Supplementary Material B. Six PCA macro-codes and sixteen PCA items

B.1 Corpus adequacy and scale development

This Supplementary material documents the qualitative foundation of Perceived Cognitive Assistance (PCA) and shows how the qualitative evidence was translated into the Supplementary material B survey blocks. It draws on two source tiers: (a) the Tier B large language model (LLM) trading corpus (eight interviews and 44 YouTube narratives) and (b) a Tier A corpus of 24 qualitative and mixed-method studies on robo-advice, social trading, copy trading, and related advisory technologies. Using reflexive thematic analysis, we reconstruct four recurring experience themes in the 8+44 corpus and translate them into six cross-cutting macro-codes for item development (C1–C6). For conceptual clarity, C1–C5 define the PCA content domain (felt cognitive support), while C6 captures autonomy and over-reliance as an adjacent risk pathway that is tracked separately to keep PCA conceptually clean (Braun & Clarke, 2006; Malterud et al., 2016; Podsakoff et al., 2016).

Supplementary material C documents the Tier A qualitative corpus (24 qualitative and mixed-method studies, 2015–2025) that complements the Tier B 8+44 LLM trading corpus and maps each study to the six PCA macro-codes (C1–C6) and to the relevant survey blocks in Supplementary material B. Supplementary material E documents the measurement architecture by specifying the canonical item "content universes" for each comparator construct and the rationale for the retained Supplementary material B blocks (Table S18). Together, Supplementary materials B–C provide the qualitative grounding and triangulation for the PCA code frame, while Supplementary material E provides the canonical-versus-retained design logic for the full instrument.

Within this evidence base, Tier A combines robo-advice and broader digital-advisory work with social- and copy-trading research. Mixed-method and user-experience studies of robo-advisors document how trust, perceived diagnosticity, transparency, and control shape adoption and continued use, while also highlighting concerns about opacity, incomprehensible explanations, and low felt control in AI-empowered advisory systems—experiences that directly inform PCA codes C1–C5. Social- and copy-trading studies complement this picture by showing how following signal providers or copying strategies reconfigures responsibility attribution, overconfidence, and perceived skill, providing rich support for C1, C3 and especially C6.

Tier B adds a smaller but conceptually important cluster of LLM-related sources. These include the Bridging study on LLMs and retail investors, which combines TPB-guided interviews with YouTube material to show how LLMs alter perceived capability, reliance on human experts, and intentions to adopt complex strategies, and the Strategic Complexity and Behavioral Distortion article, which formalises PCA and articulates the five-step research agenda for empirical testing. A growing evidence base on LLMs in finance and an asset-management–focused report on trustworthy adoption extend this corpus to institutional settings, documenting perceived assistance, governance drift, and risk channels that generalise PCA-like experiences beyond retail trading. The separate "Qualitative Insights into User Experience with AI Financial Advisors and Trading Tools" synthesis (summarised in Supplementary material C) records the setting, key user-experience themes, assigned PCA code(s) C1–C6, and the corresponding Supplementary material B question family; Supplementary material E then situates these Supplementary material B blocks within the canonical-versus-retained scale architecture (Table S18) at the level of item blocks.

While the corpus did not contain many explicit references to "information compression" mechanisms such as summarisation dashboards, it showed saturated evidence for perceived relief from information overload—participants consistently described being able to "see everything in one place" and to process multiple sources without confusion. PCA3 captures this broader experiential theme rather than any technical compression process.

We began with four experiential themes of Tier B because they summarised the qualitative material at a narrative level. We did not retain the four-theme structure as the PCA content map because two themes reliably bundled two distinct mechanisms each, and item writing requires categories that are granular enough to support single-mechanism wording and explicit boundary control. Specifically, "overwhelm to a roadmap" combined workflow structuring and path support with decision-time navigation of uncertainty and volatility, and "analytic scaffolding and stepwise decision support" combined verification and error-checking with learning-oriented explanation. Keeping the four themes as the domain map would therefore have forced double-barrelled items or an implicit, undocumented splitting during item writing, both of which reduce transparency and weaken the content-validity argument. We also separated facilitative assistance content from the autonomy and over-reliance risk pathway, because mixing assistance and risk valence inside the same content category increases construct contamination. For this reason, the final code frame uses six macro-codes, where C1–C5 define the PCA content domain and C6 is retained only as an adjacent risk code for boundary transparency. In terms of translation from themes to macro-codes, "overwhelm to a roadmap" maps to C1 (structural and path support) and C4 (navigation of uncertainty and volatility), because these are empirically distinct mechanisms that require different item families and different boundary checks. "Analytic scaffolding and stepwise decision support" maps to C3 (error-checking and verification) and C5 (learning-oriented assistance), for the same reason. "Cognitive offloading and memory support" maps directly to C2 (cognitive-load relief), because the underlying mechanism is already narrow and scross sources. "Displacement of judgement and over-reliance" maps to C6 (autonomy and over-reliance), which is treated as a boundary and risk pathway rather than PCA scale content.

B.1.1. Saturation and "information power"

For the PCA construct, saturation is defined at the level of experiential types of perceived cognitive assistance: the distinct ways in which AI-enabled tools are experienced as helping or distorting investors' cognition, organised into the six PCA macro-codes C1–C6 (Malterud et al., 2016). The analytic question is whether further qualitative material is likely to reveal new kinds of assistance experience beyond C1–C6, or mainly additional examples and nuances of already-identified types (Hennink, Kaiser, & Marconi, 2017). To address this, we treat the Tier B LLM corpus (8 interviews and 44 YouTube sources on LLM-augmented trading) and the 24-study robo-advice, social-/copy-trading and adjacent algorithm-trust/AI-service corpus as a single qualitative evidence base focused on AI-augmented retail investing. New sources were added in waves until additional material no longer generated new PCA-relevant experiential types, but instead refined or re-situated experiences already captured under C1–C6 (for example, different flavours of "road-map" support in C1 or different forms of responsibility drift in C6). This follows an "information power" approach to saturation: given a narrow analytic aim (perceived cognitive assistance in AI-augmented retail investing), relatively homogeneous populations (retail and mass-affluent users of robo-advisors, social-/copy-trading platforms and LLM tools), and strong theoretical scaffolding (for example, perceived behavioural control, self-efficacy and trust in automation), a smaller

but information-rich corpus is sufficient once the typology of assistance experiences stabilises across tools and contexts (Malterud et al., 2016; Hennink et al., 2017).

Empirically, the threshold is reached when the code frame remains across technologies and contexts. In practice, all six macro-codes appear already in the early robo-advice and social-trading literature, and the 8+44 LLM corpus reproduces the same family of assistance experiences, now expressed in conversational or generative form. By around 18–20 of the 24 core robo-/social-trading studies, supplemented by the 8+44 LLM corpus, new work mainly reinforces or slightly repositions C1–C6 in adjacent settings rather than adding distinct macro-level codes. This pattern is consistent with methodological findings that, for focused qualitative aims with relatively homogeneous participants, most code-level themes emerge within roughly 9–17 interviews or cases, with deeper meaning-level saturation typically achieved in the high-teens to low-twenties (Hennink et al., 2017; Malterud et al., 2016). For scale development, the role of this saturation argument is to underpin content validity: it supports the claim that C1–C6 provide an adequate, non-arbitrary coverage of the domain of perceived cognitive assistance experiences that PCA is intended to measure, in line with standard guidance on construct domain specification, item representativeness and qualitative underpinnings of measurement instruments (Hinkin, 1995; DeVellis, 2016; Boateng et al., 2018; Morgado et al., 2017).

In consolidating the code frame, we aimed for the smallest number of non-redundant macro-codes that still captured the recurring ways in which AI-enabled tools assist or distort investors' thinking. Collapsing structural and path support with cognitive load reduction (C1 and C2) into a single efficiency-oriented code would have blurred a distinction that matters empirically and for item design: participants frequently describe clearer, better-structured workflows without a substantial reduction in felt cognitive effort, and conversely report reduced overload without strong perceptions of path guidance. Similarly, folding bias mitigation and discipline (C3) into general trust, or splitting transparency and explainability (C5) into several narrower subcodes, would either collapse conceptually distinct experiences (guardrail-like discipline versus confidence in the tool's logic) or create thin categories with limited coverage. Autonomy, control and over-reliance (C6) was retained as a separate macro-code, even though it sits on the distortion side of the framework and is not intended to be fully captured by PCA items, because concerns about handing over judgement recur across platforms and need to be tracked systematically. On this basis, six macro-codes represent the minimal partition that is both saturated in the qualitative corpus and aligned with the theoretical comparators used in the quantitative model (perceived behavioural control, perceived usefulness and ease of use, and trust in automation).

The coverage patterns in Supplementary material C and Supplementary material E indicate that this point has been reached for all six PCA codes. Structural and path support (C1) and cognitive load reduction (C2) are already prominent in the earliest robo-advice trials, where automated profiling, diversification and rebalancing provide novice investors with a clearly signposted "road-map" and reduce felt task complexity. Later UX-oriented robo studies, social-trading work and the 8+44 LLM corpus reproduce the same underlying types of assistance – stepwise workflows, default portfolios, and simplification of macro, news and signal information – even as interfaces, channels and specific tools differ. Bias mitigation and discipline (C3) similarly emerge early, as studies document both classic "guardrail" roles (counteracting under-diversification, preventing impulsive trading, enforcing contribution or rebalancing rules) and potential amplification of herding or overconfidence when signals are followed uncritically; subsequent studies add nuance – for example, whether rules are

perceived as supportive or constraining – rather than qualitatively new ways in which tools shape discipline. Experiences of time pressure and volatility navigation (C4) show an analogous trajectory: initial robo- and social-trading studies already highlight how automated alerts, scenario tools and stress dashboards help investors cope with turbulent markets, while also raising concerns that "calm" interfaces may underplay risk during crises. Newer LLM cases extend these patterns to conversational scenario exploration and explanation of shocks, but continue to fit within the same experiential type of assistance: navigating uncertainty and stress, rather than introducing an entirely new macro-code. Trust, transparency and explainability (C5) are the most ubiquitous: across almost all studies in Supplementary material C, investors oscillate between appreciation of diagnostic signals and frustration with opaque algorithms, generic advice and black-box rationales, and additional work in adjacent AI and algorithm-trust literatures reinforces these motifs without requiring extra top-level PCA codes. Finally, autonomy, control and over-reliance (C6) are saturated in the sense that we repeatedly observe the same family of experiences: investors using decision aids as supplements rather than replacements; pockets of over-delegation or "following the leader" in social trading; and LLM users who describe a pull towards outsourcing judgement while still recognising the need for monitoring and override.

Taken together, the coverage tables and narrative summaries suggest that by roughly the fifteenth to twentieth study, new material was primarily elaborating subcodes within C1–C6 and extending them to adjacent platforms rather than adding additional macro-level types of perceived assistance. This corresponds to what the methodological literature calls code-level or typological saturation: the point at which no new codes or code families emerge, even though further nuance and meaning can still be developed (Hennink et al., 2017). For the purposes of PCA scale development, we therefore treat the six macro-codes as theoretically saturated with respect to the domain they are intended to cover: the main recurring ways in which AI-enabled tools provide, or appear to provide, cognitive assistance in retail investing. Future qualitative work – especially in more autonomous LLM brokerage environments and in documented failure cases – is likely to sharpen the internal structure of C1–C6 and may motivate additional subdimensions, but is unlikely to necessitate new top-level PCA codes. This use of saturation is directly aligned with standard recommendations for construct domain specification and content validation in scale development, in which a sufficiently rich but conceptually bounded qualitative base is required before moving to item generation and quantitative refinement (Hinkin, 1995; DeVellis, 2016; Boateng et al., 2018; Morgado et al., 2017).

B.1.2. Coverage of the PCA code frame (C1–C6)

The coding tables and summaries in Supplementary material C show that each PCA subdimension is supported by multiple studies in at least two technological contexts (robo-advice and social trading, and for several codes also LLM tools). Experiences of clearer road-maps from goals to concrete steps and of having complex tasks broken into manageable stages underpin C1; experiences of reduced mental load and friction, including "having the system do the numbers" and "hassle-free" execution, alongside worries that decisions become "too easy" once interfaces are familiar, support C2; trust-mechanism work linking outcome quality, explanations and comparative performance to perceived reliability and continued use motivates C3; expectations that automated systems will prevent large errors under stress, coupled with uncertainty about behaviour in crises, inform C4; persistent concerns about opaque "black-box" logic, loss of control and the need for override options ground C5; and finally, social- and copy-trading research documenting how investors oscillate between self-

blame and blaming platforms or leaders for losses provides strong support for C6. Through the Qualitative Insights mapping, these codes are then linked to PCA, trust, perceived usefulness, and risk-comparator items in Supplementary material B, so that no PCA subdimension is left without qualitative backing; remaining issues concern wording refinements for context-sensitive areas such as effort/friction (C2) and responsibility/regret (C6), rather than gaps in the underlying evidence base.

Table S3 summarises this mapping at the level of question blocks: for each Supplementary material B block (B1–B11 and B9a) it records the latent construct or function, links to PCA code families C1–C6, the main theoretical anchors, and the analytical role of each block in the main model.

B.2. Structured qualitative synthesis of the 8+44 LLM corpus

B.2.1 Purpose and scope

This block reconstructs the experiential backbone of Gimmelberg et al. (2025) ("IJFBS") "8 interviews + 44 YouTube" study and translates it into PCA terms. The aim is to make explicit how the four themes that run through the published thematic analysis – overwhelm→roadmap, cognitive offloading, analytic scaffolding, and displacement of judgement / over-reliance – map onto the six PCA macro-codes C1–C6:

C1: structural and path support (from idea to execution)
C2: cognitive-load relief and information compression
C3: error-checking, discipline and bias mitigation
C4: navigation of complex or volatile conditions
C5: explanations, transparency and perceived diagnosticity
C6: autonomy, responsibility and over-reliance

All excerpts and summary pointers below derive from the "8+44 corpus" article and its underlying TPB mapping document for the eight interviews and 44 YouTube sources.

For each theme, the memo (a) gives a one-paragraph definition in PCA terms, (b) states the primary C-codes involved, and (c) lists 3–5 short illustrative excerpts from interviews (INT#) and videos (YT#).

B.2.2 Theme 1: Overwhelm → roadmap

PCA-oriented definition

Overwhelm→roadmap captures situations where retail investors describe LLMs as turning an unmanageable information environment into a clearer path from question to plan and, potentially, from plan to order. At the level of PCA, this theme embodies C1 and C2: LLMs are experienced as providing structural scaffolding for the decision process (clarifying steps, organising inputs, highlighting what matters first) and reducing subjective overload by summarising and translating complex material. In this mode, LLMs do not merely answer a question; they help the investor move from "too much, too messy" to a sequence of intelligible waypoints: understanding the situation, identifying options, and sketching an implementable strategy.

Across the "8+44 corpus", LLM use is repeatedly linked to lower access barriers: translating jargon into plain language and outlining plausible steps from question to plan. Overwhelm remains a background condition, but the perceived burden of structuring and path-finding is reduced when an LLM is available.

Primary PCA codes

Main: C1 structural/path support; C2 cognitive-load relief and information compression. Secondary: C4 when the "roadmap" explicitly concerns navigating volatile periods.

Illustrative excerpts (8+44 corpus)

– "LLMs can improve the investment decision process by optimising trading strategies, as one interviewee explained: 'LLM can critique existing strategy parameters and suggest improvements.'" (INT4)

– "LLMs can make high-quality information available to a broader audience, breaking down traditional barriers to accessing expert advice. They offer opportunity to educate RIs and translate complex financial jargon into simple terms." (YT25)

– "Experts are not accessible to most retail investors. Accessibility will lead to better understanding." (INT1)

– "Democratisation of knowledge and education by translating complex terminology into plain language." (INT summary)

– "Optimisation of trading strategies: an LLM can critique existing parameters and suggest improvements, supporting portfolio management." (INT summary)

Taken together, these excerpts describe LLMs as tools that help translate dispersed information and expert discourse into a structured plan and concrete next steps.

B.2.2 Theme 2: Cognitive offloading

PCA-oriented definition

Cognitive offloading refers to investors' experience of shifting memory- and computation-heavy parts of the task onto the LLM. In PCA language, this is primarily C2: the model compresses large volumes of information, keeps track of details, and performs laborious comparisons or scans that the human would find exhausting, while leaving the investor to interpret and decide. Offloading also has a C1 aspect: by handling background aggregation and pattern-spotting, the LLM frees the investor's limited attention to focus on higher-level structuring.

In the 8+44 corpus, this offloading is visible in how participants talk about LLMs "processing vast amounts of data", "reducing time for analysis and decision making", and "filling gaps in education and experience". What changes subjectively is not only the amount of information that can be considered, but the felt effort: with LLM support, heavy analysis becomes something that can be requested on demand, rather than personally shouldered.

Primary PCA codes

Main: C2 cognitive-load relief and information compression.
Secondary: C1 where offloading is tightly bound to structuring; C4 where offloading supports fast reactions under time pressure.

Illustrative excerpts

– "LLMs can process large volumes of data quickly and accurately." (YT14; YT5; YT11)

– "Efficient analysis of historical data, reducing time spent on analysis and decision making." (INT summary)

– "LLMs can improve the investment decision process by optimising trading strategies… [and provide] an unmatched level of customisation highlighted by an expert: 'All that AI is going to do is know more about you, learn more about you, look at patterns…'" (INT4; YT9)

– "An open-source LLM can be trained to understand complex financial concepts and market dynamics, supporting tasks such as summarising reports and identifying patterns." (YT25)

– "An LLM can help fill gaps in education and experience and may reduce emotional influence, which can increase perceived confidence and control." (INT summary)

These accounts show LLMs acting as a "cognitive buffer": they hold, compress and pre-process information so that the human does not have to, which is precisely the offloading facet of PCA.

B.2.3 Theme 3: Analytic scaffolding

PCA-oriented definition

Analytic scaffolding is the theme where LLMs are described as a second pair of eyes and a guided reasoning partner. In PCA terms, it combines C3 (error-checking and discipline), C4 (guided navigation of scenarios and volatile conditions), and C5 (explanations and interpretive support). Rather than simply dumping results, the LLM helps investors think: it critiques existing strategies, highlights inconsistencies, offers counter-arguments, and supplies explanations that make complex reasoning chains more transparent.

Within the "8+44 corpus" material, analytic scaffolding appears in references to "second opinions", bias reduction, emotional dampening, and the way LLMs are used to test and refine strategies. The LLM is not only an information source; it is an interactive scaffold that supports more disciplined and reflective decision making.

Primary PCA codes

Main: C3 error-checking and discipline; C4 scenario/volatility navigation; C5 explanations and perceived diagnosticity.

Illustrative excerpts

– “LLMs can play a tri-fold role: acting as advisors by themselves, aiding in understanding expert advice and aggregating and summarising a multitude of opinions.” (INT1 summary)

– “They can help with bias and contradictory information, offer a second opinion, reduce emotional influence, and improve understanding of large volumes of information.” (INT5; YT5; YT7; YT11)

– “LLM can critique existing strategy parameters and suggest improvements.” (INT4)

– “An LLM can offer a second opinion and an alternative view.” (INT summary)

– “LLMs have the potential to significantly enhance investors’ confidence by providing tools that simplify complex financial information. As one of the experts explained: ‘More confidence, more feeling of control – LLM will empower retail investors.’” (INT3)

These examples are exactly the behaviours PCA is meant to capture under analytic scaffolding: the LLM spots issues, forces the trader to articulate assumptions, and exposes them to alternative views, while also calming some of the affective turbulence around difficult decisions.

B.2.4 Theme 4: Displacement of judgement / over-reliance

PCA-oriented definition

Displacement of judgement captures the dark side of perceived assistance: the risk that, as LLM support becomes pervasive and persuasive, investors hand over more of their judgement than they intend. In PCA language this is C6, the autonomy and over-reliance code, with a strong link to the “illusion of understanding” and “illusion of control” noted in the Strategic Complexity paper. When investors say that the LLM “controls” them, or that increased understanding may come “with a higher risk of reliance on the model instead of using one’s own brain”, they are articulating a shift from assisted agency to partial displacement of agency.

The “8+44 corpus” makes this risk explicit: participants worry that uneducated retail investors may treat LLMs as new “speaking heads”, that confident but flawed outputs will be over-trusted, and that collective over-reliance could amplify herding and volatility if many users follow similar model-generated logic.

Primary PCA codes

Main: C6 autonomy, responsibility and over-reliance.
Secondary: C5 where opaque but persuasive outputs undermine calibrated trust.

Illustrative excerpts

– “Increased understanding will come together with a higher risk of reliance on the model instead of using [one’s] own brain capacity.” (INT3)

– “LLM is controlling me, due to my little understanding of the processes and factors – so, there is no extra ‘feel of control’ on side of retail investors.” (INT4, summary)

– “You are seeing certain false facts come out as answers from Chat GPT with a very confident sounding [tone]… This might create a false sense of control.” (YT23)

– “On the market I am familiar with the investors are mostly uneducated, so investors are heavily relying on speaking heads.” (INT4)

– “I could expect more volatility going forward. If these engines run with the same assumption and a huge crowd has access to them and follow them blindly, this could lead to a large buy or sell rush.” (INT summary)

Here, PCA ceases to be purely facilitative. The same properties that make LLMs feel helpful – fluency, confidence, speed – can, in the absence of understanding and guardrails, encourage investors to substitute model output for their own reflective judgement.

B.2.5 Synthesis: from four themes to six PCA codes

Across the 8 interviews and 44 videos, the four experiential themes align cleanly with the six PCA macro-codes:

– Overwhelm→roadmap: strong C1 and C2, with some C4 when the “roadmap” concerns turbulent conditions.
– Cognitive offloading: dominant C2, with C1 as a supporting function.
– Analytic scaffolding: combined C3, C4 and C5, as investors use LLMs to check, explore, and understand.
– Displacement of judgement / over-reliance: main locus of C6, with implications for how C5 (perceived transparency) can be misleading when outputs are fluent but fallible.

Crucially, all six codes are present in the “8+44 corpus” and are supported by both interview and YouTube material. The codes do not replace the original TPB framing (attitude, subjective norms, perceived behavioural control); rather, they cut across it, carving out the specific content domain that the PCA scale needs to measure. This serves as the bridge from the published thematic analysis to the PCA codebook: it shows how the lived experiences documented in the “8+44 corpus” study instantiate the structural, load-relief, scaffolding and autonomy facets that define Perceived Cognitive Assistance.

To document coverage and traceability in a compact way, Table S4 reports a light presence–absence coding of the six PCA macro-codes (C1–C6) across the full “8+44 corpus” , and Table S5 maps each candidate PCA item (PCA1–PCA16) to the macro-codes and to illustrative “8+44 corpus” evidence pointers.

To make the qualitative-to-measurement translation explicit, Table S3 maps each Supplementary material B question block to (a) its target construct or function, (b) the relevant PCA macro-codes (C1–C6), and (c) its analytical role in the main model. The table is included for traceability of instrument design; it does not replace the qualitative synthesis reported in Section 2.

Table S3. Translation from qualitative themes and PCA macro-codes (C1–C6) to Supplementary material B question blocks and analytical role in the main model.

Block (Supplementary material B)	Latent construct or function	Links to PCA macro-codes (C1–C6)	Main theoretical anchors	Evidence notes and analytical role in the main model	Block (Supplementary material B)	Latent construct or function	Links to PCA macro-codes (C1–C6)	Main theoretical anchors	Evidence notes and analytical role in the main model
B1. Attitudes toward using AI / LLMs in trading (Q1–Q4)	Theory of Planned Behaviour (TPB) Attitude toward using AI / LLMs for investing; evaluative stance (benefits vs risks, appeal vs unease).	Co-occurs with PCA themes C1–C3 (clarity, effort, diagnosticity) and C5 (control), but not scored as PCA. In the scale-development paper, B1 functions as a background attitude toward using AI / LLMs and is analysed as a comparator; the attitude term in the main model for complexity intentions is	TPB Attitude (Ajzen); TAM perceived usefulness; robo-advisor attitude work (ambivalence between help and opacity).	Qualitative robo, social-trading, and LLM interviews show ambivalent attitudes: appreciation for discipline / noise reduction plus concerns about loss of nuance and hallucinations. B1 captures this overall	B1. Attitudes toward using AI / LLMs in trading (Q1–Q4)	Theory of Planned Behaviour (TPB) Attitude toward using AI / LLMs for investing; evaluative stance (benefits vs risks, appeal vs unease).	Co-occurs with PCA themes C1–C3 (clarity, effort, diagnosticity) and C5 (control), but not scored as PCA. In the scale-development paper, B1 functions as a background attitude toward using AI / LLMs and is analysed as a comparator; the attitude term in the main model for complexity intentions is	TPB Attitude (Ajzen); TAM perceived usefulness; robo-advisor attitude work (ambivalence between help and opacity).	Qualitative robo, social-trading, and LLM interviews show ambivalent attitudes: appreciation for discipline / noise reduction plus concerns about loss of nuance and hallucinations. B1 captures this overall

		taken from B5.4 (attitude toward complex options tactics).		evaluative stance so that PCA can be modelled as a distinct capability belief, not a generic "AI is good" factor, and enters the main model as attitude.			taken from B5.4 (attitude toward complex options tactics).		evaluative stance so that PCA can be modelled as a distinct capability belief, not a generic "AI is good" factor, and enters the main model as attitude.
B2. Subjective norms about using AI / LLMs (Q5–Q8)	TPB Subjective Norms: perceived expectations and practices in the investor's reference groups (peers, experts, community).	Not a PCA construct; provides a background subjective-norms measure for using AI / LLMs. For the main model in this paper, the subjective norms term (norms regarding complex	TPB norms; social trading and herding literature (eToro, WSB); diffusion of innovation.	Corpus shows strong, saturated evidence of peer influence and "keeping up" with perceived practice (e.g., "everyone is using AI	B2. Subjective norms about using AI / LLMs (Q5–Q8)	TPB Subjective Norms: perceived expectations and practices in the investor's reference groups (peers, experts, community).	Not a PCA construct; provides a background subjective-norms measure for using AI / LLMs. For the main model in this paper, the subjective norms term (norms regarding complex	TPB norms; social trading and herding literature (eToro, WSB); diffusion of innovation.	Corpus shows strong, saturated evidence of peer influence and "keeping up" with perceived practice (e.g., "everyone is using AI

		tactics) is taken from B5.4 (SN_complex).		now"). B2 captures this normative climate so we can separate "I feel cognitively assisted" (PCA) from "I feel social pressure to adopt AI," and include subjective norms in the main model.			tactics) is taken from B5.4 (SN_complex).		now"). B2 captures this normative climate so we can separate "I feel cognitively assisted" (PCA) from "I feel social pressure to adopt AI," and include subjective norms in the main model.
B3. Baseline perceived behavioural control, pre-PCA (Q9–Q12)	Baseline, non-AI perceived control and self-efficacy for investing (knowledge, difficulty,	Intentionally non-PCA. Provides a baseline PBC measure against which PCA's incremental effect is judged (,). Conceptually	TPB PBC; self-efficacy theory (Bandura); trading self-efficacy and fintech self-	Qualitative data show many investors feeling under-equipped unaided (knowledge gaps, overload).	B3. Baseline perceived behavioural control, pre-PCA (Q9–Q12)	Baseline, non-AI perceived control and self-efficacy for investing (knowledge, difficulty,	Intentionally non-PCA. Provides a baseline PBC measure against which PCA's incremental effect is judged (,). Conceptually	TPB PBC; self-efficacy theory (Bandura); trading self-efficacy and fintech self-	Qualitative data show many investors feeling under-equipped unaided (knowledge gaps, overload).

	unaided capability).	contrasts “what I can do alone” vs “what I can do with LLM help.”	efficacy literature.	B3 captures that internal competence. In the main model it enters the controls vector (or perceived behavioural control) so that PCA effects reflect perceived added assistance, not generic confidence.		unaided capability).	contrasts “what I can do alone” vs “what I can do with LLM help.”	efficacy literature.	B3 captures that internal competence. In the main model it enters the controls vector (or perceived behavioural control) so that PCA effects reflect perceived added assistance, not generic confidence.
B4. PCA scale (PCA1–PCA16)	Perceived Cognitive Assistance: an intentionally over-specified 16-item	Direct operationalisation of PCA codes C1–C5 (capacity/clarity; effort/friction; diagnosticity;	TPB PBC extended to human–AI teaming; TAM decision aids; trust-		B4. PCA scale (PCA1–PCA16)	Perceived Cognitive Assistance: an intentionally over-specified 16-item	Direct operationalisation of PCA codes C1–C5 (capacity/clarity; effort/friction; diagnosticity;	TPB PBC extended to human–AI teaming; TAM decision aids; trust-	

	candidate pool capturing felt expansion of capability at decision time when an LLM is available (handling more information, structuring trades, error-checking, scenario navigation, and metacognitive learning). Codes C1–C5 (capacity / clarity, effort / friction, diagnostici	robustness/error-checking; control under complexity). C6 (responsibility/ regret) is handled outside the PCA item set.	in-automation and cognitive offloading. B4 is the only fully new construct built top-down from PCA codes C1–C6 and bottom-up from the qualitative corpus. PCA is the central mediator / moderator in the main model and the main mechanism input for the.			candidate pool capturing felt expansion of capability at decision time when an LLM is available (handling more information, structuring trades, error-checking, scenario navigation, and metacognitive learning). Codes C1–C5 (capacity / clarity, effort / friction, diagnostici	robustness/error-checking; control under complexity). C6 (responsibility/ regret) is handled outside the PCA item set.	in-automation and cognitive offloading. B4 is the only fully new construct built top-down from PCA codes C1–C6 and bottom-up from the qualitative corpus. PCA is the central mediator / moderator in the main model and the main mechanism input for the.	

	ty / discipline, robustness under volatility, control under complexity) define the content domain; C6 (responsibility / regret) is handled elsewhere. In the scale-development plan, all 16 items are used in Wave-1 EFA, with a shorter form expected for Wave-2 CFA and structural modelling.					ty / discipline, robustness under volatility, control under complexity) define the content domain; C6 (responsibility / regret) is handled elsewhere. In the scale-development plan, all 16 items are used in Wave-1 EFA, with a shorter form expected for Wave-2 CFA and structural modelling.			

B5.1 Perceived usefulness of LLM for trading (PU1–PU3)	Generic performance expectancy: belief that using LLMs improves trading outcomes and planning quality.	Neighbouring but distinct from PCA. Relates loosely to C1/C3 and is used as a convergent-validity comparator and for incremental-prediction checks.	Technology Acceptance Model perceived usefulness (Davis 1989); digital-finance adoption literature.	Robo-advisor and LLM studies repeatedly show "usefulness" language (better discipline, acceptable performance, convenience). B5.1 adapts this standard TAM block to LLM-augmented trading. In the main model, PU enters controls; PCA must outperform it to support.	B5.1 Perceived usefulness of LLM for trading (PU1–PU3)	Generic performance expectancy: belief that using LLMs improves trading outcomes and planning quality.	Neighbouring but distinct from PCA. Relates loosely to C1/C3 and is used as a convergent-validity comparator and for incremental-prediction checks.	Technology Acceptance Model perceived usefulness (Davis 1989); digital-finance adoption literature.	Robo-advisor and LLM studies repeatedly show "usefulness" language (better discipline, acceptable performance, convenience). B5.1 adapts this standard TAM block to LLM-augmented trading. In the main model, PU enters controls; PCA must outperform it to support.

B5.2 Perceived ease of use (PEOU1–PEOU3)	Perceived ease / difficulty of using LLM tools for trading tasks.	Overlaps conceptually with PCA's effort/friction code (C2), but focuses on interacting with the tool rather than felt cognitive capacity at decision time. Used as a convergent yet distinct comparator.	TAM perceived ease of use; usability and fintech-adoption work.	Qualitative corpus shows both "frictionless" experiences and onboarding / complexity complaints. B5.2 captures this system-use dimension so that PCA is not mistaken for "LLM is easy to use." In the main model it is part of the comparator set; PCA should add explanatory power	B5.2 Perceived ease of use (PEOU1–PEOU3)	Perceived ease / difficulty of using LLM tools for trading tasks.	Overlaps conceptually with PCA's effort/friction code (C2), but focuses on interacting with the tool rather than felt cognitive capacity at decision time. Used as a convergent yet distinct comparator.	TAM perceived ease of use; usability and fintech-adoption work.	Qualitative corpus shows both "frictionless" experiences and onboarding / complexity complaints. B5.2 captures this system-use dimension so that PCA is not mistaken for "LLM is easy to use." In the main model it is part of the comparator set; PCA should add explanatory power

				beyond PEOU.					beyond PEOU.
B5.3 Trust in LLM / automation (TR1–TR3)	Trust in the reliability, predictability, and benevolence of LLM-based assistance for trading.	Draws on experiences coded under C3–C5 (diagnosticity, robustness, transparency/control) and is conceptually related to responsibility (C6). Used as a comparator to demonstrate that PCA is distinct from trust and to probe whether mismatches between perceived assistance and trust relate to distortion risk.	Trust-in-automation scales (Jian, Lee & See); robo-advisor trust research; AI-ethics "trustworthy AI."	Many studies stress that reliance on automation depends on perceived reliability, explanations, and kill-switches. B5.3 imports short trust scales to sit alongside PCA; in validity tests PCA must not collapse into trust. In the main model trust is a control	B5.3 Trust in LLM / automation (TR1–TR3)	Trust in the reliability, predictability, and benevolence of LLM-based assistance for trading.	Draws on experiences coded under C3–C5 (diagnosticity, robustness, transparency/control) and is conceptually related to responsibility (C6). Used as a comparator to demonstrate that PCA is distinct from trust and to probe whether mismatches between perceived assistance and trust relate to distortion risk.	Trust-in-automation scales (Jian, Lee & See); robo-advisor trust research; AI-ethics "trustworthy AI."	Many studies stress that reliance on automation depends on perceived reliability, explanations, and kill-switches. B5.3 imports short trust scales to sit alongside PCA; in validity tests PCA must not collapse into trust. In the main model trust is a control

				and potential moderator for governance interpretations.					and potential moderator for governance interpretations.
B5.4 Attitudes and norms toward complex tactics (ATT_complex, SN_complex)	Evaluative stance and perceived social support for using more complex options strategies (multi-leg, volatility tactics).	Not PCA per se. Provides the attitude and subjective norms components specifically for complex strategy adoption, important for – where PCA predicts intentions conditional on TPB covariates.	TPB applied to adoption of complex financial products; strategic complexity framework.	Qualitative work shows mixed views and peer influence around multi-leg / options complexity. These items ensure the main model does not attribute to PCA what is actually driven by liking or social approval	B5.4 Attitudes and norms toward complex tactics (ATT_complex, SN_complex)	Evaluative stance and perceived social support for using more complex options strategies (multi-leg, volatility tactics).	Not PCA per se. Provides the attitude and subjective norms components specifically for complex strategy adoption, important for – where PCA predicts intentions conditional on TPB covariates.	TPB applied to adoption of complex financial products; strategic complexity framework.	Qualitative work shows mixed views and peer influence around multi-leg / options complexity. These items ensure the main model does not attribute to PCA what is actually driven by liking or social approval

				of complexity. They enter attitude and subjective norms in models of vignette-based complexity intentions and later outcomes.					of complexity. They enter attitude and subjective norms in models of vignette-based complexity intentions and later outcomes.
B5.5 Financial risk tolerance (RT1–RT3)	Dispositional willingness to bear financial risk in investing.	Orthogonal to PCA codes C1–C6 by design. Used as a control and robustness moderator for – (e.g., does PCA predict complexity beyond risk appetite?).	Grable–Lytton risk-tolerance scale and derivatives.	Risk-tolerance measures are well-validated; qualitative corpus confirms heterogeneity in comfort with risk. B5.5 enters controls so PCA	B5.5 Financial risk tolerance (RT1–RT3)	Dispositional willingness to bear financial risk in investing.	Orthogonal to PCA codes C1–C6 by design. Used as a control and robustness moderator for – (e.g., does PCA predict complexity beyond risk appetite?).	Grable–Lytton risk-tolerance scale and derivatives.	Risk-tolerance measures are well-validated; qualitative corpus confirms heterogeneity in comfort with risk. B5.5 enters controls so PCA

				effects are not confounded with simple risk-seeking. It also supports segmenting high-vs low-tolerance traders in analysis.					effects are not confounded with simple risk-seeking. It also supports segmenting high-vs low-tolerance traders in analysis.
B5.6 Trading self-efficacy (TSE1–TSE5)	Domain-specific self-efficacy for key options-trading tasks (mapping narratives to structures, reading IV, spotting assignment risk, etc.).	Conceptually adjacent to PCA's capability language but deliberately unaided ("what I can do myself"). Used as a discriminant comparator and for incremental-prediction checks.	Self-efficacy theory; domain-specific trading efficacy; fintech self-efficacy.	Interviews show some traders already feel highly capable, others rely heavily on tools. B5.6 quantifies that baseline domain expertise. In the main model,	B5.6 Trading self-efficacy (TSE1–TSE5)	Domain-specific self-efficacy for key options-trading tasks (mapping narratives to structures, reading IV, spotting assignment risk, etc.).	Conceptually adjacent to PCA's capability language but deliberately unaided ("what I can do myself"). Used as a discriminant comparator and for incremental-prediction checks.	Self-efficacy theory; domain-specific trading efficacy; fintech self-efficacy.	Interviews show some traders already feel highly capable, others rely heavily on tools. B5.6 quantifies that baseline domain expertise. In the main model,

				TSE is a comparator: PCA must predict intentions beyond TSE; their interaction is informative for distortion risk (high TSE + high PCA vs low TSE + high PCA).					TSE is a comparator: PCA must predict intentions beyond TSE; their interaction is informative for distortion risk (high TSE + high PCA vs low TSE + high PCA).
B6. Overconfidence calibration (OC1–OC6)	Trait-like calibration of probability judgments (general over- or under-confidence), scored via Brier error.	Connects conceptually to responsibility and regret (C6) but is intentionally measured outside PCA. Used as a discriminant comparator and as a candidate	Calibration and overconfidence literature (Brier; Lichtenstein et al.; Barber & Odean; Glaser & Weber).	Qualitative work highlights miscalibration and "illusion of understanding" under AI guidance. This block	B6. Overconfidence calibration (OC1–OC6)	Trait-like calibration of probability judgments (general over- or under-confidence), scored via Brier error.	Connects conceptually to responsibility and regret (C6) but is intentionally measured outside PCA. Used as a discriminant comparator and as a candidate	Calibration and overconfidence literature (Brier; Lichtenstein et al.; Barber & Odean; Glaser & Weber).	Qualitative work highlights miscalibration and "illusion of understanding" under AI guidance. This block

		moderator in robustness checks (PCA × overconfidence).		yields an overconfid ence measure as a standalone control, helping distinguish global miscalibrat ion from PCA-specific distortions. It is included as a control and used in interaction checks.			moderator in robustness checks (PCA × overconfidence).		yields an overconfid ence measure as a standalone control, helping distinguish global miscalibrat ion from PCA-specific distortions. It is included as a control and used in interaction checks.
B7. LLM-usage intensity for trading (IU1–IU5; LLM_idx)	Objective exposure to LLM tools for trading over the last week (days, minutes,	Operationalises LLM use in the main model. Not a PCA code, but necessary for –: tests whether PCA moderates the effect of	Technolog y-use intensity indices; digital-finance adoption metrics.	Corpus and prior work show variation in how intensively traders use LLMs. B7 quantifies	B7. LLM-usage intensity for trading (IU1–IU5; LLM_idx)	Objective exposure to LLM tools for trading over the last week (days, minutes,	Operationalises LLM use in the main model. Not a PCA code, but necessary for –: tests whether PCA moderates the effect of	Technolog y-use intensity indices; digital-finance adoption metrics.	Corpus and prior work show variation in how intensively traders use LLMs. B7 quantifies

	sessions, step-share), combined into an index LLM_idx.	LLM exposure and mediates LLM → complexity intentions.		that variation. LLM_idx_z enters the main model as LLM use and interacts with PCA (). It also serves as the exposure variable in mediation analyses (LLM → PCA → complexity intentions).		sessions, step-share), combined into an index LLM_idx.	LLM exposure and mediates LLM → complexity intentions.		that variation. LLM_idx_z enters the main model as LLM use and interacts with PCA (). It also serves as the exposure variable in mediation analyses (LLM → PCA → complexity intentions).
B8. Decision vignettes on tactic complexity (vignettes 1–2, C1–C3 coding)	Stated intention to choose higher-complexity tactics (defined-risk or more	Primary outcome block for PCA propositions –. PCA and LLM use predict the probability of choosing C2 / C3 vs C1,	Strategic structural complexity framework; design; vignette-based intention	Qualitative material shows that "moving up the complexity ladder" is the key behavioura	B8. Decision vignettes on tactic complexity (vignettes 1–2, C1–C3 coding)	Stated intention to choose higher-complexity tactics (defined-risk or more	Primary outcome block for PCA propositions –. PCA and LLM use predict the probability of choosing C2 / C3 vs C1,	Strategic structural complexity framework; design; vignette-based intention	Qualitative material shows that "moving up the complexity ladder" is the key behavioura

	advanced multi-leg) vs low-complexity or do-nothing options; provides Y_iv and I_iv outcomes.	controlling for attitudes, norms, risk tolerance, TSE, etc. Not a PCA code; this is where PCA "earns its keep."	methods in behavioural finance.	l risk / opportunity. B8 turns that into a measurable outcome (complexity choice and ratings). It is the dependent variable behaviour in the main model for the scale-development paper and a building block for later measures.		advanced multi-leg) vs low-complexity or do-nothing options; provides Y_iv and I_iv outcomes.	controlling for attitudes, norms, risk tolerance, TSE, etc. Not a PCA code; this is where PCA "earns its keep."	methods in behavioural finance.	l risk / opportunity. B8 turns that into a measurable outcome (complexity choice and ratings). It is the dependent variable behaviour in the main model for the scale-development paper and a building block for later measures.
B9. Experience and context (EX1–EX6; demographics)	Trading experience bands, options experience, platform, region,	Not PCA-coded; defines controls and grouping variables for invariance tests and moderation	Standard demographic and experience covariates in trading	Used to stratify samples (e.g., EX2 bands), run measurem	B9. Experience and context (EX1–EX6; demographics)	Trading experience bands, options experience, platform, region,	Not PCA-coded; defines controls and grouping variables for invariance tests and moderation	Standard demographic and experience covariates in trading	Used to stratify samples (e.g., EX2 bands), run measurem

	education, income.	(e.g., PCA effects across experience strata).	and HCI research.	ent invariance for PCA, and adjust regressions. Ensures PCA findings generalise across realistic trader segments and that PCA ≠ "just more experienced people."		education, income.	(e.g., PCA effects across experience strata).	and HCI research.	ent invariance for PCA, and adjust regressions. Ensures PCA findings generalise across realistic trader segments and that PCA ≠ "just more experienced people."
B9a. Wave-2 exploratory modules A–C and vignette V (motives, cost salience, mechanics, regime–strategy fit)	Exploratory moderators and segmentation variables: option-use motives, cost / friction salience, mechanics comprehen	Not PCA items; loosely related to PCA codes (e.g., cost salience vs C2; mechanics comprehension vs C3/C4) but treated as external moderators in the wider study, not part	Behavioural-finance work on motives, frictions, and knowledge; options-mechanics pedagogy; regime-strategy matching.	These modules support richer segmentation (e.g., whether PCA effects differ by mechanics knowledge or cost awareness)	B9a. Wave-2 exploratory modules A–C and vignette V (motives, cost salience, mechanics, regime–strategy fit)	Exploratory moderators and segmentation variables: option-use motives, cost / friction salience, mechanics comprehen	Not PCA items; loosely related to PCA codes (e.g., cost salience vs C2; mechanics comprehension vs C3/C4) but treated as external moderators in the wider study, not part	Behavioural-finance work on motives, frictions, and knowledge; options-mechanics pedagogy; regime-strategy matching.	These modules support richer segmentation (e.g., whether PCA effects differ by mechanics knowledge or cost awareness)

	sion, and regime–strategy recognition.	of the PCA construct.		and inform later behavioural-shift analyses. They are not included in the PCA measurement model, but can be interacted with PCA and LLM_idx in exploratory analyses.		sion, and regime–strategy recognition.	of the PCA construct.		and inform later behavioural-shift analyses. They are not included in the PCA measurement model, but can be interacted with PCA and LLM_idx in exploratory analyses.
B10–B11. Attention check and debrief	Data-quality gate (attention check) and ethical debrief with completion code.	No direct PCA linkage; indirectly supports all propositions by securing data quality.	Survey-methodology standards for panel work.	Attention checks and debrief do not enter the main model but protect the integrity of PCA measurement and outcome variables	B10–B11. Attention check and debrief	Data-quality gate (attention check) and ethical debrief with completion code.	No direct PCA linkage; indirectly supports all propositions by securing data quality.	Survey-methodology standards for panel work.	Attention checks and debrief do not enter the main model but protect the integrity of PCA measurement and outcome variables

				by flagging low-effort responses and ensuring ethical treatment of participants.					by flagging low-effort responses and ensuring ethical treatment of participants.

B.3. Coverage and item traceability checks in the 8+44 corpus.

Section 3 provides two compact checks on the qualitative backbone of the scale. Table S4 summarises how widely each PCA macro-code (C1–C6) appears across the 8 interviews and 44 YouTube narratives, and Table S5 provides an item-level traceability map linking PCA1–PCA16 to the code frame and to illustrative excerpts.

B.3.1. Coverage snapshot: presence–absence coding of PCA macro-codes (Table S4)

Table S4. Presence (0/1/2) of PCA macro-codes in the 8+44 corpus (interviews I1–I8; YT videos Y1–Y44) 0 = absent; 1 = weak / tangential; 2 = clearly present and central.

Case	**Source type**	**C1**	**C2**	**C3**	**C4**	**C5**	**C6**
I1	Interview	0	2	1	0	1	1
I2	Interview	0	1	2	0	1	0
I3	Interview	0	2	0	1	0	2
I4	Interview	1	2	2	1	2	2
I5	Interview	1	2	2	1	1	1
I6	Interview	1	2	1	0	2	0
I7	Interview	1	2	1	1	2	1
I8	Interview	2	2	1	1	1	1
Case	**Source type**	**C1**	**C2**	**C3**	**C4**	**C5**	**C6**
Y1	Video	2	2	1	0	1	0
Y2	Video	2	2	0	0	1	0
Y3	Video	2	2	1	1	0	0
Y4	Video	2	2	0	1	1	0
Y5	Video	2	2	1	0	0	0
Y6	Video	2	2	0	0	1	0
Y7	Video	2	2	1	0	1	1
Y8	Video	2	2	0	0	0	1
Y9	Video	2	2	1	0	1	0
Y10	Video	2	1	0	0	1	0
Y11	Video	2	1	1	0	0	0
Y12	Video	1	2	0	0	1	0
Y13	Video	1	2	1	0	0	0
Y14	Video	0	2	0	1	0	0
Y15	Video	1	2	0	1	0	0
Y16	Video	0	2	1	0	1	0
Y17	Video	1	2	0	0	0	1
Y18	Video	0	2	0	0	1	1
Y19	Video	1	2	1	0	1	0
Y20	Video	0	2	0	0	0	0
Y21	Video	1	2	1	1	0	0
Y22	Video	1	2	0	0	1	0
Y23	Video	1	1	2	1	2	0
Y24	Video	0	1	2	0	2	1
Y25	Video	1	0	2	1	1	0

Y26	Video	0	1	2	1	2	0
Y27	Video	1	1	2	0	1	1
Y28	Video	0	0	2	1	2	0
Y29	Video	1	1	2	1	1	0
Y30	Video	0	1	2	1	2	1
Y31	Video	1	0	2	0	1	0
Y32	Video	0	1	2	1	1	0
Y33	Video	1	1	2	1	2	1
Y34	Video	0	1	1	0	1	2
Y35	Video	1	1	0	0	1	2
Y36	Video	0	1	0	1	0	2
Y37	Video	0	0	1	0	1	2
Y38	Video	1	1	1	1	1	2
Y39	Video	0	0	0	1	1	2
Y40	Video	1	1	1	0	0	2
Y41	Video	0	1	0	1	1	2
Y42	Video	1	0	1	0	1	2
Y43	Video	0	1	0	0	1	2
Y44	Video	1	0	0	1	0	2

To make the link between the “8+44 corpus” thematic analysis and the six PCA macro-codes explicit, the eight interviews and 44 YouTube narratives reported in Gimmelberg et al. (2025) were re-read with a narrow focus on how LLM use helps or hinders investors’ thinking. Each case was rated on six content codes (C1 structural/path support; C2 cognitive-load relief and information compression; C3 error-checking and discipline; C4 navigation of complex or volatile conditions; C5 explanations and diagnosticity; C6 autonomy, responsibility and over-reliance) using a three-point scale: 0 = absent, 1 = mentioned briefly or tangentially, 2 = clearly present and central. For the interviews, ratings were based on the question-by-question summaries and quotations in the Theory of Planned Behaviour (TPB) mapping document used in the 8+44 corpus study. For the YouTube narratives, ratings were based on the question-level summaries used in the “8+44 corpus” analysis, because full verbatim transcripts were not consistently available. Table S4 provides a coverage snapshot showing that all six PCA codes are represented across the Tier B corpus, with C1–C3 most pervasive and C6 concentrated in cases that explicitly discuss over-reliance, displacement of judgement, or responsibility drift.

B.3.2. Item-level traceability: mapping PCA1–PCA16 to macro-codes and illustrative evidence pointers (Table S5)

Table S5. Traceability map for the Perceived Cognitive Assistance (PCA) candidate items (PCA1–PCA16): alignment to macro-codes (C1–C6) and illustrative evidence pointers to the Tier B “8 interviews + 44 YouTube” (8+44) corpus. PCA = Perceived Cognitive Assistance; LLM = large language model; INT# = interview source identifier (1–8); YT# = YouTube source identifier (1–44). C1 = structural and path support (from idea to execution); C2 = cognitive-load relief and information compression; C3 = error-checking, discipline and bias

mitigation; C4 = navigation of complex or volatile conditions; C5 = explanations, transparency and perceived diagnosticity; C6 = autonomy, responsibility and over-reliance.

Item	Statement	Primary macro-code(s)	Secondary macro-code(s)	Illustrative "8+44 corpus" evidence pointers			
PCA1	When I use an LLM during a trade decision, the path from idea to order becomes clearer.	C1	C2	INT4; YT25	INT4; YT25		
PCA2	The LLM helps me break a complex trade into concrete, manageable steps.	C1, C2	C5 (weak)	INT4; YT25	INT4; YT25		
PCA3	With an LLM, I can handle more information without feeling overloaded.	C1, C2	—	YT14; YT5; YT11; YT25	YT14; YT5; YT11; YT25		
PCA4	The LLM makes it easier to structure multi-leg or conditional orders when needed.	C1, C2	C5 (secondary)	INT4	INT4		
PCA5	Using an LLM, I spot inconsistencies and avoid simple mistakes I might otherwise miss.	C3	C4	INT5; INT4	C3	C4	INT5; INT4
PCA6	The LLM helps me compare alternative tactics side-by-side and	C1, C3	—	INT5	INT5		

	choose efficiently.				
PCA7	When markets move fast, the LLM helps me keep my bearings and act deliberately.	C1, C2	C5 (partial)		
PCA8	The LLM improves my ability to translate an overall view into a precise trade plan.	C1, C3	C5 (secondary)	INT4; YT25	INT4; YT25
PCA9	LLM support helps me notice when numbers, dates, or assumptions in my plan do not line up.	C3	C4	INT5	INT5
PCA10	LLM support helps me identify potential failure points in a trade before I commit capital.	C4	C3	INT4; INT5	INT4; INT5
PCA11	LLM support helps me compare alternative tactics for the same view side by side.	C1, C3	C5 (weak)	INT5	INT5
PCA12	LLM support helps me think through ‘what-if’ scenarios and plan for different market paths.	C4	C1		
PCA13	LLM support helps me keep my bearings	C1	C4, C5		

	and stay organised when markets move quickly.				
PCA14	LLM support helps me understand why a strategy fits my view, not just which strategy to use.	C1	C3, C5		
PCA15	With LLM support, I feel able to explore trading approaches I would not consider on my own.	C1	C5; C6 interface (risk tracked elsewhere)	INT1; YT25	INT1; YT25
PCA16	LLM support helps me reflect on and refine my own decision process across different trades.	C1	C3, C5	YT5	YT5

B.4 Tier B Corpus: YouTube Sampling Procedure

Channel sampling frame. Seventy-eight YouTube channels were selected using purposive sampling (Patton, 2015). The sampling frame was constructed in two layers. The first layer comprised major mainstream financial media channels (e.g., Bloomberg, CNBC, Yahoo Finance) included to ensure coverage of widely disseminated expert commentary. The second layer comprised smaller independent channels identified through platform recommendations and the research team's domain expertise in the retail investment media landscape. Each channel was required to regularly publish substantive content on financial markets, investment strategy, or financial technology. Channels focused exclusively on product promotion, course selling, or entertainment were excluded. The final set comprised 4 large channels (more than 1 million subscribers) and 74 smaller channels. Geographical distribution was concentrated in the United States (n = 67), with remaining channels based in Australia (n = 2), the UAE (n = 2), Singapore (n = 1), New Zealand (n = 1), the United Kingdom (n = 1), and four channels with unspecified location (Gimmelberg et al., 2025).

Transcript collection and dataset size. All publicly available transcripts from the 78 channels uploaded during Q2 2023 and Q2 2024 were collected, yielding 1,617 and 4,513 transcripts respectively (6,130 total). The combined dataset comprised approximately 20

million tokens (approximately 15 million words), with an average transcript length of 3,271 tokens (approximately 2,452 words).

Computational relevance filtering. Transcripts were processed through the AESTIMA tool using a Retrieval-Augmented Generation (RAG) pipeline (Gao et al., 2024; Yu, 2022). Textual content was converted into dense embeddings using OpenAI's text-embedding-ada-002 model and indexed via Pinecone.io cosine similarity search. Filtering proceeded in four stages: (i) identification of any LLM-related discussion (160 sources retained: 98 from Q2 2023, 62 from Q2 2024); (ii) narrowing to sources discussing LLMs in asset management (51 sources: 32 from Q2 2023, 19 from Q2 2024); (iii) application of 10 substantive questions aligned with the TPB interview schedule, extracted via ChatGPT-4o-mini and ChatGPT-4o (44 sources retained: 27 from Q2 2023, 17 from Q2 2024). The 10 extraction questions covered trust in LLM-generated advice, expert endorsement effects, LLM capability relative to human experts, replacement potential, rationality of LLM-informed decisions, adoption barriers, market dynamics impact, evolving professional roles, regulatory/ethical considerations, and feedback incorporation. Full question wordings are available from the corresponding author.

Validation. A random sample of 130 responses (25% of total extracted responses) was independently reviewed by three researchers against original source transcripts. Responses were classified as confirmed (clear match to source content), partially confirmed (meaning preserved but incomplete), or unconfirmed (no alignment with source). Results showed 76% full accuracy and 95% subject-level accuracy. Discrepancies were resolved through discussion until consensus (Gimmelberg et al., 2025).

Bias mitigation and limitations. Channel selection was purposive rather than probabilistic, which limits generalisability to the broader population of financial YouTube content (Malterud et al., 2016). However, three design features reduce the impact of selection bias on the final corpus. First, the computational filtering operated on all transcripts from selected channels, not on a manually curated subset, removing researcher judgment from transcript-level inclusion. Second, the filtering pipeline did not use view count, subscriber count, or engagement metrics at any stage, preventing popularity-driven selection bias. Third, the final 44-source corpus includes both favourable and critical perspectives on LLM use in investment, as reflected in the thematic results (e.g., over-reliance concerns, hallucination risks, scepticism about trust). Remaining limitations include English-language restriction, geographical concentration in US-based channels, and the absence of channels from Africa, South America, and parts of Asia, which may exclude regionally specific investor perspectives (Gimmelberg et al., 2025).