hesitations, mistakes, that contributes to the formation of these parasocial bonds and, by extension, to trust in the content (Maloney Yorganci & McMurtry, 2024; Nadora, 2019). Machine-generated podcasts, produced, for instance, using generative AI tools such as Google’s NotebookLM, lack this quality, even as synthetic voices grow increasingly adept at mimicking natural human speech patterns.

The propensity of AI-generated content to contain errors is well documented, including the phenomenon of “hallucination”, the generation of content that sounds plausible but does not hold up under scrutiny (Huang et al., 2025; Xu et al., 2025). Studies show that awareness of this risk produces strikingly divergent responses among listeners: from broad uncertainty and deliberate avoidance of generative AI at one extreme, to a largely uncritical acceptance driven in no small part by convenience at the other (Shao, 2025; Gnambs et al., 2025; Schäfer et al., 2024).

Against this backdrop, the present exploratory study investigates whether and how knowledge of an audio contribution’s provenance, as a product of human authorship or as a machine-generated artefact, affects trust in its content. Methodologically, the investigation draws on a quasi-experimental between-subjects design based on a random sample.

We first discuss podcasting as a subject of inquiry by tracing its technological development and terminological origins before describing the current podcast ecosystem. We then narrow the focus to educational podcast use and examine, on the one hand, how generative AI tools are transforming podcast production, up to and including quasi-automated content generation, and, on the other hand, how listeners’ prior attitudes and experiences with AI come into play during reception, drawing on the concept of parasocial relationships. This theoretical groundwork sets the stage for the presentation and discussion of our own study, including its design, conduct, and findings.

Podcasting

Technology and Origins

Podcasting refers to the production and distribution of serial audio or audiovisual content delivered online and available via RSS feed subscription (Kluckhohn, 2009). Several defining characteristics set podcasts apart from other audio formats. Their serial structure creates continuity and fosters ongoing audience engagement. Online distribution enables global reach without recourse to traditional broadcast infrastructure. And the RSS subscription model means that new episodes are delivered automatically to subscribers, encouraging regular listening habits (Linares et al., 2018).

Podcasts can be consumed across virtually any device (desktop computers, laptops, tablets, and smartphones), with mobile devices offering the greatest flexibility.

Smartphones in particular enable on-the-go, incidental listening unconstrained by time or place: while exercising, doing housework, traveling, or commuting through the city (Hebbel-Seeger, 2009).

Although podcast-like concepts can be traced back to audio tutorials distributed on portable cassette players such as the Sony Walkman (Palenque, 2016), the modern term derives from Apple's "iPod," a portable digital audio player launched in 2001. This etymology reflects the medium's roots in the convergence of portable digital audio technology with internet distribution at the turn of the millennium. "Pod" functions (among other readings) as an acronym for "play on demand," while "cast" comes from the second syllable of "broadcasting" (Berry, 2006, 2016a).

The technological developments that have progressively lowered the barriers to podcast production, a process now further accelerated by AI tools, have, in conjunction with the global reach of online distribution, fuelled both steady growth in supply and continued expansion of the audience (McKenzie, 2019; Rime et al., 2022; Ali et al., 2025). McKenzie's (2019) analysis of global science podcast output between 2004 and 2018 documents exponential growth in both the number of shows and the total number of episodes; a trend that has continued across genres ever since.

The podcasting ecosystem today is dominated by a relatively small number of major platforms. Spotify and Apple Podcasts together command roughly two thirds of the global market, each holding approximately one third of listeners (Gohil, 2025), effectively constituting a duopoly. Figures for the German market confirm Spotify's leading position, with Apple Podcasts, the media libraries of Germany's public broadcasters, and Audible trailing at some distance (Gattringer, 2024). Data from our own podcast production studies broadly mirror this distribution, with Spotify and Apple Podcasts running neck by neck (Hebbel-Seeger, 2021; Hebbel-Seeger & Strauß, 2022; Hebbel-Seeger & Horky, 2026).

For producers, podcasts offer a wide range of value-creation and impact possibilities, from brand communication and marketing to educational media. Beyond content distribution, they also serve as platforms for self-presentation and, much like social media, for interaction with audiences and even co-creation:

"Podcasting is not simply a new way to distribute audio recordings; it's a form of expression, interaction, and community building" (Geoghegan & Klass, 2007, p. 167).

For listeners, podcasts accordingly address needs for entertainment and learning, and provide a platform for identification. Contemporary podcasting is thus far more than a distribution channel; it creates spaces for identity formation, community building, and sustained engagement between producers and audiences alike.

Podcasts as an Educational Medium

The relative ease of podcast production and distribution, combined with the flexibility of how and when they can be consumed, makes the medium attractive for informal learning as much as for organized educational settings. This flexibility extends beyond time and place: listeners can consume content whole or in segments, replay episodes as often as needed, and adjust playback speed; an affordance that matters

because learners vary considerably in the time and repetition they require to acquire and consolidate material (Hebbel-Seeger, 2021).

By allowing learners to set their own pace, podcasts support self-regulation and self-efficacy (Erabo et al., 2024). According to Enríquez et al. (2023), podcast-based learning enables students to take ownership of their learning process (p. 898); a shift that the authors associate not only with greater autonomy but also with motivational benefits.

Evans (2008) finds that students particularly value the time savings when reviewing material and the situational flexibility that podcasts afford. A study by Meden et al. (2024) suggests these potential benefits do translate into practice: analyzing survey data from over 600 podcast listeners, the authors conclude that more than half of respondents learn in a self-directed and deliberate way, while around a third learn incidentally. Yet podcasts do not realize their educational potential automatically. Engzell et al. (2025) show that successful integration into higher education depends substantially on how well educators embed podcasts within their pedagogical frameworks. Drawing on their findings, the authors also offer practical guidance on podcast design: effective educational podcasts are entertaining, run roughly 20–30 minutes, and combine analytical commentary with concrete, application-oriented examples; observations that align with Blumers's (2025) expert-interview-based study.

These advantages notwithstanding, podcasts in educational contexts come with real limitations. Unlike written texts, they do not incorporate inline references or citations; sources are typically relegated to show notes or supplementary materials, making it difficult for listeners to evaluate the quality of evidence in real time. Moreover, podcasts generally lack the quality assurance mechanisms of established educational media (peer review, editorial oversight and so on) owing to their informal character (Shaw et al., 2025). This places a particular responsibility on educators who use third-party podcasts in their teaching to vet the material carefully.

Accessibility presents a further challenge. The frequent absence of transcripts or audio descriptions effectively excludes listeners with hearing impairments or those who prefer or need text-based formats (Chelsey, 2021; Mohale, 2024); though this limitation is gradually being eroded by the improving quality and accessibility of external speech-to-text tools in the generative AI era. Additionally, research points to the risk that podcasting may reinforce digital inequalities, as access and use vary with technological literacy, internet connectivity, and device availability (Kakhki et al., 2025; Galily et al., 2024).

Beyond reception, there is growing interest in podcast production by students as a pedagogical activity in its own right. Research shows that such settings productively develop communication skills, teamwork, research competencies, and media literacy, while simultaneously encouraging more active engagement with the subject matter at hand (Hernandez-Lopez & Mendoza-Jimenez, 2025). This dimension is, however, beyond the scope of the present study, which is exclusively concerned with reception effects.

Emotional Proximity and Trust

Because the spoken word in a podcast lacks the source references and citations that anchor credibility in academic written texts, the affective dimension of trust takes

on particular weight in how listeners evaluate what they hear. Alongside institutional signals, such as a podcast's affiliation with a university, listeners may be guided by recommendations from people they know or from influencers (Ao et al., 2023; Sardar & Vijay, 2025), by ratings and reviews (Sung et al., 2023; Wang et al., 2022), or by algorithmically curated suggestions (Li et al., 2024). But above all, trust in a podcast is anchored in its host (Mayer et al., 2024).

Repeated exposure, perceived similarity, and personal self-disclosure on the part of the host generate an emotionally meaningful bond with listeners, who come to attribute friendship, trustworthiness, and credibility to that person (Brinson & Lemon, 2022). When hosts share personal thoughts and emotional reactions, or directly address their audience, they activate a relational schema that fosters intimacy and trust. The perceived authenticity and communicative immediacy of the podcast medium encourages listeners to identify with the host and reduces the sense of social distance, even in the absence of any real feedback channel. Hosts' emotional states and dispositions are experienced vicariously by listeners, who gradually begin to perceive them not as distant media personalities but as people they personally know.

Repeated listening, particularly to high-frequency podcasts, integrates the podcast into everyday life, turning the host into an omnipresent social reference point, and deepening what listeners experience as an emotional bond, accompanied by growing trust (Dibble et al., 2016).

Vilceanu (2025) provides systematic evidence for the central role of the podcast host in building and sustaining parasocial familiarity and trust, and for the relationship between listening time and bond strength, through an analysis of over 12,000 listener reviews. Despite the inherently one-sided, mediated communication structure, listeners consistently develop feelings of closeness, trust, and connection toward hosts.

Strikingly, it is not flawless articulation that makes a host seem credible or trustworthy, but rather the authenticity conveyed by uncertainty, verbal slips, and mistakes. Maloney Yorganci & McMurtry (2024) identify seven authenticity markers that promote parasocial relationships, one of which is explicitly "Imperfection": hosts who stumble, digress, or err come across as more human and approachable. Spontaneity and personal admissions create the sense of really knowing the host, while ordinariness and perceived similarity foster feelings of friendship.

Nadora (2019) grounds this in the podcast's character as a "de-professionalized" medium (*ibid.*, p. 4) in which hosts function as "ordinary experts" (*ibid.*, p. 17; Tolson, 2010) whose imperfections are precisely what make them relatable. Moments of ambiguity, self-correction, and reflexivity are not merely tolerated by listeners; they actively strengthen the parasocial bond.

The choice of listening device also plays a role. Consuming podcasts through headphones heightens the sense of immersion and presence compared with loudspeaker listening (Liebermann et al., 2022), which in turn affects perceived proximity to the host. And the content of a podcast rarely feels entirely foreign to listeners: it tends to connect with their personal experience; whether through intrinsic interest or through an extrinsic framing in a school or university context:

“Podcasts are listened to in an intimate setting (headphones), utilising an intimate form of communication (human speech). Furthermore, in many cases podcasts are presented by people from within a listener’s own community of interest or by people s/he may already have a relationship with via social media, and are frequently recorded in a podcaster’s own personal or domestic space” (Berry, 2016b, p. 666).

Podcast Production and AI

Conventional podcast production typically involves four phases: (1) concept development and scripting, (2) audio recording, either in-person or online, (3) post-production, which may involve editing the raw recording as well as integrating external material, and (4) distribution, including publishing on relevant platforms and promoting via social media.

Since 2022, the rapid advancement of generative AI has substantially expanded what is possible in podcast production (Desmedt et al., 2025). The range of available tools runs from AI-assisted aids that support human creators at specific points in the workflow all the way to fully autonomous systems capable of generating audio content from text prompts alone.

Post-production, phase 3, typically accounts for a disproportionate share of the overall effort: beyond content-level editing, it encompasses noise reduction, level adjustment, and speech optimization (Hebbel-Seeger & Strauß, 2022). AI-powered tools have made these steps considerably faster and more accessible in recent years (Bruns, 2008; Zou et al., 2026).

Tools such as Adobe’s Podcast Enhanced Speech¹ and Descript² employ deep neural networks, based in particular on transformer architectures and diffusion models, to eliminate background noise in real time and bring speech recordings close to studio quality. Descript additionally enables transcript-based editing: audio is automatically transcribed, and changes made at the text level are applied directly to the audio file. Further powerful features include automatic loudness normalization, AI-assisted removal of filler pauses, and voice-cloning functionality that allows corrections to be made without re-recording (the so-called “Overdub” feature).

Beyond recording and post-processing, large language models such as ChatGPT, Claude, or Gemini can support the substantive and editorial work surrounding an episode, conducting research on topics or guests, drafting interview questions, generating show notes and episode descriptions, and producing visuals for promotional use.

Text-to-speech tools such as ElevenLabs³ can then be used to give a voice to LLM-generated scripts (Hasanabadi, 2023), making it possible to render a dialogue between multiple AI-generated conversational voices as a finished audio file.⁴ The quality of synthetic voices has now reached a level at which lay listeners can barely

¹<https://podcast.adobe.com/>

²<https://web.descript.com/>

³<https://elevenlabs.io/>

⁴Episode 101 of the German-language podcast “LectureCast” illustrates this approach: text outputs generated via prompts from the LLMs ChatGPT, Claude, and Gemini were each assigned synthetic voices, and the resulting audio files were combined into a multi-voiced discussion in a single audio file (<https://open.spotify.com/episode/5kFIqI0ICWKdBcZNCsh8W3?si=qDRzhb-CT6SjUiIOCgpMaw>).

distinguish them from human recordings; a phenomenon discussed in the literature through constructs such as “naturalness,” “human-likeness,” and “eeriness” (Abdulrahman & Richards, 2022).

The integration of LLMs into editorial workflows is not without controversy, however. It sits in tension with journalistic quality standards: while generative AI tools promise efficiency gains and scalability, they also increase the risk of homogenized content and factual hallucinations (Bender et al., 2021), demanding at minimum a critically curated approach from producers.

The current high-water mark of AI integration in podcast production is represented by systems that generate complete, dialogic podcast episodes from prompts or linked documents. The most prominent example is Google’s “NotebookLM,”⁵ first released in mid-2024 and continuously updated since. NotebookLM analyzes user-supplied documents, local or linked, combines them with information from its trained knowledge base, and produces a two-person dialogue between a synthetic female and a corresponding male voice in podcast format, simulating natural conversational dynamics including follow-up questions and moderating transitions.

This kind of content generation draws on several interlocking technologies: an LLM to produce a script from the source documents (Sigurgeirsson & King, 2023), a text-to-speech system to synthesize natural-sounding voices (Hasanabadi, 2023), and prosodic models to simulate conversational features such as emphasis, pauses, and affective tone (Wang & Székely, 2024).

The output of such a process need not stand alone; it can also be combined with recordings of real speakers in a mash-up format.⁶

Trust in Artificial Intelligence

“The rise of generative artificial intelligence (GenAI) marks a turning point in the development of human cognition. Machines are no longer limited to retrieving information; they simulate thinking, produce fluent responses, and generate seemingly original content—tasks that were once considered uniquely human” (Gerlich, 2025, p. 932; Translation from German).

As the surface differences between human and AI-produced podcasts continue to blur, the question of trust takes on new dimensions. It confronts podcast producers who draw on generative AI to varying degrees, and it confronts listeners for whom the reception of partially or fully AI-generated content brings an additional variable into play: their own prior experiences with and attitudes toward artificial intelligence. What was once a matter of trust rooted in human authenticity now competes with a closeness technically constructed by generative AI, and with a broader stance toward the technology itself.

⁵<https://notebooklm.google/>

⁶Episode 114 of the German-language podcast “LectureCast” provides an example: a conversation recorded between two human participants was supplemented by dialogue passages generated with NotebookLM (<https://open.spotify.com/episode/61TyUZRYq5f4EHvPPN4o5w?si=XOY a5K2jRxKxevnbI26qTw>).

This is a particularly interesting tension because, as noted above, trust in human communicators does not appear to hinge on the absence of error, quite the contrary. With AI-generated content, the dynamics are different. LLMs are prone to hallucination: they produce output that sounds coherent and plausible but cannot be factually verified (Huang et al., 2025). This appears to be a structural feature rather than an incidental flaw. Xu et al. (2025) demonstrate that LLMs cannot learn all computable functions and will therefore inevitably hallucinate when used as general-purpose problem-solvers. Shao (2025) puts the hallucination rate in benchmark tests at between 5% for general queries and up to 29% for specialized technical questions.

Whether this represents a high or low error rate relative to human performance is a question we set aside here. What seems safe to assume is that all users of generative AI have encountered incorrect outputs at some point, from which, in principle, a healthy skepticism ought to follow. In practice, however, generative AI tools are comparatively easy to access and use, and the fluency and apparent authority of their responses make the output convincing enough for many users to encourage surface-level processing at the expense of epistemic vigilance.

“What intensifies this development further is the growing trust in GenAI systems. This trust, which is often characterized by speed, coherence, and convenience, is rarely based on a thorough understanding of how these systems function. The more users trust the results of GenAI, the less inclined they are to question them. In this way, they relinquish the burden of verification and assessment, replacing analysis with acceptance. Under these conditions, trust is not the result of careful deliberation, but a shortcut around it” (Gerlich, 2025, p. 933; translation from German).

People tend to use AI systems even when they harbour diffuse anxieties or general unease about them. Chen et al. (2025) describe this as “intermittent neutralization”, a process in which acceptance, rejection (driven, for example, by data protection concerns), and information illusions blend into a cognitive dissonance that progressively dampens initial skepticism.

Survey data bear this out. A survey of over 30,000 consumers and employees across eleven European countries finds that non-users are markedly more skeptical than users: 74% of non-users express concerns about data misuse, compared with 62% of users; and while 57% of generative AI users trust it to deliver reliable results, only 33% of non-users share that confidence (Corduneanu et al., 2024).

A similar disconnect between self-assessed knowledge and actual attitudes emerges in Scantamburlo et al. (2023). Drawing on survey data from 4,006 people across eight European countries (Germany, France, Italy, the Netherlands, Poland, Romania, Spain, and Sweden), they find that half of respondents hold a positive attitude toward AI despite consistently low self-rated AI competence. This aligns with Gnambs et al. (2025), whose analysis of 1,098 German adults finds broad acceptance of AI use even among people with little hands-on experience.

A more differentiated picture emerges when AI use is examined by context. Schäfer et al. (2024), drawing on telephone interviews with 1,037 young adults in Germany, find, in contrast to the studies cited above, considerable skepticism toward generative AI in scientific contexts, even as general trust in science stands at a comparatively high 56%. Lack of accuracy, bias, and the potential for abuse in

pseudoscientific applications are identified as the main risks associated with using generative AI in scientific communication.

How Knowledge of AI Provenance Shapes Podcast Reception

Credibility and trust in podcasts rest, beyond external markers, primarily on the perceived authenticity of the host: imperfections, spontaneity, and personal self-disclosure that listeners read as signs of approachability and humanity. These authenticity signals stabilize parasocial relationships and constitute the foundation of trust in the content conveyed. Generative AI can now simulate these signals convincingly enough that listeners often cannot distinguish synthetic voices from human recordings.

Fully or partially machine-generated podcasts are no longer a future scenario, they are a present media reality. The question of how listeners respond when they know (or do not know) that a podcast was machine-generated is therefore less a technological matter than one belonging to educational media studies and communication psychology.

Human expertise as expressed in educational podcasts now competes directly with AI. Tolmeijer et al. (2022) find that participants rated AI as more capable than humans on the tasks studied, and that with experience participants grew more willing to follow an AI's advice or accept its decisions over a human expert's:

“Our results indicate that people perceive AI to be more capable than humans for the given tasks, but place somewhat higher moral trust in humans. The capable trust in AI is apparent in participant reliance behavior: as they do more missions, they are more likely to take an AI's advice or accept an AI's decision than a human expert's” (Tolmeijer et al., 2022, p. 2).

At the same time, studies point to context-specific skepticism toward generative AI outputs that appears particularly pronounced in scientific and educational settings (Schäfer et al., 2024). Prior experience with AI moderates attitudes toward it, suggesting that reactions to AI-generated educational podcasts are shaped less by the technology per se than by listeners' prior knowledge and expectations.

Two further theoretical considerations sharpen the framing of our research questions. First, trust in AI-generated content is unlikely to be uniform across genres: whether a podcast primarily conveys analytical commentary or relies on strict factual accuracy may itself shape how much verification listeners feel is warranted, and how readily they extend trust to a synthetic host (cf. Li et al., 2025, on the moderating role of content type). The present study is deliberately scoped to a single genre, an informational, lower-stakes higher-education topic, and the findings should be read as bounded by that choice rather than as evidence about AI-generated podcasts in general. Second, the very fact that listeners form impressions of synthetic hosts' trustworthiness and expertise at all points to what Balkin (2017) calls the “homunculus fallacy”: the tendency to attribute human-like intention, competence, and character to non-human agents. That listeners in our sample evaluated AI-generated voices in much the same relational terms they would apply to human hosts, see Section 4.3, suggests that this fallacy operates even under conditions of full AI disclosure, and that it may be a more productive lens for future hypothesis development than a simple human-versus-AI binary.

Against this background, the present study addresses two research questions:

- (1) Does it matter to podcast listeners whether they know the episode was produced by a human or by AI; and if so, does this knowledge alter perceived credibility, trust, and evaluations of content quality?
- (2) To what extent do individual AI attitudes and usage experience moderate these effects?

Drawing on the theoretical considerations and empirical findings reviewed above, we derive two research hypotheses:

- (H1) Knowledge of a podcast's provenance as human-made or AI-generated affects credibility and trust in content quality: listeners who know the content is AI-generated will attribute higher credibility to it and trust it more than listeners who believe they are consuming a human-produced episode.
- (H2) The extent of participants' own AI use and their personal attitude toward generative AI tools modulates the strength of the effect posited in H1: the greater the prior experience and the more positive the baseline attitude toward AI, the greater the trust in content identified as AI-generated.

Methods

To test these hypotheses we opted for a quasi-experiment (Behi & Nolan, 1996) structured as a between-subjects design (Charness et al., 2012). Participants were randomly assigned to one of two groups, both of which received the same audio contribution as the treatment stimulus but differed in what they were told about its origin. One group was explicitly informed that the episode was AI-generated; the other received no such information and was left to assume they were listening to a conventionally produced, human-authored audio contribution. Group assignment was carried out using the randomization function of the survey platform, which allocated each participant to one of the two conditions with equal probability at the point of survey entry.⁷

The audio contribution was generated using Google's NotebookLM on the basis of a text by Hebbel-Seeger & Vohle (2026). It was subsequently manually trimmed to 5 minutes and 28 seconds to minimize participant burden and reduce dropout risk. The resulting episode features a "human-sounding" synthetic female speaker and a comparably natural-sounding synthetic male speaker discussing the content of the source publication; a format modeled on the expert-dialogue podcast structure common in educational podcasting.

Following the stimulus, participants completed a questionnaire capturing four outcome constructs: perceived host trustworthiness, perceived host expertise, confidence in the accuracy of the content, and the perceived need to fact-check it. Each construct was rated on a 5-point Likert scale; demographic items (gender, age

⁷Data collection and preparation were carried out with the assistance of *Aylin Smokovic* and *Benita Off*.

band, highest educational qualification) and self-reports of AI familiarity, AI usage frequency, and personal interest in AI were collected alongside the outcome measures.

The random sample was drawn from the personal networks of student assistants. The topic selected for the audio contribution, AI-generated videos in higher education, was chosen because it is directly relevant to the participants' lived experience and can reasonably be assumed to hold some intrinsic interest for them. This convenience-sampling approach was chosen for reasons of feasibility within the time frame available for this exploratory pilot study; it was considered acceptable given the study's hypothesis-generating rather than hypothesis-confirming purpose, though, as discussed below, it limits the representativeness of the sample.

The quasi-experiment was conducted online over 23 days at the end of 2025. Participants accessed it individually, at a time and place of their choosing, and were asked to listen to the audio contribution and then complete a questionnaire about it.

Results

A total of 60 valid responses were collected: 38 from women and 22 from men, the majority of whom were under 30 years of age. Fifty participants held at least a general university entrance qualification. This educational profile likely accounts for the high level of familiarity with generative AI that participants reported: 52 rated their familiarity as "high" or "very high," while only 7 expressed uncertainty or said they were "not very familiar," and just 1 indicated no familiarity at all. Interest in AI was widespread: 42 participants described themselves as personally interested to varying degrees, 13 were undecided, 4 reported little interest, and only 1 reported none (see Table 1).

Table 1. *Sample Characteristics (N = 60)*

Characteristic	Category	n	%
Gender	Women	38	63.3
	Men	22	36.7
Age	Majority < 30 years	38	63.3
Education	General university entrance qualification or higher	50	83.3
	Other / lower qualification	10	16.7
AI familiarity	High / very high	52	86.7
	Uncertain / not very familiar	7	11.7
	None	1	1.7
Interest in AI	Interested (varying degrees)	42	70.0
	Undecided	13	21.7
	Little interest	4	6.7
	None	1	1.7

The two experimental groups were comparable on age, gender, educational level, AI familiarity, and personal interest in AI. Differences in response patterns can therefore be attributed primarily to whether participants knew (Group 1) or did not know (Group 2) that the audio they had listened to was AI-generated or not.

Differences first emerged on the question of host trustworthiness. Participants who believed they had listened to a human-produced episode rated the hosts as less

trustworthy than those in Group 1 who knew the voices were synthetic. A similar pattern appeared in assessments of expertise: participants aware of the AI provenance rated the hosts' expertise higher than those who assumed human speakers. Across both groups, women were more skeptical than men, consistently rating hosts as less trustworthy and less knowledgeable.

Participants who assumed real human protagonists also expressed lower confidence in the accuracy of the content and a greater perceived need to verify it compared with Group 1. The gender difference recurred here as well: women attributed less trust to the content and saw a greater need for fact-checking.

Table 2. Mann–Whitney U Tests Comparing Group 1 (AI-Disclosed) and Group 2 (No Disclosure)

Measure	G1 M (SD)	G2 M (SD)	U	z	p	r
Host trustworthiness	3.11 (0.99)	2.68 (1.09)	480.5	1.508	.134	.202
Host expertise	3.43 (0.96)	3.14 (0.85)	459.5	1.173	.244	.157
Confidence in content accuracy	3.00 (1.05)	2.57 (0.84)	492.5	1.734	.084	.232
Need to fact-check	2.00 (1.00)	2.57 (1.07)	262.0	−2.079	.038*	.280

scale 1 = fully agree, 5 = fully disagree. Missing values were excluded listwise. (Trustworthiness: n = 56; Expertise: n = 56; Accuracy: n = 56; Necessity: n = 55). U is the U-value for Group 1 vs. Group 2. $r = |z| / \sqrt{N}$ as Effect size (Cohen: .10 = small, .30 = medium, .50 = large). * $p < .05$.

Because the data were not normally distributed, group differences were tested using the non-parametric Mann–Whitney U test (see Table 2). Only one of the four comparisons reached significance: participants who had been told the content was AI-generated indicated a significantly lower need to verify it ($U = 262.0$, $z = -2.079$, $p = .038$, $r = .280$). Given that this is the only significant result among four tested comparisons, it should be read as tentative, exploratory evidence for H1, the hypothesis that listeners attribute greater credibility to audio identified as AI-generated, rather than as a confirmed effect; the possibility that this single finding reflects chance (a Type I error) cannot be ruled out and underscores the need for replication with a larger sample. Exploration of H2 was not possible given the near-absence of within-group variance on AI attitudes and experience.

Discussion

The finding that knowing a podcast is AI-generated reduces rather than increases the perceived need to verify its content appears, at first glance, to cut against the AI disclosure literature. Based on a meta-analysis of 13 studies spanning multiple domains, Schilke & Reimann (2025) conclude that disclosing AI use consistently reduces trust, even if the effect is attenuated for people with positive attitudes toward AI or a basic trust in generative technology, it does not disappear. Wang & Huang's (2024) meta-analysis of experimental studies in news research similarly finds that articles labeled as AI-generated are perceived as less credible than human-authored ones. Koning & Voorveld (2025) add evidence from advertising research showing that AI disclosure significantly reduces trust and brand attitudes.

Our result diverges from this first cluster of findings but connects with a second, more nuanced one. Huschens et al. (2023) show that LLM-generated and human-authored texts are rated as equivalent in terms of competence and trustworthiness, with AI-generated texts even perceived as clearer and more appealing; a finding that converges with ours, despite the difference in media format. Li et al. (2025) further demonstrate that AI labels have no significant effect on perceived credibility in their experiment, and that the effect is strongly moderated by prior experience and content type. Particularly relevant for educational contexts is their finding that participants with more prior AI experience rate non-commercial, informational, and educational content as highly credible even when it carries an AI label; consistent with the education level and AI affinity of our sample.

Gallegos et al. (2026) add that AI labels change perceptions of authorship without reducing the persuasive impact of content, supporting the view that AI disclosure does not automatically trigger a trust deficit but operates in a context-dependent way. In educational settings in particular, an AI label may actually activate a kind of “competence heuristic”, as generative AI tools are widely perceived by younger, more highly educated users as capable and factually reliable (Corduneanu et al., 2024). Gerlich (2025) puts this dynamic succinctly:

“Trust in GenAI ... reflects a profound shift in epistemic preferences. Users are increasingly willing to trust AI, not merely because it performs well, but because it appears less prone to error than human alternatives” (p. 935; translation from German).

Taken together, our findings point to an experience-driven evaluation of AI outputs that aligns with Tolmeijer et al. (2022), whose work informed our hypothesis formulation, and resonates with Alieto et al. (2025), who similarly find that people tend to regard generative AI output as a less error-prone knowledge source and engage with it accordingly uncritically. Beyond general AI-affinity, two more specific features of the sample plausibly underpin this competence heuristic. First, the high proportion of participants with hands-on generative-AI experience (52 of 60 rating their familiarity “high” or “very high”) means that, for most respondents, exposure to fluent, competent-sounding AI output already preceded the study; repeated positive encounters of this kind are precisely the mechanism through which Tolmeijer et al. (2022) and Gerlich (2025) locate the growth of capability-based trust. Second, the topic of the stimulus, AI tools in higher education, sits squarely within participants’ own field of competence as students, which may have lowered the perceived stakes of being wrong and correspondingly the perceived need to verify; this would be consistent with the genre-dependency argument introduced in Section 4. We note, however, that the present design cannot separate these two accounts from one another or from the homunculus fallacy discussed above, since AI experience, topic familiarity, and anthropomorphic trust attribution were not manipulated independently. Disentangling them is a natural next step for the replication study proposed below.

The cross-group pattern of greater skepticism among women, lower trust in hosts, higher perceived need for fact-checking, mirrors findings in the literature on gender differences in AI perception. Borwein et al. (2026), drawing on a large North American sample, show that women consistently rate AI risks higher than men, a pattern attributed to general risk aversion and lower exposure to the purported benefits

of AI. Aldasoro et al. (2024) similarly report, on the basis of the Survey of Consumer Expectations, that the “Gen-AI gender gap” in generative AI use is explained primarily by lower self-assessed knowledge, reinforced by greater data protection concerns and lower baseline trust in AI providers. This structural female skepticism toward generative AI offers a plausible account of the pattern observed here, though the sample size precluded a formal significance test.

Limitations

The present study offers initial empirical evidence of differential reception effects tied to awareness of AI provenance in an educational podcast, but several methodological constraints must be kept in mind when interpreting the results.

The sample size of $n=60$ imposes strict limits on statistical power. Several substantively plausible group differences, on host trustworthiness and perceived expertise, for instance, fail to reach significance, which is most readily explained by the combination of a small sample and modest effect sizes. A larger replication study is needed before robust conclusions can be drawn. The present sample size reflects the time and personnel constraints of a student-supported pilot data collection rather than a deliberate methodological choice; recruiting and retaining a substantially larger, more diverse sample within the available 23-day window was not feasible. We treat this explicitly as a constraint on the present study rather than as a finished design, and accordingly frame our contribution as a first, exploratory step that motivates rather than substitutes for the larger replication called for throughout this section.

The sample is also markedly skewed in terms of age and educational attainment, which is reflected in the high levels of AI familiarity participants reported. This limits the external validity of the findings and leaves open whether the observed patterns would generalize to other populations, older listeners or those less comfortable with AI, for example, especially given that prior research consistently shows AI-related attitudes and trust judgments to be strongly moderated by experience and background knowledge.

The homogeneity of the sample on AI familiarity and attitudes also meant that H2, the hypothesis that AI usage and attitudes moderate trust in AI-generated content, was simply not empirically tractable. Testing it properly would require a sample with substantially greater variance on these dispositions. We want to be explicit on this point: since RQ2 was one of the study’s two central research questions, its untested status is not a peripheral caveat but a central limitation of the present study, and the empirical contribution offered here should be read as addressing RQ1 only.

Finally, because this is a quasi-experiment, complete randomization in group assignment was not achieved. While the groups appear comparable on the key variables, a selection effect cannot be ruled out.

Conclusion and Outlook

Starting from the question of how awareness or ignorance of AI provenance shapes the reception of an educational podcast, the present study brings together three research domains, podcast research, educational media, and generative AI, that have not previously been examined in this combination. As an exploratory study, it delivers an initial empirical foothold, even if most of the measured items produce group differences that fall below the threshold of significance. We take this less as a limitation than as an invitation to pursue the questions it raises. We deliberately frame the contribution this way rather than as a finished test of both hypotheses: given the constraints discussed in Section 4.4, a single small-sample study could not by itself deliver a robust answer to RQ2, and we judged it more useful to report the present, theoretically grounded pilot transparently and to specify concretely what a properly powered follow-up would need to look like, rather than to delay publication until a larger study could be mounted.

The most pressing next step is a replication with substantially larger and more diverse samples, one that would permit reliable significance testing and enable the investigation of moderating effects. Of particular interest is the question of whether and how individual AI attitudes, prior usage experience, and general trust in science shape the reception of AI-generated educational content; a question that the present study could not address due to insufficient variance and sample size.

The quasi-experimental design also invites extension in multiple directions. The degree of AI involvement could be varied systematically, from fully AI-generated audio, to AI-written scripts with human voices, to entirely manually produced podcasts, to test whether trust judgments respond more to the binary “human or AI” distinction or to the degree of AI involvement. In this regard, the model of “calibrated trust” developed by Hebbel-Seeger & Vohle (forthcoming), with reference to Lee & See (2004) and Ayoub et al. (2022) in the context of academic video use, offers a potentially useful methodological orientation. A longitudinal design would also be valuable for examining whether trust judgments shift over time as exposure to AI-generated educational content accumulates, including in light of the “intermittent neutralization” dynamic discussed in Section 3 (Chen et al., 2025).

More broadly, the study suggests that the provenance of an educational medium is not a neutral piece of information. It shapes listeners’ attitudes and expectations in ways that have practical implications for how AI-supported educational offerings are designed and communicated. For higher education didactics, a set of concrete follow-up questions arises: how does transparency about AI involvement affect learning behavior; the depth at which content is processed, the willingness to engage with it critically, the nature and extent of that engagement? These questions are as important as they are underexplored.

References

- Abdulrahman, A., & Richards, D. (2022). Is Natural Necessary? Human Voice versus Synthetic Voice for Intelligent Virtual Agents. *Multimodal Technologies and Interaction*, 6(51). DOI: 10.3390/mti6070051.
- Aldasoro, I., Armantier, O., Doerr, S., Gambacorta, L., & Oliviero, T. (2024). The gen AI gender gap. *Economics Letters*, 241, Article 111814. DOI: 10.1016/j.econlet.2024.111814.
- Ali, S.N., Immorlica, N., Jagadeesan, M., & Lucier, B. (2025). Flattening Supply Chains: When do Technology Improvements lead to Disintermediation? *arXiv*. DOI: 10.48550/arXiv.2502.20783.
- Alieto, E.O., Dumagay, A.H., Serdenia, J.R.C., Labad, E.M., Galang, S.K., & Vallejo, R.G. (2025). Attitude Toward Artificial Intelligence Among Teacher Aspirants in an Emerging AI Landscape: A Gender-Based Analysis. In R. González Vallejo, G. Moukhliiss, E. Schaeffer, & V. Paliktzoglou (Eds.), *The Second International Symposium on Generative AI and Education (ISGAIE'2025)* (pp. 499–512). Cham: Springer. DOI: 10.1007/978-3-031-98476-1_39.
- Ao, L., Bansai, R., Pruthi, N., & Khaskheli, M. B. (2023). Impact of Social Media Influencers on Customers Engagement and Purchase Intentions: A Meta-Analysis. *Sustainability*, 15(3), Article 2744. DOI: 10.3390/su15032744.
- Ayoub, J., Avetisyan, L., Makki, M., & Zhou, F. (2022). An Investigation of Drivers' Dynamic Situational Trust in Conditionally Automated Driving. *IEEE Transactions on Human-Machine Systems*, 52(3), 501–511. DOI: 10.1109/THMS.2021.3131676.
- Balkin, J. M. (2017). The Three Laws of Robotics in the Age of Big Data. *Ohio State Law Journal*, 78, 1217–1241, Forthcoming, Yale Law School, Public Law Research Paper No. 592, Available at SSRN: <https://ssrn.com/abstract=2890965> [22.06.2026]
- Behi, R., & Nolan, M. (1996). Quasi-experimental research designs. *British Journal of Nursing*, 5(17), 1079–1081. DOI: 10.12968/bjon.1996.5.17.1079.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?* FAccT'21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 610–623). New York: ACM. DOI: 10.1145/3442188.3445922.
- Berry, R. (2006). Will the iPod kill the radio star? Profiling podcasting as radio. *Convergence: The International Journal of Research into New Media Technologies*, 12(2), 143–162. DOI: 10.1177/1354856506066522.
- Berry, R. (2016a). Podcasting: Considering the evolution of the medium and its association with the word 'radio'. *The Radio Journal: International Studies in Broadcast & Audio Media*, 14(1), 7–22. DOI: 10.1386/rjao.14.1.7_1.
- Berry, R. (2016b). Part of the establishment: Reflecting on 10 years of podcasting as an audio medium. *Convergence: The International Journal of Research into New Media Technologies*, 22(6), 661–671. DOI: 10.1177/1354856516632105.
- Blumers, C. L. (2025). *Die Zukunft von Edutainment-Podcasts in Deutschland: Narrative Strategien im Wandel*. Unpublished master's thesis, Macromedia University, Campus Berlin.
- Borwein, S., Magistro, B., Alvarez, R. M., Bonikowski, B., & Loewen, P. J. (2026). Explaining women's skepticism toward artificial intelligence: The role of risk orientation and risk exposure. *PNAS Nexus*, 5(1), pgaf399. DOI: 10.1093/pnasnexus/pgaf399.
- Brinson, N. H., & Lemon, L. L. (2022). Investigating the effects of host trust, credibility, and authenticity in podcast advertising. *Journal of Marketing Communications*, 29(6), 558–576. DOI: 10.1080/13527266.2022.2054017.
- Bruns, A. (2008). *Blogs, Wikipedia, Second Life and Beyond: From Production to Produsage*. New York: Peter Lang.

- Carvalho, A. A., Aguiar, C., Santos, H., Oliveira, L., Macques, A., & Maciel, R. (2009). *Podcasts in Higher Education: Students' and Lecturers' Perspectives*. In A. Tatnall & A. Jones (Eds.), *Education and Technology for a Better World* (pp. 417–426). Berlin, Heidelberg: Springer. DOI: 10.1007/978-3-642-03115-1_44.
- Charness, G., Gneezy, U., & Kuhn, M. A. (2012). Experimental methods: Between-subject and within-subject design. *Journal of Economic Behavior & Organization*, 81(1), 1–8. DOI: 10.1016/j.jebo.2011.08.009.
- Chelsey, A. (2021). *Is There a Transcript? Mapping Access in the Multimodal Designs of Popular Podcasts*. Proceedings of the 39th ACM International Conference on Design of Communication (pp. 46–53). New York: ACM. DOI: 10.1145/3472714.3473622.
- Chen, J., Xie, W., Xie, Q., Hu, A., Qiao, Y., Wan, R., & Liu, Y. (2025). A Systematic Review of User Attitudes Toward GenAI: Influencing Factors and Industry Perspectives. *Journal of Intelligence*, 13(78), 1–46. DOI: 10.3390/jintelligence13070078.
- Corduneanu, R., Winters, S., Machalski, J., Horton, R., Sahu, R. K., & Echániz, J. (2024). Europeans are optimistic about generative AI but there is more to do to close the trust gap. *Deloitte Insights*. <https://www.deloitte.com/us/en/insights/topics/digital-transformation/trust-in-generative-ai-in-europe.html> [28.02.2026].
- Desmedt, C. P., Budts, W., De Vos, M., & Moons, P. (2025). Artificial intelligence-generated podcasts open new doors to make science accessible: a mixed-method evaluation of quality and the potential for science communication. *European Journal of Cardiovascular Nursing*, 24(6), 885–895. DOI: 10.1093/eurjcn/zvaf074.
- Dibble, J. L., Hartmann, T., & Rosaen, S. F. (2016). Parasocial interaction and parasocial relationship: Conceptual clarification and a critical assessment of measures. *Human Communication Research*, 42(1), 21–44. DOI: 10.1111/hcre.12063.
- Edison Research (2024). *The Podcast Consumer 2024*. <https://www.edisonresearch.com/wp-content/uploads/2025/07/The-Podcast-Consumer-2024-FINAL.pdf> [24.02.2026].
- Engzell, J., Norrman, C., Norberg, A. & Lundvall, C. (2025). Soundwaves of knowledge: using podcasts to facilitate learning in higher education. *Educational Media International*, 62(3), 361–379. DOI: 10.1080/09523987.2025.2533120.
- Enríquez, I.G., Cutuli, M.S., & Mancha-Cáceres, O.I. (2023). Enhancing Collaborative Learning in Higher Education through Podcast Production. *Education Sciences*, 13, 898–914. DOI: 10.3390/educsci13090898.
- Erabo, D.D., Dela Rosa, A., & Gonzales, L.J.M. (2024). Optimizing differentiated podcasts to promote students' self-regulation and engagement, self-efficacy and performance in asynchronous learning. *Journal of Research in Innovative Teaching & Learning*, 17(2), 368–390. DOI: 10.1108/JRIT-02-2024-0039.
- Evans, C. (2008). The effectiveness of m-learning in the form of podcast revision lectures in higher education. *Computer & Education*, 50, 491–498. DOI: 10.1016/j.compedu.2007.09.016.
- Galily, Y., Laor, T., & Azran, T.S. (2024). Behavioral and Psychosocial Dynamics of Engagement: The Digital Divide in Artificial Intelligence [AI]-Driven Sports Podcasts. *Behavioral Sciences*, 14, 911–923. DOI: 10.3390/bs14100911.
- Gallegos, I. O., Shani, C., Shi, W., Bianchi, F., Gainsburg, I., Jurafsky, D., & Willer, R. (2026). Labeling messages as AI-generated does not reduce their persuasive effects. *PNAS Nexus*, 5(2), pgag008. DOI: 10.1093/pnasnexus/pgag008.
- Gattringer, K. (2024). Ergebnisse der ARD/ZDF-Medienstudie 2024. Podcastnutzung 2024: Konsolidierung von Nutzungsgewohnheiten. *Media Perspektiven*, 27, 1–5.
- Geoghegan, M., & Klass, D. (2007). *Podcast Solutions: The Complete Guide to Audio and Video Podcasting* (2nd ed.). Berkeley, CA: Apress. DOI: 10.1007/978-1-4302-0473-2.
- Gerlich, M. (2025). Kognitives Offloading und die Erosion des kritischen Denkens im Zeitalter der generativen KI. In D. Schulze Heuling, A. P. Jakobi, G. S. Schaal, & M. Gerlich (Eds.),

- Generative KI und (politik)wissenschaftliches Schreiben. *Politische Vierteljahresschrift*, 66, 913–940. DOI: 10.1007/s11615-025-00631-9.
- Gnambs, T., Stein, J.-P., Zinn, S., Griesse, F., & Appel, M. (2025). Attitudes, experiences, and usage intentions of artificial intelligence: A population study in Germany. *Telematics and Informatics*, 98, 102265. DOI: 10.1016/j.tele.2025.102265.
- Gohil, S. (2025). *Podcast Statistics: Global Listeners, Demographics, & More in 2025*. <https://meetanshi.com/blog/podcast-statistics/> [24.02.2026].
- Hasanabadi, M. R. (2023). An overview of text-to-speech systems and media applications. *arXiv*. DOI: 10.48550/arXiv.2310.14301.
- Hebbel-Seeger, A. (2009). Wissensdistribution und Wissenschaftsmarketing via PodCast. In C. Igel & A. Baca (Eds.), *Update eLearning* (pp. 101–110). Hamburg: Czwilina.
- Hebbel-Seeger, A. (2021). Technologien und Techniken in der (Online-)Lehre – Video, PodCast und Partizipation. In U. Dittler & C. Kreidl (Eds.), *Wie Corona die Hochschullehre verändert* (pp. 259–280). Berlin, Heidelberg: Springer VS. DOI: 10.1007/978-3-658-32609-8_17.
- Hebbel-Seeger, A. & Horky, T. (2026). From Knowledge Exchange to Brand Engagement: The Strategic Use of Podcasts in Sports Management by the example of the FISU World University Games Rhine-Ruhr 2025. *Athens Journal of Mass Media and Communications*, 12(2), 93–114.
- Hebbel-Seeger, A. & Strauß, A. (2022). *Konzeption von Lehre und didaktisch-methodische Qualifikation von Lehrenden unter der Maßgabe digitaler Transformation*. In U. Dittler (Ed.), *E-Learning* (pp. 101–117). Berlin: De Gruyter Oldenbourg. DOI: 10.1515/9783110754728-004.
- Hebbel-Seeger, A., & Vohle, F. (2026). „Echt jetzt?!“ Inwieweit verändert sich die Haltung zu und die Nutzung von Videos in Forschung und Lehre durch generative KI. In U. Dittler & C. Kreidl (Eds.), *Hochschulentwicklung der Zukunft. Lehre - Prüfungen - Qualitätsmanagement* (S. 221–241). Stuttgart: Schäffer-Poeschel.
- Hernandez-Lopez, M., & Mendoza-Jimenez, J. (2025). Podcasts Created by University Students: A Way to Improve Subject Understanding, Connection with Peers, and Academic Performance. *Education Sciences*, 15(3), 284. DOI: 10.3390/educsci15030284.
- Horton, D., & Wohl, R. (1956). Mass Communication and Para-Social Interaction: Observations on Intimacy at a Distance. *Psychiatry*, 19(3), 215–229.
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., & Liu, T. (2025). A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Transactions on Information Systems*, 43(2), Article 42, 1–55. DOI: 10.1145/3703155.
- Huschens, M., Briesch, M., Sobania, D., & Rothlauf, F. (2023). Do You Trust ChatGPT? - Perceived Credibility of Human and AI-Generated Content. *arXiv*. DOI: 10.48550/arXiv.2309.02524.
- Kakhki, S.K., Aghebati, N., & Moonaghi, H.K. (2025). Exploring the impact, challenges, and integration of podcasts in patient education: a systematic review. *BMC Medical Education*, 25(690), 1–11. DOI: 10.1186/s12909-025-07217-4.
- Kluckhohn, K. (2009). *Podcasts im Sprachunterricht am Beispiel Deutsch*. Berlin: Langenscheidt.
- Koning, B., & Voorveld, H. A. M. (2025). Disclaimer! This content is AI-generated: How AI-disclosures influence trust in advertisements and organizations. *Journal of Interactive Advertising*, 25(3), 240–253. DOI: 10.1080/15252019.2025.2554149.
- Lee, J.D., & See, K.A. (2004). Trust in automation: designing for appropriate reliance. *Human Factors*, 46(1), 50–80. DOI: 10.1518/hfes.46.1.50_3039.
- Li, Y., Deng, X., Hu, X., & Liu, J. (2024). The Effects of E-Commerce Recommendation System Transparency on Consumer Trust. *Journal of Theoretical and Applied Electronic Commerce Research*, 19(4), 2630–2649. DOI: 10.3390/jtaer19040126.

- Li, F., Yang, Y., & Yu, G. (2025). Nudging Perceived Credibility: The Impact of AIGC Labeling on User Distinction of AI-Generated Content. *Emerging Media*, 3(2), 275–304. DOI: 10.1177/27523543251317572.
- Liebermann, A., Schroeder, J., & Amir, O. (2022). A voice inside my head: The psychological and behavioral consequences of auditory technologies. *Organizational Behavior and Human Decision Processes*, 170, Article 104133. DOI: 10.1016/j.obhdp.2022.104133.
- Linares, D., Fox, N., & Berry, R. (2018). *Introduction: podcasting and podcasts - parameters of a new aural culture*. In D. Linares, N. Fox & R. Berry (Eds.), *Podcasting: New aural cultures and digital media* (pp. 1–13). Cham: Palgrave Macmillan. DOI: 10.1007/978-3-319-90056-8.
- Maloney Yorganci, K. T., & McMurtry, L. G. (2024). “One of us”: Examining the authenticity and parasocial relationships of stand-up comedian podcast hosts. *Journal of Radio & Audio Media*, 1–21. DOI: 10.1080/19376529.2024.2432023.
- Markman, K. M. (2012). Doing radio, making friends, and having fun: Exploring the motivations of independent audio podcasters. *New Media & Society*, 14(4), 547–565. DOI: 10.1177/1461444811420848.
- Mayer, F., Auschel, E., Baumbach, M., & Schindler, K. (2024). Corporate Communication Through Podcasts: The Impact of Voice-Fit and Attention During Listening on Parasocial Interaction with the Host and Persuasive Effects. *Journal of Radio & Audio Media*, 1–28. DOI: 10.1080/19376529.2024.2396840.
- McClung, S., & Johnson, K. (2010). Examining the Motives of Podcast Users. *Journal of Radio & Audio Media*, 17(1), 82–95. DOI: 10.1080/19376521003719391.
- McKenzie, L.E. (2019). Science podcasts: analysis of global production and output from 2004 to 2018. *Royal Society Open Science*, 6(1), 1–18. DOI: 10.1098/rsos.180932.
- Meden, E., Radovan, M., & Stefanc, D. (2024). Podcasts and Informal Learning: Exploring Knowledge Acquisition and Retention. *Education Sciences*, 14(10), 1129. DOI: 10.3390/educsci14101129.
- Mohale, N.E. (2024). Accessibility Challenges of Using Podcasts and Vodcasts in a South African Distance Learning University. *Research in Social Sciences and Technology*, 9(3), 93–109. DOI: 10.46303/ressat.2024.49.
- Nadora, M. (2019). *Parasocial Relationships with Podcast Host*. Portland State University. DOI: 10.15760/honors.789.
- Palenque, S.M. (2016). The Power of Podcasting: Perspectives on Pedagogy. *Journal of Instructional Research*, 5, 4–7. DOI: 10.9743/JIR.2016.1.
- Rime, J., Pike, C., & Collins, T. (2022). What is a podcast? Considering innovations in podcasting through the six-tensions framework. *The International Journal of Research into New Media Technologies*, 28(5), 1260–1282. DOI: 10.1177/13548565221104444.
- Sardar, S., & Vijay, T. S. (2025). Social Media Influencers and Purchase Intention: A Review and Research Agenda. *International Journal of Consumer Studies*, 49(3). DOI: 10.1111/ijcs.70046.
- Scantamburlo, T., Cortés, A., Foffano, F., Barrué, C., Distefano, V., Pham, L., & Fabris, A. (2023). Artificial Intelligence across Europe: A Study on Awareness, Attitude and Trust. *IEEE Transactions on Artificial Intelligence*, 6(2), 477–490. DOI: 10.1109/TAI.2024.3461633.
- Schäfer, M. S., Kremer, B., Mede, N. G., & Fischer, L. (2024). Trust in science, trust in ChatGPT? How Germans think about generative AI as a source in science communication. *Journal of Science Communication*, 23(9), A04. DOI: 10.22323/2.23090204.
- Schilke, O., & Reimann, M. (2025). The transparency dilemma: How AI disclosure erodes trust. *Organizational Behavior and Human Decision Processes*, 188, Article 104405. DOI: 10.1016/j.obhdp.2025.104405.

- Shao, A. (2025). Beyond Misinformation: A Conceptual Framework for Studying AI Hallucinations in (Science) Communication. *arXiv*. DOI: 2504.13777v1.
- Shaw, M.P., Fernandes, J.F.T., McGawley, K., Bell, L., & McNamara, S. (2025). “You’re the Gatekeeper”: Exploring Open-Access Podcast Creation in Sport and Exercise Sciences. *Journal of Radio & Audio Media*, 1–14. DOI: 10.1080/19376529.2025.2495349.
- Sigurgeirsson, A. T., & King, S. (2023). Controllable Speaking Styles Using a Large Language Model. *arXiv*. DOI: 10.48550/arXiv.2305.10321.
- Sung, E., Chung, W. Y., & Lee, D. (2023). Factors that affect consumer trust in product quality: a focus on online reviews and shopping platforms. *Humanities & Social Sciences Communication*, 10, Article 766. DOI: 10.1057/s41599-023-02277-7.
- Tolmeijer, S., Christen, M., Kandul, S., Kneer, M., & Bernstein, A. (2022). *Capable but Amoral? Comparing AI and Human Expert Collaboration in Ethical Decision Making*. CHI ’22: Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, Article 160, 1–17. DOI: 10.1145/3491102.3517732.
- Tolson, A. (2010). A new authenticity? Communicative practices on YouTube. *Critical Discourse Studies*, 7(4), 277–289. DOI: 10.1080/17405904.2010.511834.
- Vilceanu, M. O. (2025). Parasocial intimacy, change, and nostalgia in podcast listener reviews. *Media and Communication*, 13, Article 9059, 1–19. DOI: 10.17645/mac.9059.
- Wang, S. & Huang, G. (2024). The impact of machine authorship on news audience perceptions: A meta-analysis of experimental studies. *Communication Research*, 51(7), 815–842. DOI: 10.1177/00936502241229794.
- Wang, J., Shahzad, F., Ahmad, Z., Abdullah, M., & Hassan, N. M. (2022). Trust and Consumer’s Purchase Intention in a Social Commerce Platform: A Meta-Analytic Approach. *SAGE Open*, 12(2), 1–15. DOI: 10.1177/21582440221091262.
- Wang, S., & Székely, É. (2024). Evaluating Text-to-Speech Synthesis from a Large Discrete Token-based Speech Language Model. *arXiv*. DOI: 10.48550/arXiv.2405.09768.
- Xu, Z., Jain, S. & Kankanhalli, M. (2025). Hallucination is Inevitable: An Innate Limitation of Large Language Models. *arXiv*. DOI: 10.48550/arXiv.2401.11817.
- Zou, T., Shi, Z. (J.), & Wu, Y. (2026). EXPRESS: Welfare Implications of Democratization in Content Creation: Generative AI and Beyond. *Journal of Marketing Research*. DOI: 10.1177/00222437261423540