International Journal of Computing and Artificial Intelligence

E-ISSN: 2707-658X
P-ISSN: 2707-6571
Impact Factor (RJIF): 5.57
www.computersciencejournals.com/ijcai
IJCAI 2025; 6(2): 133-138
Received: 20-06-2025
Accepted: 25-07-2025

Ravi Palwe
Department, Information Technology (User Experience Design), USA

Corresponding Author:
Ravi Palwe
Department, Information Technology (User Experience Design), USA

Designing the Invisible: UX strategies for agentive AI in 2025 and beyond

Ravi Palwe

DOI: https://www.doi.org/10.33545/27076571.2025.v6.i2b.189

Abstract
The UX design should be ready to a future where the interface becomes less visible as AI systems move beyond the reactive assistant models to autonomous agent models. Financial planners, content schedulers and productivity agents are automated and act on behalf of users and are dependent on stated goals. Decisions under such systems are made independently and with little or no input of the user. The theme of this white paper is the use of UX specialists to create intuitive, trustful, ethically accountable interfaces of agentive AI. It discusses basics in delegation, calculation of trust, and feedback loops and ethical safeguards. It provides practical underpinning to creating a real experience in which users feel confident, educated, and in control, even when they are not in the loop, through practical examples of real-world experiences and changing patterns of design.
With the shift of AI systems after being reactive assistants to autonomous agents, the UX design will need to be prepared to handle a future where interfaces are more invisible. Autonomous financial planners, content schedulers, and smart productivity agents of the agentive AI systems act on behalf of users based on predetermined goals. These systems take autonomous decisions with little or no direct user control.
This white paper presents a discussion of how UX practitioners can create intuitive, trustworthy and morally responsible interface to agentive AI. It overviews the basic concepts in ethical protections, feedback loops, calibration of trust, and delegation. It provides a practical base on which to develop experiences in which users feel secure, informed and empowered, even without being in the loop, through real-life examples and changing patterns of design.

Keywords: Agentive AI, UX design, trust calibration, ethical safeguards

1. Introduction
Artificial intelligence that possesses agentive ability has transformed the interaction of the user with computers. In contrast to reactive AI systems, which have to be explicitly given command, proactive AIs systems continually act on behalf of their owners. Therefore, autonomy adds extra complexity to UX because conventional interface models must be substituted with systems that work almost silently. More industries are moving to agentic AI, whether in finance or in health and the creative industry. It demands that UX designers should break the paradigms of their past frameworks and work as part of multidisciplinary teams, including ethicists, AI engineers, and behavioral psychologists. This part seeks a brief overview of the traps and hopes in this new UX environment to precondition a deeper discussion of design principles on autonomous AI agents.

2. The Shift to Delegation-Based UX
The software tool interaction paradigm is radically being reorientated. Customers used to type in commands, keyboard information into input forms and press buttons to use systems. The agentive AI operates through delegation: the user enters macro-goals, and the AI calculates how to best act to fulfil the goals. This is not a purely technical change, but a change in conception- users will need to acquire new mental models, e.g. that using the system to get the correct decision is more appropriate, that system output is suggestions rather than orders, that it is all right to lose microanual control.

Facilitating the Agentive AI Shift through User Experience
The scaffolding of user understanding will have to go on purpose, the machine agency will have to be presented as collaborative, system behavior will be truthful with users,

and metaphorical design will have to correspond mental models to unlock the promise of agentive AI which is autonomous systems operating on behalf of people. This chapter outlines four ways of enabling that change, which are informed by recent human-centered design studies.

1. Structured Onboarding of Capabilities and Limitations: A properly written onboarding series establishes the right user expectations. The first laid down are the first-level introductions of the most important competencies of the AI, which are put aside into doable steps, and then there are specific disclaimers about cases when human help is required. Interactive walkthroughs, as an example, emphasize the occurrence of such failures - i.e. when the agent fails to interpret idioms - as it walks the user through didactic activities - i.e. automatic summarizing or template-based writing. To supplement the tutorials, capability matrices are provided which relate user goals (e.g. organize email, create report outline) with capabilities that are supported and the gaps that are documented, thus allowing the user to self-allocate tasks that are within the capacity of the agent. Lastly such teachings can be reinforced by just-in-time micro-notifications: Contextual tool instructions on initial use of a feature prompt consumers to always verify agent advice.

2. Writing Delegation as Cooperative Partnership

Blind faith in autonomous systems compromises system integrity as well as user agency. Interfaces must therefore set the artificial intelligence as a co-creator, not as an opaque actor in both language and the graphical display. Dual-pane workstations are built with inline annotations to justify every idea, using both user created material and AI generated ideas. Conversational voice, first-person and or otherwise described as I suggest making this paragraph more readable, etc., serves to assign agency appropriately and to create a bond. In addition, offering gradient interaction functions (like a "Suggest only" mode, in which the agent focuses on potential changes, or a "Draft Mode" mode, in which the agent creates entire paragraphs) gives users control over the intensity of involvement. Most importantly, All proposed improvements possess accept/ reject affordances, thus enhancing user control.

3. Early Communication to Control Anticipations

It can behave contrary to the expectations of the user when the autonomy of agentive artificial intelligence is greater than their own knowledge. Pre-action summaries (e.g., I will archive 15 newsletters tagged "Promo" - proceed?) must be used in systems, real-time progress displays must include confident estimates (e.g., 75 percent confidence this rephrase improves readability). This assists in reducing cognitive dissonance. Quickly follow polite failure messages (I'm unsure how to interpret this request--could you clarify?) accompanied with inline remedial options, when the agent encounters unresolvable uncertainty, or execution limits. Always accompanied by a clear opt-out, anticipatory cues can also recognize emergent user patterns ("You have already seen three summaries today--do you want me to automatically summary new documents in the future?).

4. Mental-Model Alignment: Metaphorical Design

Human beings superimpose new technologies on precedent patterns in order to make sense of them. Designers can quickly assist users to form the right mental model by selecting and frequently using a metaphor- partner, team member, or coach. By positioning the AI as a team member, such as in the form of teamwork graphics (avatar icons with team in their name, your teammate suggests..., workflow parallels (drag and drop task assignment)). Reflecting questions and progress badges, on the other hand, could be used by a coach. Above all, the chosen metaphor should be present throughout language, imagery, and patterns of interaction; a mix of metaphors (e.g., executing a program as robot, brainstorming as a hint as a human being) risks becoming confusive. Allow users to toggle metaphors-between coach and assistant mode - and retain that choice to more closely tailor the experience and model the support correctly.

5. Agentive AI Core UX Design Guidelines

The UX design of agentive artificial intelligence systems, which autonomously execute tasks to achieve user goals, should be carefully considered through the lens of autonomy, trust, and adaptability. Drawing from research in human-computer interaction (HCI), we outline six mutually dependent design principles to guide UX professionals in creating transparent, controlled, predictable, ethical, context-sensitive, and continuously evolving agentive experiences.

3.1 Open Independence: Autonomy that is not accompanied by visible reasoning can lead to user mistrust. Therefore, it is critical for agentive systems to make the rationale behind their actions clear to the user, promoting confidence and a sense of control. Autonomous AI actions must be accompanied by explanation layers so users can comprehend what's happening behind the scenes.

Describe reasonable paths of action

To foster transparency, provide lightweight "why this action?" affordances, such as hover tooltips or inline annotations, which clearly explain the data inputs, heuristics, or decision-making logic driving each action. These affordances should be contextually triggered and designed to help users easily understand the agent's reasoning process. For example, when an AI selects an action like a financial transaction, a small popup could briefly explain the rationale behind that decision, helping users see the logic in real-time.

Expanded Example: In a personal finance app, if an AI agent suggests reallocating an investment portfolio, it should include an explanation like, "Reducing tech stocks by 5% due to market volatility. This action is expected to reduce potential risk by 2% over the next quarter."

Visibility of States: It's essential to surface the ongoing processes within the agent's system through progress indicators or status updates. These visual cues (e.g., "Analyzing 4 of 12 files") help users track what the agent is currently doing and gauge how long it will take to complete the task. This could also include time-based indicators, such as "85% complete," giving users a real-time view of the

system's progress and enabling them to make informed decisions about whether to intervene or trust the system's completion timeline.
Expanded Example: For a content scheduling tool, users might see something like "Analyzing audience engagement metrics for the last 7 posts," helping them understand the current operation before scheduling new content.

3.2 Control and Adaptability
True collaboration between the user and the agent requires the user to maintain a level of personal control. As such, agentive systems must allow users to adjust the degree of autonomy granted to the system, ensuring that the user can always intervene when necessary and manage the system's actions effectively.

Dynamic Autonomy Spectrum
Provide users with the flexibility to adjust the level of autonomy through a sliding scale or preset options (e.g., "Review Only," "Suggest and Edit," or "Fully Automated"). This dynamic autonomy allows users to determine how much control they want the system to have at any given time, accommodating different user preferences and task requirements. For instance, a user may want the system to generate a report but manually edit its content, while in other instances, they might allow the system to fully automate the report creation.

Expanded Example
In a smart home system, users could choose from settings such as "Manual Control" (where they make all decisions), "Smart Assist" (where the system suggests changes but requires confirmation), or "Fully Automated" (where the system independently adjusts settings like temperature, lighting, etc.).

Flip and Revert Mechanisms
Empower users by integrating familiar controls, such as Stop, Undo, Revert, and Control buttons, allowing them to halt system actions or undo changes quickly. This ensures that if the system makes an unexpected or incorrect decision, the user can easily revert the process without losing data or progress. For example, if an AI system reorders a user's playlist incorrectly, the user should have an immediate, intuitive option to undo the change with a single click, providing assurance that they are always in control.

Expanded Example
In a financial planning app, if the AI recommends purchasing stocks based on certain criteria, users can click Undo if they disagree, and the system will instantly revert to the previous state, ensuring that no irreversible decisions are made without the user's full consent.

3.3 Contextual Adjustability: Failing to consider the user's current context can lead to information overload or insufficient assistance. A system that adapts to its user's situation will provide a more effective, less disruptive experience. UX must be dynamic, responsive to the situation at hand, and sensitive to the user's evolving needs.

Creating User Profiles and Expert Knowledge
A key aspect of contextual adjustability is personalizing the user experience based on individual preferences and expertise levels. For new users, provide easy-to-understand prompts, hints, and guided steps to help them navigate the system effectively. For experienced users, allow access to more advanced features and controls that can optimize the system's functionality for their needs. For example, a novice user might receive simple tooltips explaining features, while an expert might have the ability to customize those features or access more complex settings.

Expanded Example
In an online photo editor, a beginner might see tooltips guiding them through basic editing features like crop and color correction, while a professional user would be able to quickly access advanced tools like layer blending modes or RAW file adjustments.

Environmental Sensitivity
The system should recognize and adapt to the user's environment (e.g., mobile vs. desktop), device, workload, and current task. For example, on a mobile device, interactions might need to be more concise, with shorter alerts and fewer notifications, while a desktop version could accommodate full dashboards and detailed notifications. Additionally, during high-focus tasks, such as drafting a report, non-urgent recommendations should be delayed or minimized to avoid distractions. This level of adaptability ensures that the system remains helpful and unobtrusive in any context.

Expanded Example
On a smartphone, a personal assistant AI might only give the user essential updates like reminders or calendar events, whereas on a desktop, it might provide more detailed suggestions or handle more complex tasks, such as managing emails or project timelines.

3.4 Continuous Commentary and Learning
In order to maintain user trust and ensure that the system is consistently improving, a **continuous feedback loop** between the user and the agent must be established. This loop not only helps users stay informed about system performance but also fosters a sense of partnership and ongoing collaboration.

Systems of In-Context Feedback
Integrating micro-surveys or quick feedback requests directly within the system's interface allows users to provide context-specific feedback without disrupting their workflow. For instance, after an AI suggests an edit or decision, users could be asked, "Was this suggestion helpful?" This real-time feedback allows the system to learn from user interactions and adjust its performance accordingly. The system could then use this data to refine its algorithms or adapt to the user's preferences over time.

Expanded Example
After an AI tool in a design software suggests color palette changes, the user could be prompted with "How helpful was this suggestion?" with the option to rate it. The system could then use this rating to refine future suggestions for the user.

Performance Dashboards
Incorporating performance dashboards that display metrics like time saved, error rates, or the ratio of accepted

recommendations helps users evaluate how well the system is performing and provides transparency into its effectiveness. These dashboards serve as both a tool for users to assess the AI's impact and a means of building trust. For example, a user could see a summary of how much time the system has saved them over the past month, or how accurately it's predicted certain outcomes. By offering real-time feedback on system performance, users feel more confident in the system's capabilities and its ability to meet their goals.

Expanded Example: In a productivity tool, a dashboard might show that the AI has helped save the user 10 hours per week by automatically organizing tasks, along with a satisfaction metric showing how often the user accepts the AI's suggestions.

4. Case Studies & Design Explorations

Case study 01: Agentive Financial Adviser

In periods of high market volatility, investors require both nimbleness and transparency. Conventional robo-advisors might make clients feel out of place and not understand the reasoning behind the choices they have made even with their automatic portfolio adjustments. This case analysis examines the design and structure of an artificial intelligence application known as Agentive Financial Advisor- which autonomously reallocated funds without harming consumers who remained informed, in control, and empowered.

Objective Statement for Problems: Lack of Transparency: Users reject automatic deals since they cannot understand their justification.

- Once rebalancing runs, it is difficult to undo or change behavior.
- Dense performance reports cause cognitive overload for non-experienced investors.

Research & Persona

Who: Considering Samantha Lee our persona she is 42, Marketing Manager Investment Experience: Moderate (retirement portfolio); Tech Comfort: Regular user of consumer financial apps, unfamiliar with professional trading tools

Why: Samantha wants her portfolio to react quickly to market volatility, but she also needs clear explanations and the power to change or undo choices to feel comfortable.

What: An artificial intelligence-driven adviser who lets Samantha specify

1. Risk tolerance and target areas.
2. Shows instantaneous reasoning and confidence.
3. Creates visual comments and narrative reports.
4. Provides quick-undo together with a full audit trail.

Design Goals

1. **Transparency:** Reveal "what" the AI did and "why."
2. **Control:** Enable easy adjustments and reversions.
3. **Clarity:** Present data and reports in everyday language with visual cues.
4. **Trust:** Build confidence through continuous feedback loops.

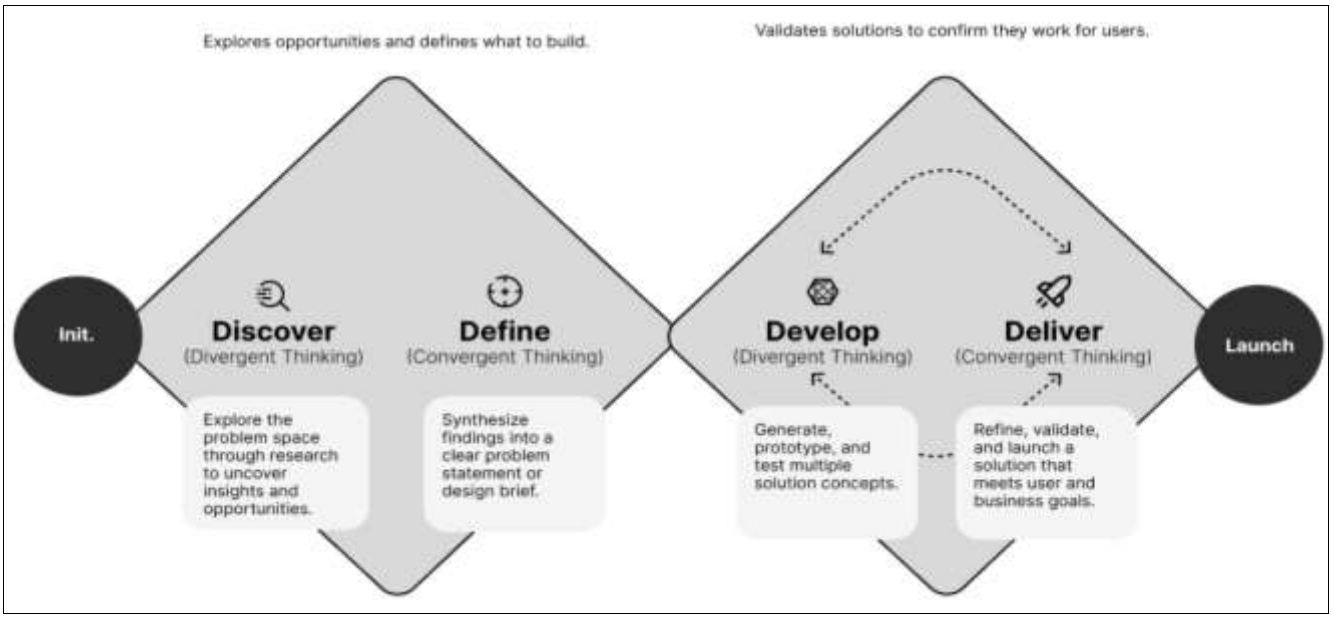

Fig 1: Double diamond UX design process

Design Process Map 1.0

1. Design Process: Ideation & Sketching

Rapid sketches of a dual-slider goal panel, trust dashboard, and audit feed.

2. Wireframes & Prototypes

- Low-fidelity wireframes to validate flows for setting goals, viewing live metrics, and undoing trades.
- Mid-fidelity prototypes in Figma to test dashboard layouts and narrative report styles.

3. Usability Testing

- Five sessions with investors like Samantha.
- Key feedback: need for concise "why" bullets, a clear undo affordance, and animated visuals for before/after allocations.

4. Iteration
- Refined the confidence score display into a simple gauge.
- Simplified the narrative report into a three-section digest ("What Happened," "Why It Happened," "Next Steps").
- Added a prominent, time-bound "Undo" button next to each trade entry.

Solution

1. Interface for goal definition
- **Dual Sliders:** "Time Horizon" (Short ↔ Long) and "Risk Tolerance," Conservative - Aggressive.
- **One-click simulations** for "Recession, Bull Market, and Inflation Surge."
- **Live preview** graph instantly sees expected volatility and returns as sliders move.

Knowing the AI will stay inside her selected limits, Samantha boldly sets her comfort level.

2. Real-Time Dashboard of Trust
- Based on model confidence and backrest data, confidence falls between 0 and 100%. "Best,," "Base,," and "Worst" case tiles refresh dynamically with streaming market data.
- Reasoning Shortcuts: In line bullets like "Reducing Tech by 5%—VIX up 12% in 2 days."

Impact: Samantha sees choices come to pass and is not surprised by unexpected reallocations.

3. Three-section digest for the post-action narrative report
- What happened that "$10k moved from Fund A to B on May 14"?
- Why It Happened "Tech volatility crossed your 8% threshold" ?
- What's Next? "Expect ~2.3% uplift next quarter if stable."
- Animated pie charts show changes both before and after.

Impact: Samantha understands right away the results of AI actions without reading spreadsheets.

4. Quick-undo & Audit Trail
- Every trade shows data inputs, model version, and time from a timed transaction feed.
- Reversing individual or batch trades within 24 hours will instantly restore the portfolio.

Impact

Samantha feels always in control; low-confidence or mistake may be corrected in a few seconds.

User Journey

Table 1: User journey, touchpoints, pain points, improvements, confidence

Stage	User Action	System Touchpoint	User Thought / Feeling	Pain Point	Opportunity
1. Morning Adjustment	Logs in; tweaks risk slider from 6 → 5	Risk slider control; "Save" button	"I want to be more conservative today."	Slider default too coarse; no quick preset for conservative	Add one-click presets (Conservative/Moderate/Aggressive) next to slider
2. Scenario Testing	Loads "Inflation Surge" scenario; reviews 5% bond reduction simulation	Scenario dropdown; results panel	"What happens if bonds drop even more?"	Results cluttered; hard to see impact on overall allocation	Show scenario impact overlay directly on portfolio chart
3. Go Live	Hits "Rebalance"; dashboard flashes 85% confidence and updated projections	"Rebalance" button; confidence badge; projection graphs	"Looks good high confidence."	Confidence metric ambiguous; no explanation of calculation	Tooltip or link to "How confidence is computed"
4. Coffee Review	Reads narrative report; watches animated pie-chart transitions	Narrative text panel; animated pie charts	"Nice to have the story behind the numbers."	Animation too slow; narrative too verbose	Allow pausing animation; offer "summary" vs. "detailed" report toggle
5. Evening Audit	Spots a 60% confidence trade; taps "Undo" and confirms rollback	Low-confidence highlight; "Undo" prompt	"Good catch I want to reverse that."	Undo confirmation feels modal-heavy	Inline "undo" with minimal confirmation (e.g., toast action)

Results & Learnings
1. **Increased Engagement:** As users investigated scenarios, daily log-ins climbed by forty percent.
2. **Reduced Anxiety:** Undo use accounted for just 5% of trades, but it raised general satisfaction levels by 25%.
3. **Inspired Cooperation:** Users said they felt more like active participants than like docile onlookers.

Main Learning: Combining transparent thinking, easy controls, and strong security turns automated investing into a trustworthy, team player experience. Agentive artificial intelligence may provide both speed and peace of mind by giving user agency first priority at every touchpoint**.**

Context
1. Reacting to market volatility, artificial intelligence system independently re-allocated funds.
2. Slides and scenario templates help to define goals in the user experience.

3. A real-time confidence level, anticipated outcome, and rationale trust dashboard
4. Reports on post-action analysis covering every asset change, market trigger, and performance criteria.
5. Every transaction has an audit trail; there is a quick-undo option with reversal capability.

Wireframes Concepts

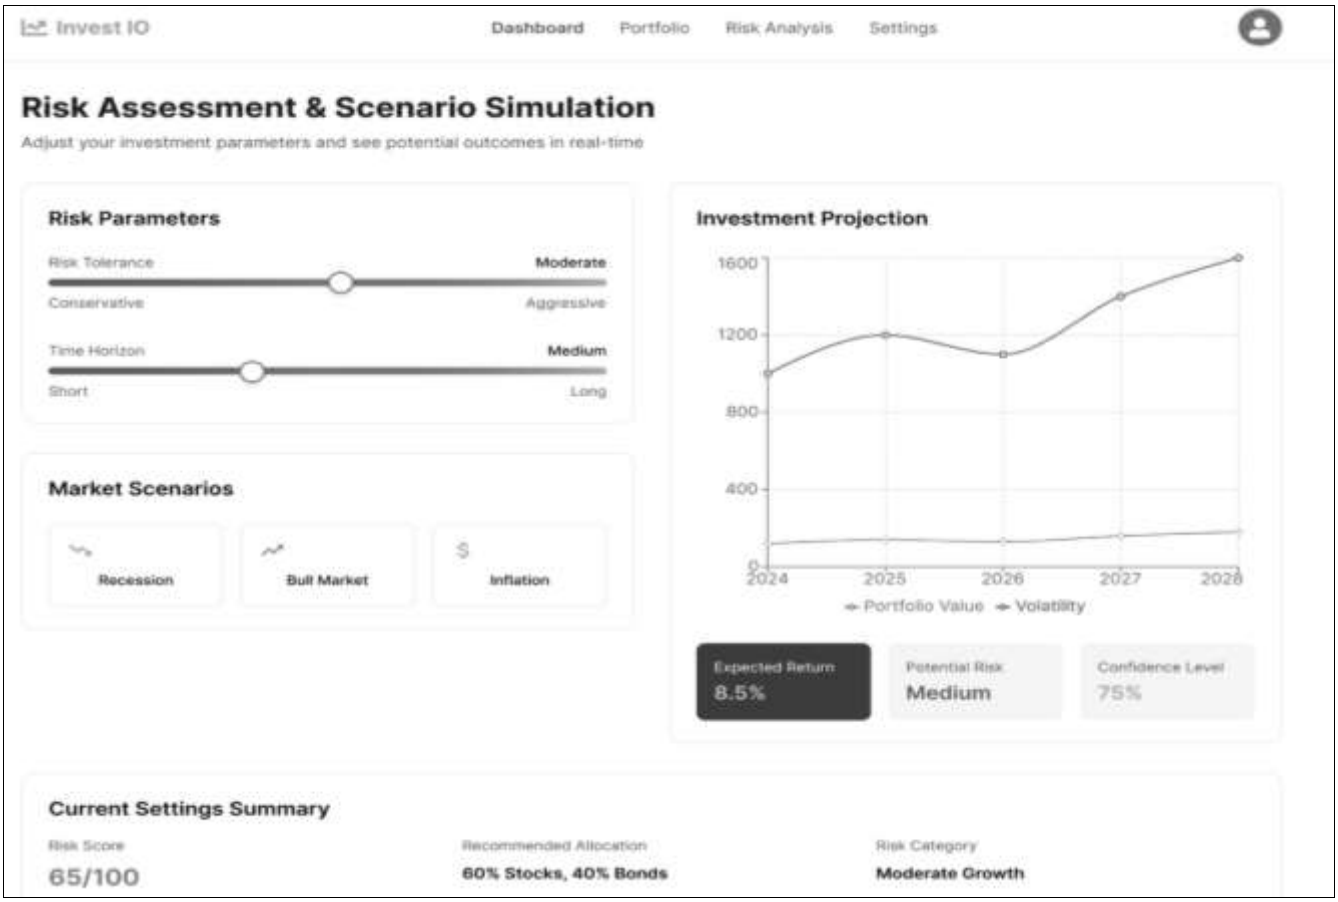

Fig 2: Risk assessment and scenario simulation

Conclusion

Agentive artificial intelligence marks a paradigm change in the field of user experience. Designers today have to create systems that, although functioning behind the scenes, remain transparent, trustworthy, and human-centered. Every action has to be reversible; autonomy has to be justified and delegation has to be earned.

Success of agentive artificial intelligence depends on users' confidence as much as on technological competence. UX designers may contribute to build an ethical, efficient future for artificial intelligence that greatly benefits people even if it runs without direction by giving trust, explainability, and control top priority.

UX for agentive artificial intelligence is about making the unseen visible but also about making it consistent with human ideals, compassionate, and accountable.

References

1. Norman DA. Designing the agentive technology. Interactions. 2017;24(3):12-19.
2. Sundar SS. Human-computer interaction: Interacting with intelligent machines. Annu Rev Psychol. 2020;71:255-278.
3. Google AI Blog. How Google Assistant anticipates your needs [Internet]. 2023 [cited 2025 Sep 17]. Available from: https://ai.googleblog.com/
4. Tesla. Tesla Autopilot user experience [Internet]. 2023 [cited 2025 Sep 17]. Available from: https://www.tesla.com/autopilot
5. Apple Inc. Apple privacy overview [Internet]. 2023 [cited 2025 Sep 17]. Available from: https://www.apple.com/privacy/
6. Spotify. Spotify's personalization and UX design [Internet]. 2023 [cited 2025 Sep 17]. Available from: https://newsroom.spotify.com/personalization
7. Replika. Replika emotional AI design principles [Internet]. 2023 [cited 2025 Sep 17]. Available from: https://replika.ai/blog