EY Studio+

Shape the future with confidence

From AI to Agentic Experience:

The Pragmatic Logic Behind Agentic AI Acceptance

Human Signals

Dutch Edition

May 2026

EY Studio+, The Netherlands

The better the question. The better the answer.
The better the world works.

Contents

EY Studio+

Powering transformation with insight and empathy

EY Studio+ creates transformative experiences that move people and shape markets. Our solutions work at the intersection of creative design and commercial insight, powered by cutting-edge technology, to drive real business value in a captivating way. We reimagine customer centricity and shape new possibilities at scale, drawing on EY's full spectrum of services, sector insights, and diverse ecosystem partners.

Welcome to the latest edition of Human Signals

To begin, a provocative question: would you trust AI to make the most important decisions in your life?

Most people would have said 'no'. In moments that carry personal or financial weight, the natural inclination is to seek human expertise: someone to listen, to provide perspective, and to confirm that an important decision has been considered properly. Yet artificial intelligence is becoming an increasingly visible part of the services that govern how individuals manage their money, their plans and their decisions.

This edition of Human Signals takes that tension into the Dutch market. The original UK research introduced the concept of Empathy Demand: the desire to feel emotionally recognized, particularly in moments where money intersects with significant life events. The question explored here is whether the same logic governs acceptance in a different cultural context, or whether Dutch consumers draw the line elsewhere.

Preface

From AI to AX in the Dutch financial landscape

Dutch consumers prefer human interaction in principle, but use AI in practice when it is useful, controllable and clearly bounded. This tension, between stated preference and actual behavior, runs through the findings of this report and shapes how AI is accepted in customer facing financial services.

This research does not argue for faster or broader AI deployment. It argues for more precise deployment. In this report, that shift is described as the move from AI to AX: from deploying artificial intelligence for automation, to deliberately designing augmented experiences where AI supports human led service rather than replaces it. Within this frame, "agentic services" refer to AI systems that move beyond explanation or advice and are authorized to act, decide or execute on the user's behalf. This logic translates into a small set of concrete implications that should guide decisions about where and how (agentic) AI is introduced.

AI should be deployed first in low stakes, high friction moments because this is where trust is earned. When AI consistently delivers correctness, speed and convenience in supportive roles, it establishes credibility without risk. These early successes create the conditions under which acceptance of more agentic behavior can grow.

As users experience tangible value in bounded AI powered services, their willingness to extend AI's role increases. Autonomy becomes acceptable not by design principle, but by accumulated proof. Introducing agentic AI too early, in irreversible heavy moments, breaks this progression. In such contexts, autonomy is experienced as overreach rather than support and trust. Once this is lost, it is hard to rebuild. In these contexts, leading with automation is interpreted as overreach, regardless of technical capability. Trust cannot be accelerated by deploying AI earlier than it is culturally and emotionally acceptable.

Automation creates the most value when it is applied to repeatable, bounded and document heavy tasks that sit around the decision rather than replacing it. Intake, triage, summarization, document preparation and administrative follow up reduce effort while preserving user agency. In these tasks, automation is not threatening; it is relieving.

Human involvement should remain central where judgment, exceptions, escalation and final decisions are required. The research shows that resistance to AI is rarely about competence and consistently about the loss of an intervention moment. Where discretion, interpretation and responsibility matter, users expect a human presence that AI supports rather than replaces.

Finally, the hand off between AI and humans must be designed as a first order part of the experience. Clear escalation paths, visible user control and explicit accountability shape whether AI is perceived as assistance or deflection. Trust is not granted upfront; it is built when users can see where AI stops, who takes over and how responsibility is handled if something goes wrong.

Read through this lens, the report is not a case for automating more of the front office. It is a case for making sharper choices about role, timing and boundaries.

The organizations that succeed will not be those that deploy AI most aggressively, but those that deploy it exactly where it is welcomed, and deliberately restrain it where it is not.

Section 1

Designing for moments that matter

Studying how consumers might engage with AI in high-stakes moments is a difficult research problem. Asking participants to evaluate a technology they have not yet experienced in a complex personal context tends to produce abstract responses, often disconnected from how they would actually behave. This section describes how the research was designed to overcome that limitation.

A mixed-methods approach

The study combines qualitative depth with quantitative validation. Qualitative interviews carry the weight of interpretation, and surface the motivations, hesitations and contradictions that shape how users perceive AI. The survey, fielded to a representative Dutch sample of more than 2,000 respondents, tests whether the patterns observed in the interviews hold at population scale. Read together, the two waves distinguish strongly held minority views from broadly shared positions, and preserve the human texture of the research without losing statistical traction.

Grounding insights in lived experience

The qualitative component consists of 32 in-depth interviews with Dutch consumers who had recently experienced financially complex situations: moments characterized by uncertainty, stress or significant personal impact. These are precisely the situations in which the value, and the risks of AI become most visible. They are also the situations in which abstract opinions tend to give way to specific ones, anchored in what participants had actually felt and needed.

A semi-structured interview format was used, combining a consistent framework with the flexibility to follow individual experiences in greater depth. This balance enables comparability across participants while leaving room for unexpected insights to emerge.

So, we built an AI prototype

When AI is discussed only in abstract terms, responses remain hypothetical. To shift the conversation onto more concrete ground, EY Studio+ developed a visual prototype: an AI-supported financial assistant named "***Robin***". The prototype was designed not as a product proposal but as a research instrument, providing participants with a concrete experience to react to, rater than a theoretical concept.

Robin was deliberately built to operate at varying levels of autonomy: from explaining and clarifying, to advising, to acting on the user's behalf. Making these distinctions explicit, proved to be one of the most useful design choices in the study. Concerns that participants struggled to articulate in general terms became immediately accessible once the question shifted from wheter AI was acceptable, to wheter it was acceptable in a specific role.

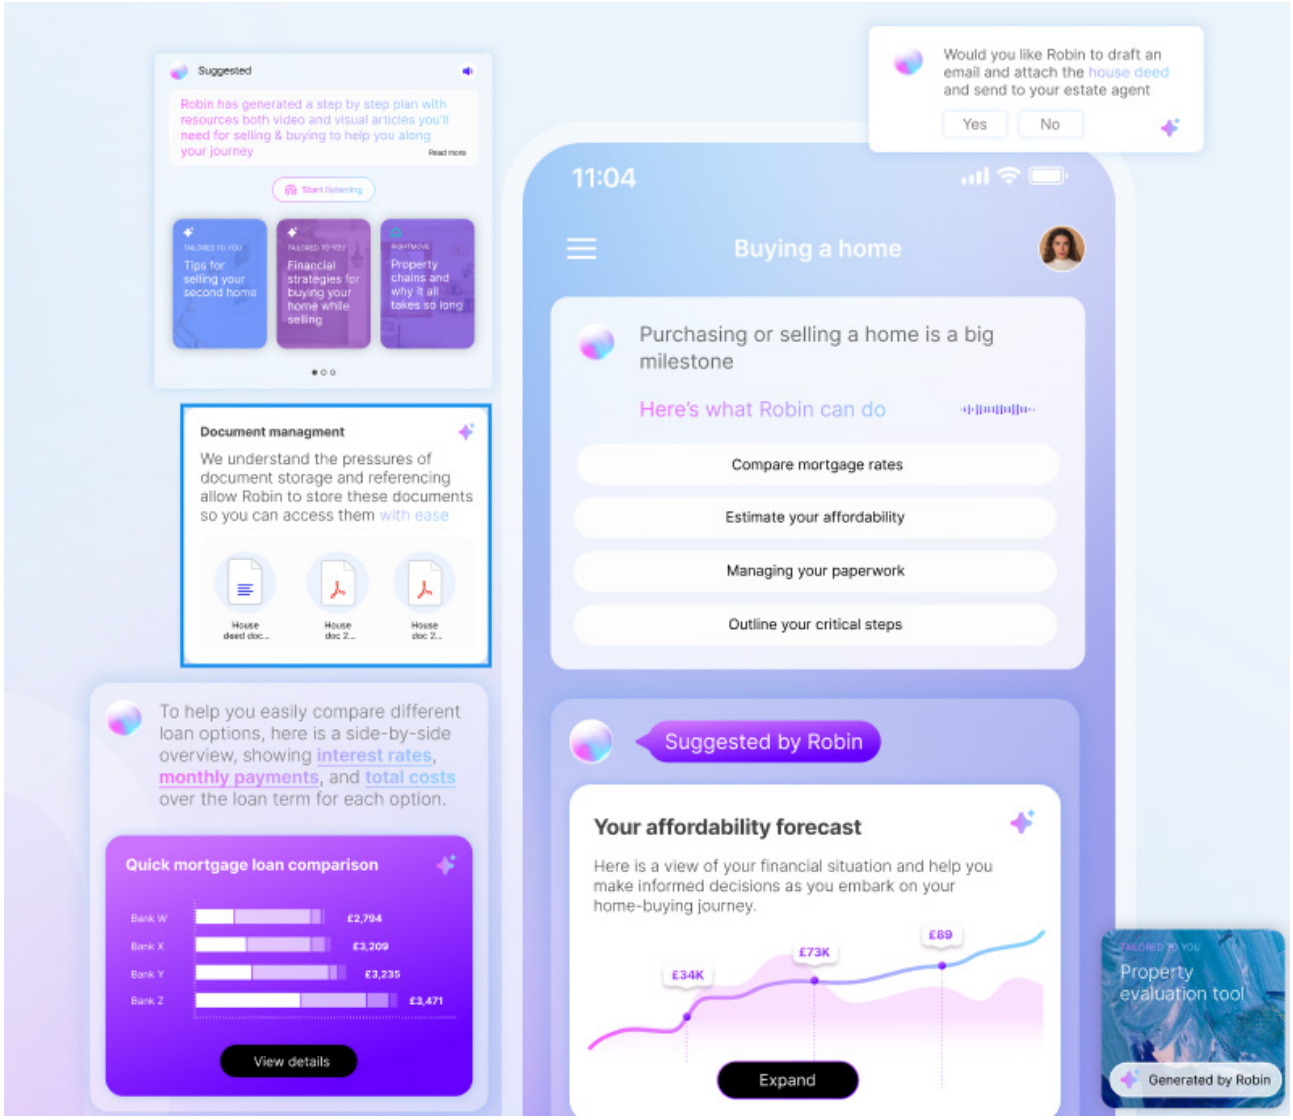

Robin prototype screens illustrating AI supported guidance in a home buying scenario

A scenario-based interaction

Each interview opened with a financial situation the participant had personally experienced: financial difficulty, buying a home, bereavement, retirement planning or complaint handling. Robin was then introduced as a potential support tool within that scenario. The sequencing was deliberate. By starting with a real experience and only then introducing AI, the conversation remained anchored in the texture of an actual moment rather than drifting into hypothesis. Responses reflected what participants had genuinely felt and needed, not what they imagined they might feel in some future moment of stress.

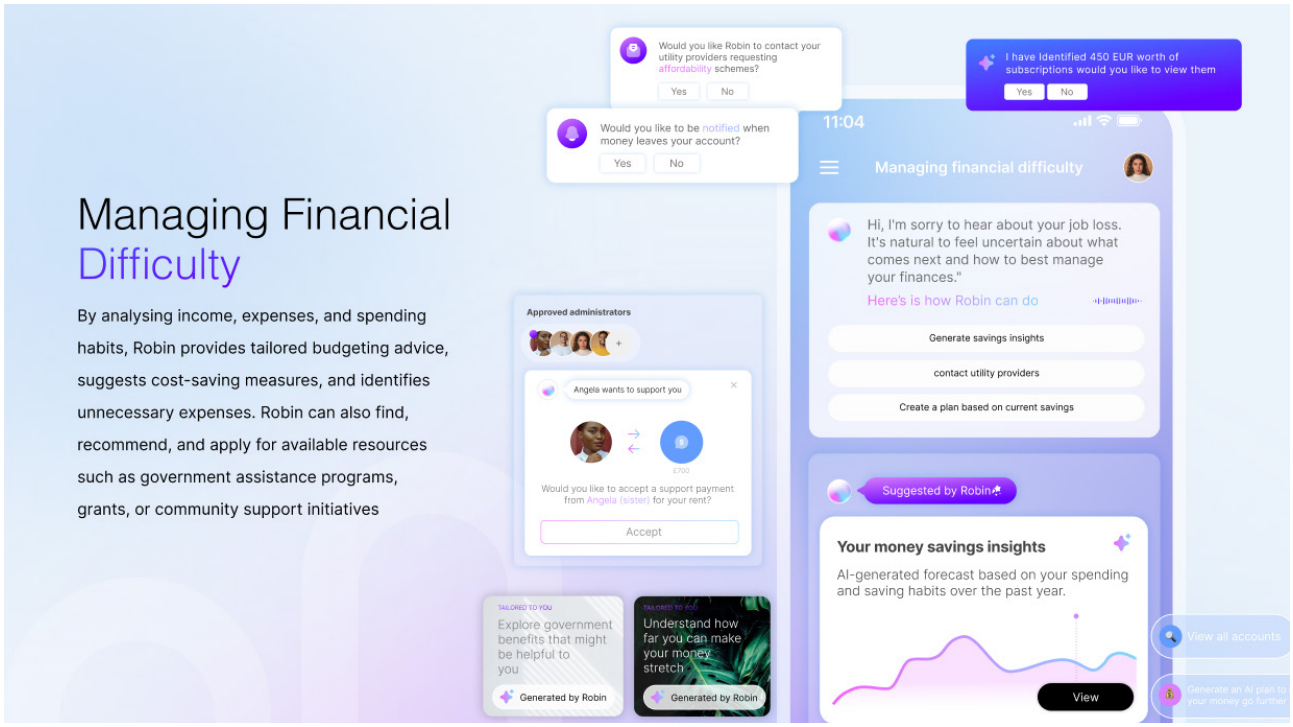

Robin prototype screen illustrating AI supported guidance in situations of financial difficulty, used as stimulus during qualitative interviews.

Section 2

The customer perspective: themes and insights

What needs to be in place for Dutch consumers to accept AI when decisions really matter? The interviews and survey point to a shared conclusion—one that challenges common assumptions in the technology debate. Acceptance is not driven by excitement or resistance. Instead, it depends on specific expectations that people apply deliberately and situation by situation.

These expectations are strikingly consistent. Regardless of age, financial confidence, or previous experience with AI, Dutch consumers tend to assess AI's role in similar ways. The following results set out the conditions under which AI is considered acceptable, where its role is limited, and how responsibility and control are expected to be shared when stakes are high.

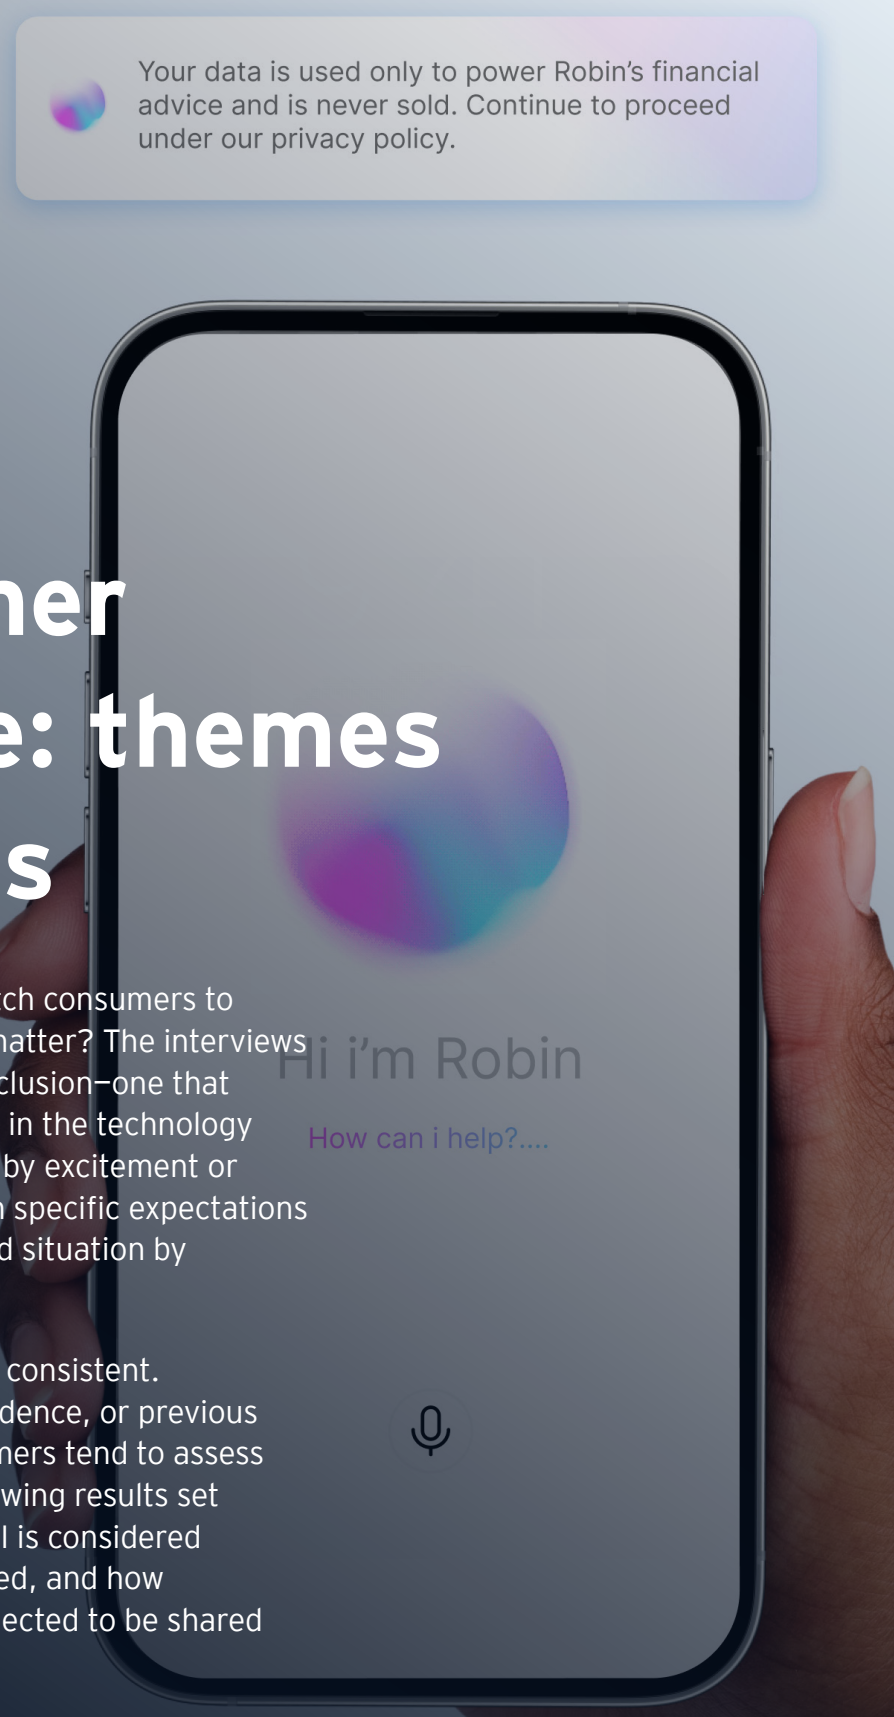

Four themes that define financial AI acceptance

Across the findings, four themes consistently emerge. Each is unpacked in detail later in this section, with supporting survey results, charts and direct quotes from interview participants.

Conditional acceptance, calibrated to autonomy

Acceptance is not driven by who the user is, but by what AI is asked to do. It is broad when AI explains, narrows when AI advises, and falls sharply when AI is expected to act or decide on the user's behalf.

01

The pragmatism paradox

Stated preference and actual behavior diverge. Dutch consumers prefer human contact in principle and use AI in practice when it offers tangible value, producing an implicit division of labor between the two.

02

Barriers are universal, but framed differently

The intensity of concern is comparable across generations, but the lens differs. Younger users worry about errors and privacy; older users worry about empathy and loss of control.

03

A gap between adoption and value

Adoption begins where AI feels safe to try, while perceived value sits where AI matters most. The two locations rarely overlap.

04

01 Conditional acceptance, calibrated to autonomy

Across the interviews, Dutch consumers approached AI with neither hostility nor enthusiasm. The recurring position was more measured: AI is a tool whose acceptability depends on what it is asked to do. Within that conditional stance, a clear and defining pattern emerged. Acceptance is broad when AI explains information or clarifies a complex situation, becomes more conditional when AI offers advice, and is largely withdrawn when AI is expected to act or decide on the user's behalf.

Average acceptance (1 = low, 5 = high), by AI role and generation. n = 2,000+

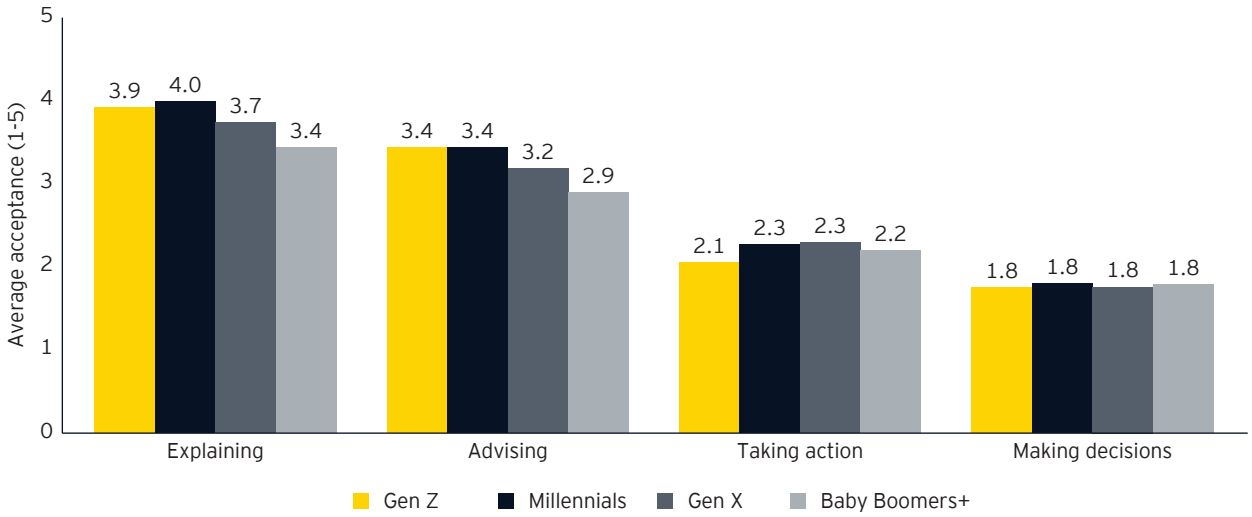

Source: EY Studio+ — Human Signals, Dutch Edition (May 2026). Quantitative survey, n = 2,000+ Dutch consumers.

Figure 1. Acceptance of AI roles across levels of autonomy in financial decision making (n = 2,000+).

This pattern likely stems from a distinction users draw, often implicitly, between reversible and irreversible interactions. Explanation is reversible: the user retains the option to ignore the explanation, request another, or step away. Action is not. Once a message has been sent, a form submitted or a payment authorized, the user is no longer evaluating advice but managing the consequences of a decision that has, in functional terms, already been taken. What participants resist is not AI in general, but the disappearance of the deliberation moment in which they can still decline.

What this quote adds is the framing of AI as a partner in preparation rather than execution. The participant welcomes AI in a role that reduces stress and structures information, but stops short of granting it authority over the action itself. The willingness is real, but bounded.

> “Yes, I think it could be very helpful. These kinds of things often cause people a lot of stress. You have to remember things, check everything, prepare documents, and not everyone is equally strong legally. So I think this could be a very good tool to support that.
>
> **Female, 26**

Trust is more solid than expected, but conditional

Across the survey sample, baseline trust in a financial AI assistant to perform correctly is meaningfully positive. 60% of respondents express high or very high trust (33% “to a very high degree”, 27% “to a high degree”), 20% are neutral and 20% express limited or no trust. This is a more encouraging starting point than commentary on AI scepticism often suggests, and it points to a population that is open to engagement rather than pre-positioned against it.

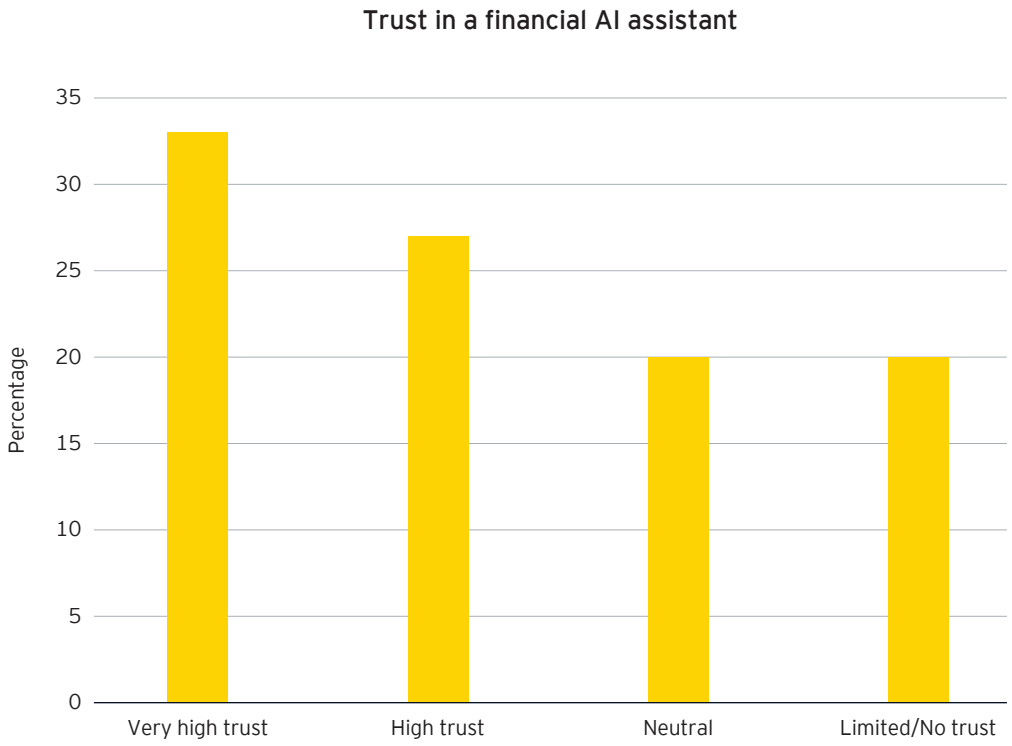

The level of trust, however, is less informative than its shape. Initial trust does not equate to unconditional approval; it represents a willingness to engage. The decisive question is therefore not whether trust exists at the baseline, but how it behaves under pressure. The survey provides a clear answer: it falls quickly as the role assigned to AI moves from supportive to autonomous.

A consistent autonomy gap across generations

When respondents rated their acceptance of different AI roles on a 1–5 scale, a clear and consistent hierarchy emerged across all generations. AI is most widely accepted as an explainer, with 60–80% of respondents indicating acceptance.

Acceptance softens when AI takes on an advisory role, and declines sharply when AI is positioned to act or decide on the user's behalf: only 9–17% accept AI taking action, and just 4–6% accept AI making decisions on their behalf.

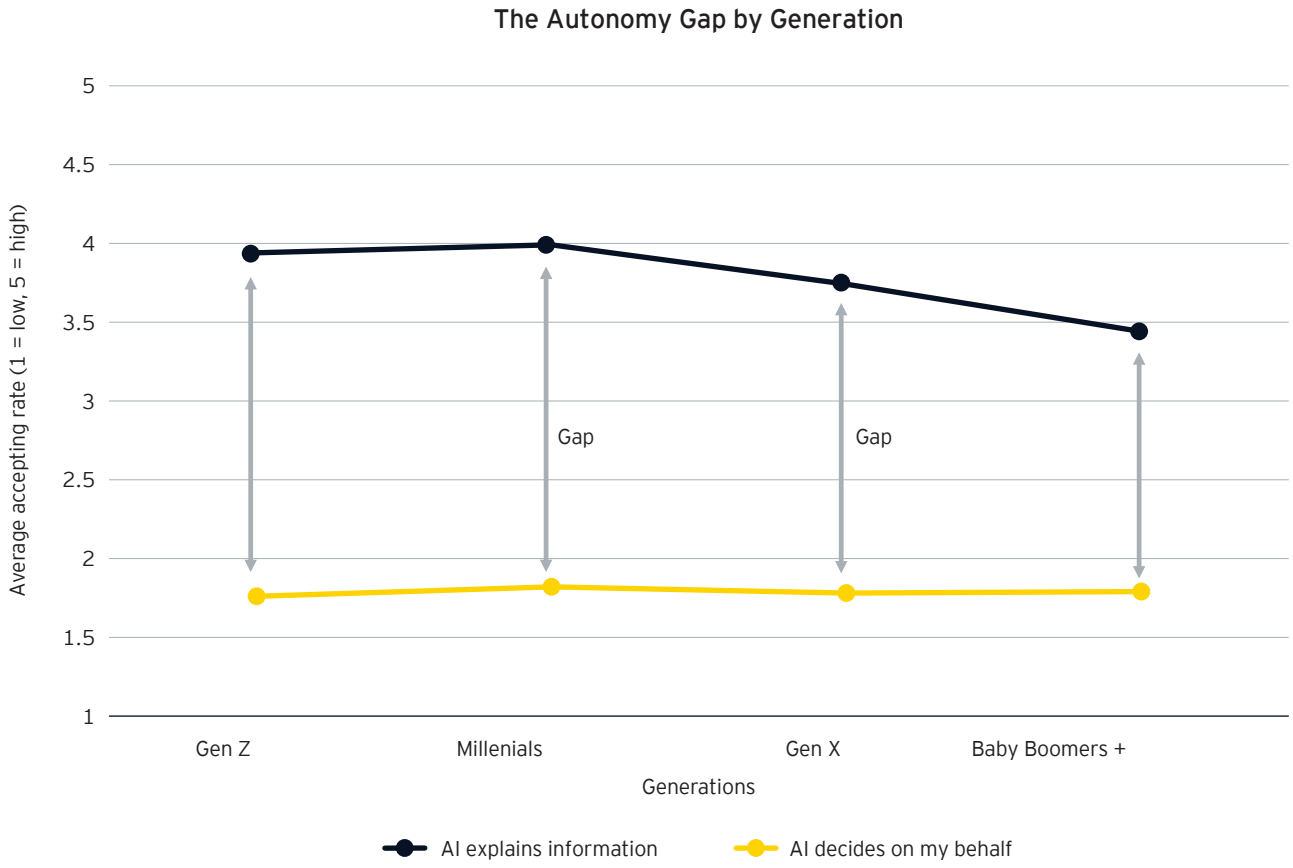

Figure 2. The autonomy gap by generation: average acceptance of AI as explainer vs. AI as decision-maker.

The generational difference is more subtle than the headline figures suggest. It lies primarily in the starting point rather than the slope. Younger respondents begin from a higher level of acceptance for explanatory AI, which produces a steeper visible decline as autonomy increases. Older respondents begin from a lower point and travel a shorter distance, but arrive at a comparable position. At the highest levels of autonomy, resistance converges across age groups and becomes near-universal.

> “Because it remains something that is generated, I feel you always have to check, in some way, whether it is actually correct.
> **Male, 61**

This convergence is a meaningful finding. It suggests that resistance to AI as a decision-maker is not a generational holdover that will dissolve as digital natives age into financial decision-making. Even respondents who use AI tools daily and trust their outputs in lower-stakes contexts continue to draw the same line in higher-stakes ones. The boundary appears to track the nature of the decision rather than the user's familiarity with the technology.

The phrasing here is precise. The participant does not reject AI; they accept it as a producer of useful output, but treats verification as a non-negotiable step. The instinct to retain a checking role is not driven by distrust of any specific output, but by the awareness that the system itself can be wrong in ways the user cannot always detect.

02 The pragmatism paradox

One of the most striking findings of the study is the gap between what users say they prefer and what they are willing to do in practice.

When asked directly, Dutch consumers express a clear preference for human interaction in complex financial moments. The same respondents accept AI in supportive and advisory roles at significantly higher rates, and report meaningful baseline trust in its ability to perform correctly.

This apparent inconsistency may be best understood as two valid answers to two different questions. Stated preferences reflect ideal expectations: how participants would prefer the world to be arranged if cost, time and effort were no constraint. Behavior reflects practical reality, shaped by speed, convenience, accessibility and the cost of pursuing a more cumbersome alternative. Dutch consumers do not need to prefer AI in order to use it. They need it to be useful, controllable and appropriately bounded.

The result is an implicit division of labor. AI handles the first layer of the interaction: structuring the situation, surfacing options, drafting routine communications. Human contact is reserved for situations that are genuinely complex, sensitive or carry significant personal consequences. Few participants articulate this division explicitly, but most operate by it in practice, treating channel choice as a question of fit rather than loyalty.

Embedded in this short statement is a definition of what AI is good for in the participant's view: rapid orientation.

The user is not asking AI for empathy or judgment in this moment, but for a clear sequence. Where AI delivers that, it earns its place in the workflow.

> "I want to quickly see: this step leads to the next step.
>
> Male, 59

Human preference is dominant, but not exclusive

When asked whether they expect to feel better understood by a human or by an AI in complex financial situations, respondents show a clear preference for human interaction. Approximately one-third strongly prefer a human advisor, while openness to AI in this emotionally charged framing ranges from 15% to 28%. Notably, Gen Z shows the lowest openness to AI in this specific question, which challenges the assumption that younger consumers are inherently more accepting of automation.

Gen Z's lower openness in this question may reflect not a rejection of AI in general, but a sharper awareness of its limits. Younger respondents have grown up with AI tools and are familiar with their failure modes; they may be more willing to use AI in routine contexts precisely because they are clearer about which contexts those are. Familiarity does not always translate into broader trust. Sometimes it produces a more discriminating one.

Having grown up in a tech driven world, **Gen Z** is often more **aware** of the limitations of modern systems

Read alongside the behavioral data, the pragmatism paradox reasserts itself. In principle, users prefer humans.

In practice, that preference does not prevent them from using AI when AI offers clear and tangible value. The two findings are not in tension; they describe different layers of the same position.

This quote captures the division of labor in a single sentence. AI is used as preparation; the human is retained for the parts of the situation that benefit from interpretation. The two are not in competition; they occupy different points in the same journey.

What surfaces here is a calibration the user makes almost automatically: AI for general matters, humans for safety. The participant does not articulate this as a principle, but applies it consistently. It is the kind of decision rule that organizations designing AI services will need to learn to read.

> Yes, I use AI also as a supporting tool. Because if you are trying to get a mortgage and you need help with that, it can still be nice to talk to a human, someone who can look a bit further and think along with you.
>
> **Female, 29**

> When it comes to safety, I simply feel safer with a person than with an AI. But when it comes to more general things, it doesn't really matter that much to me.
>
> **Female, 45**

03 Barriers are universal, but framed differently

Privacy and control emerge consistently as fundamental conditions for AI adoption. Users are not opposed to sharing personal or financial data in principle, but their willingness depends on how that data is handled and what risks they associate with it. The concerns voiced are concrete: fear of hacking or data breaches, uncertainty about where data is stored and processed, and unease about applications having access to information that feels deeply personal.

> “That’s private. I am willing to enter some information, but letting an app log into my bank? I would not trust that.
>
> Male, 36

Privacy in this context is not a single issue but a combination of security, transparency, consent and ongoing control. This composite character helps explain why simple reassurances rarely resolve the concern. A user worried about what happens to their data is asking several different questions at once, and an answer that addresses one of them while leaving the others unaddressed produces partial reassurance at best.

The boundary here is drawn not between data and no data, but between data the user actively provides and data the system retrieves on its own. The first feels controlled; the second feels surveilled. The same person willing to disclose information in a conversation may resist granting an application persistent access to the same information through a different mechanism.

Layered on top of these privacy concerns is a more nuanced finding around empathy. Dutch users acknowledge that empathy matters, but do not expect it as a default mode of interaction. The strong preference is for a factual, direct and structured communication style focused on clarity and concrete next steps. Empathy is welcomed in specific contexts, particularly during financial stress or major life events. Applied universally, it tends to feel unnecessary or even insincere.

This may stem from how Dutch users interpret the function of empathy in a service interaction. Empathy that arrives unbidden, in moments where the user has come for clarity rather than reassurance, can read as a stand-in for substance. Where empathy and clarity are both delivered, the empathy registers; where empathy substitutes for clarity, it tends to undermine confidence in the interaction it is meant to support.

The quote describes empathy as a context-sensitive choice rather than a default register. The participant is not opposed to warmth, but to warmth in the wrong moment.

Frustration, in her account, is resolved by progress on the underlying issue, not by acknowledgement of the feeling.

> “
>
> It depends. If you are frustrated, you don’t need to be coddled. I would much rather just get results.
>
> **Female, 42**

Same level of concern, different risk lens

When asked to rank reasons for not using AI, the survey data shows that the overall level of concern is comparable across generations. The framing of those concerns, however, differs meaningfully. Among Gen Z, 32.2% cite fear of errors as their primary barrier and 30.3% point to privacy concerns. Among Baby Boomers and older respondents, the most frequently cited barriers are the lack of human empathy (26.6%) and loss of control (19.0%).

Generation	Top barrier	Second barrier
Gen Z	Fear of errors — 32.2%	Privacy concerns - 30.3%
Baby Boomers+	Lack of empathy — 26.6%	Loss of control — 19.0%

The difference is not in the strength of concern but in its orientation. This may stem from the lens each generation brings to AI in the first place. Younger users approach AI through the experience of digital natives, where errors and data leaks are familiar risks framed in those terms. Older users approach AI through decades of human service relationships, where the salient questions are about recognition, accountability and the locus of decision-making. Both groups arrive at AI with caution, but the caution is shaped by different histories.

The implication for organizations is that there is no single barrier to AI adoption and no single message that addresses all of them simultaneously. A communication strategy that emphasizes data security will land for younger users while leaving older users untouched, and vice versa.

Why these barriers are perceived as risky

Although users cite different reasons for hesitating to engage with AI, these barriers should not be ranked against one another. They reflect parallel ways of interpreting and managing the same underlying uncertainty about a system whose inner workings are largely opaque to the user.

For some, that uncertainty circles around output: how the AI arrives at its conclusions, when those conclusions might be wrong, and how a user would recognize an error. In an information environment in which confidently stated falsehoods circulate easily, errors that present as authoritative are particularly difficult to detect. For others, the uncertainty circles around input: who has access to the data, how it is processed, and what happens if that custody fails. For others still, it manifests as a question of presence: whether anyone, human or otherwise, is genuinely attending to the situation.

These are different vocabularies for the same underlying anxiety. They shape how risk is interpreted rather than how strongly it is felt. A communication strategy that addresses only one of them will resonate with one segment of the audience and miss the others.

What this quote demonstrates is that scepticism does not always translate into rejection. The participant has developed a verification routine that allows her to engage with AI while preserving her own critical reading. Her trust is not in the system, but in her ability to check the system. That distinction matters for design: the AI services that earn engagement from sceptical users are those that make verification easy, not those that demand trust upfront.

This quote shows the inverse pattern: a user for whom emotional register matters more than verification. In her case, the AI's tone is part of what makes the interaction tolerable. The two participants are not contradicting each other; they are describing different conditions under which engagement becomes possible.

> “
>
> No, I would keep the privacy aspects restricted. That way I can still ask my question and get help, but without sharing any personal data.
>
> **Female, 29**

> “
>
> It is simply the unknown.
>
> **Male, 55**

> “
>
> I am quite skeptical. I realize that not everything is necessarily reliable. I do use it in a way where I ask it to support its answers with scientific articles, so I can check the sources myself. That way, I can quickly fact-check whether what it says is actually correct.
>
> **Female, 26**

> “
>
> I am very sensitive, and I can even end up yelling at a computer. Then I appreciate it when it approaches me like a human.
>
> **Female, 60**

04 A gap between adoption and value

Participants draw a clear distinction between the situations in which they would be willing to start using AI and those in which they perceive its greatest value. AI is readily accepted for simple, low-impact tasks, while its greatest potential is consistently described in complex and high-impact situations such as financial difficulty, bereavement or retirement planning.

This pattern likely stems from how trust is built rather than declared. Users want to test a new tool somewhere safe before extending it into the situations that matter most. Low-stakes contexts function as proving grounds. The bridge between the two is not bridged by capability; it is bridged by accumulated experience. As AI demonstrates reliability and value in routine contexts, the willingness to share more data and to extend its role in higher-stakes ones grows in step.

Trust, in this sense, behaves less like a setting that can be configured and more like a relationship that develops over time. The quote frames trust in AI as inseparable from trust in the institution behind it. This is consistent with the broader finding: AI is rarely evaluated on its own terms. It is evaluated as an extension of the organization that deploys it, and the credibility of that organization shapes the user's willingness to engage.

> “For me, it depends strongly on my relationship with the bank. With a positive relationship, using an app feels fine; with negative past experiences, I would be far less inclined to rely on one.
>
> Male, 37

Complaint handling is the universal entry point

When asked which scenario they would first consider using AI for, complaint handling is the dominant entry point across generations, selected by 62% to 73% of respondents. The harder, higher-value scenarios such as bereavement, retirement planning and home buying sit considerably further behind in terms of first-choice adoption.

First choice for AI use by financial scenario

Share of respondents selecting each scenario as their first choice, by generation.

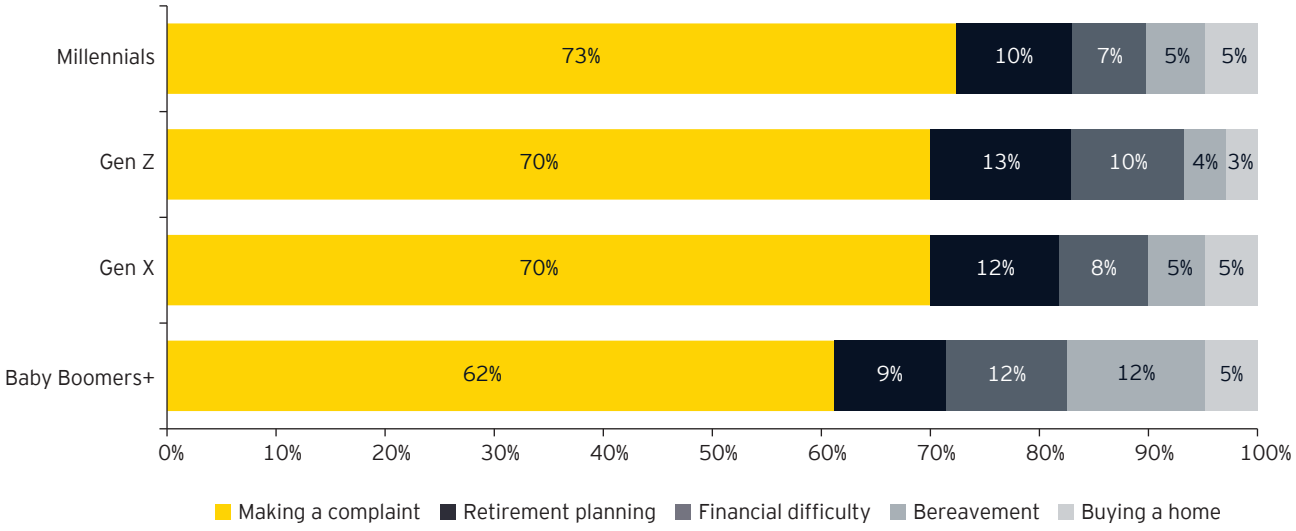

Source: EY Studio+ — Human Signals, Dutch Edition (May 2026). Quantitative survey, n = 2,000+ Dutch consumers.

Figure 3. First-choice scenario for using AI, by generation.

The dominance of complaint handling can be explained by the structure of the task itself. The user is irritated rather than vulnerable; the worst outcome of a poorly handled AI interaction is mild additional frustration rather than a serious financial setback; and the work involved, gathering facts, articulating the issue and proposing a remedy, aligns well with what AI does competently. The harder, higher-value moments demand the qualities that consumers remain uncertain AI can deliver: judgment, sensitivity, and a clearly assigned point of accountability if something goes wrong.

> “Yes, and I think there is a lot to gain in all these life-changing events, where there is so much documentation and hassle. This can really support you, so you can focus on what truly matters.
>
> **Male, 37**

For organizations, the gap between where adoption begins and where value resides is not a problem to be solved with a single product launch. It is a sequencing problem. Trust earned in low-stakes contexts is the route to engagement in higher-stakes ones.

Leading with the most autonomous, consequential capabilities is the most reliable way to lose the audience before the relationship has been established.

Read together, these two quotes illustrate the gap from both sides. The first articulates the value AI could deliver in life-changing events; the second explains why that value is not yet enough to overcome hesitation in those same events. The willingness exists in principle. What is missing is the demonstrated capability that would make it actionable in practice.

> “
>
> Yes, when it comes to something like handling complaints, the fact that it can already do things for you and take so much off your hands really appeals to me. And if it were well developed, I would certainly consider using it for mortgages as well. But based on what I have seen so far, I feel it still does not offer enough support in the areas where I personally need it most.
>
> Female, 29

A consistent pattern emerges

Read individually, these four themes can appear as four distinct findings. Read together, they describe a coherent set of conditions under which Dutch consumers are willing to engage with AI. Several strands run through the four themes and reinforce each other.

The most visible strand is **control**. Acceptance of AI is consistently tied to how much room the user retains to verify, override and disengage. The autonomy ladder is the most explicit expression of this principle: acceptance falls as the user's scope to intervene falls. The pragmatism paradox is its quieter expression: users accept AI in practice when they retain the ability to steer, even when they would prefer human contact in principle. The barriers, generationally varied as they appear, all reduce to versions of the same concern. Privacy is control over data; accountability is control over outcomes; the option to disengage is control over the relationship itself.

Alongside control, a second strand runs through the findings: trust as a developing relationship rather than a fixed setting. Baseline trust is meaningfully positive, but it is not unconditional approval. It is a willingness to engage that grows or contracts based on what users experience over time. The gap between adoption and value is the practical consequence of this dynamic. Users start where the cost of a poor experience is bounded, and extend their engagement as the relationship accumulates evidence of competence. Organizations that treat trust as a precondition to be argued for, rather than an outcome to be earned through demonstrated value, tend to misjudge the sequencing required.

A third strand concerns empathy, and reframes its role in the Dutch context. The findings do not show that empathy is unimportant; they show that it is conditional. Empathy is welcomed in specific moments, particularly during financial stress or major life events, but is read as performative when applied universally.

The implication is that empathy in AI services should be treated as a context-sensitive design choice rather than a default register, and that the system should be capable of switching modes as the user's situation requires.

A fourth strand emerges in the generational data and complicates the assumption that AI acceptance will simply rise as digital-native cohorts age. The intensity of concern is comparable across generations; what differs is the lens. Younger users are more attentive to errors and privacy; older users to recognition and the locus of decision-making. As Gen Z and Millennials move into the higher-stakes financial moments that older respondents currently occupy, their concerns may not fade so much as shift in vocabulary. Designing for one generation's risk frame is therefore unlikely to be a durable strategy. Designing for both, simultaneously, is the more demanding but more lasting approach.

Taken together, these four strands point to a single conclusion. The conditions Dutch consumers place on financial AI should not be read as obstacles to adoption. They constitute a specification of how AI must behave in order to be accepted: with control retained by the user, with trust earned through demonstrated value, with empathy applied where the situation calls for it, and with the recognition that different segments of the audience read the same risks through different lenses. Treated as a brief rather than a barrier, most of these conditions are designable.

> In the Netherlands, financial AI is trusted only as far as users remain in **control**, and only as long as that control is matched by demonstrated value, contextual empathy and a clear sense of who is accountable.

Section 3

Implications for organizations

Acceptance of AI is conditional rather than categorical, this is not a limitation; it is a design specification. The findings point to a market where meaningful adoption is possible, provided AI is introduced with the right sequencing, role definition and service design. The commercial opportunity lies not in deploying more AI, but in deploying it where it compounds trust rather than erodes it. The implications below describe a way of positioning AI that aligns with how Dutch users actually want to engage with it.

01 Reframe the question

Rather than asking whether customers accept AI, the more useful question is in which specific moments, and for which specific tasks, AI is accepted. Users do not respond to technology in isolation, they respond to the role it plays within their decision-making process. The same AI capability may be welcomed in one scenario and rejected in another by the same user. Strategy should therefore be built around situational relevance rather than generic adoption metrics, with roadmaps grounded in the moments that matter to customers rather than the features available to be deployed.

02 Design for need states, not just use cases

Use cases describe what gets done. Need states describe why the customer is there in the first place, and how the moment feels to them. In high-emotion, high-complexity moments, customers come to their financial provider with both functional and emotional needs: the task at hand, and the underlying desire to feel less overwhelmed, less alone with the situation, and clearer about what comes next.

The strongest AI experiences are designed against the underlying need state, not only against the task. A complaint flow that begins by helping the customer structure what happened, before asking them to file anything, is responding to a need state. A flow that begins with a form is responding only to a use case. The difference is small in implementation and significant in how it is experienced.

03 Build trust through demonstrated value

Trust is not a precondition for adoption; it is a result of interaction. Users build trust gradually, as they experience consistent value, reliability and transparency. The implication for implementation is direct: do not lead with the most autonomous capabilities. The most reliable way to lose a hesitant user is to ask them to entrust a sensitive decision to a system that has not yet earned a smaller one.

Start with low-risk, high-clarity applications, such as explaining a product, summarizing a statement or surfacing the next step, and allow the relationship to deepen as positive experiences accumulate. The temptation to skip this stage is real, particularly when leadership is focused on visible and ambitious deployments. The cost of skipping it is paid in adoption that never converts into engagement.

04 Match the level of autonomy to the level of stakes

Acceptance follows autonomy. AI as an explainer is broadly welcomed. AI as an advisor is conditionally welcomed. AI that acts or decides on the user's behalf is rejected at scale. The design implication is to calibrate the role of AI to the emotional and financial weight of the situation: keeping AI in a supportive layer where stakes are high, and granting greater agency only where stakes are low and outcomes are easily reversible.

What this asks of designers is the willingness to let AI do less than it is technically capable of doing. That restraint should not be read as a weakness. It is the surface on which trust accumulates.

05 Design the human–AI relationship explicitly

Dutch users do not expect, or want, full automation. They expect a system in which they remain in control and a human is reachable when needed. AI should not be designed as a standalone solution but as part of a broader service ecosystem. The clarity of the pathway between AI and human support, including when it triggers, how it triggers and who the human is, is as important as the AI capability itself, and becomes more important as complexity or emotional intensity increases.

An AI service whose escalation path is buried, slow or conditional on completing further AI steps is read as a service designed to deflect. That perception, once formed, is difficult to repair. The reverse is also true: A service whose path to a human is clear, fast and unconditional is read as confident enough in its own capability to acknowledge its limits, which paradoxically increases trust in the AI itself.

06 Make accountability explicit

Across the interviews, acceptance of AI was closely tied to the presence of clear accountability. Users are more willing to engage with AI when there is a clearly identifiable party, whether human or organizational, that can be held responsible if something goes wrong. This does not require constant human involvement, but it does require clarity: who is accountable for outcomes, errors and unintended consequences.

Designing AI-enabled services therefore includes designing accountability. Where responsibility is diffused across vendors, models and teams, users tend to sense the diffusion even when they cannot articulate it. Where responsibility is named and visible, the AI itself carries less of the trust burden, because users understand that the trust they extend rests, ultimately, with the institution behind it.

07 Make transparency and control visible by default

Across every finding, transparency and control surface as fundamental conditions for adoption. Users want to understand how AI arrives at its outputs, how their data is used, and what they can switch off. These should not be buried in settings; they should be visible parts of the experience. Even in lower-risk contexts, perceived control shapes user confidence more than the underlying performance of the technology.

The implication is to surface controls as part of the experience rather than as a compliance afterthought, and to design them so that exercising them feels routine rather than exceptional. A user who feels obliged to apologize for opting out of personalization is being given the wrong control.

08 Tailor barriers, not benefits

The barriers to adoption differ in framing rather than in intensity. Younger users worry about errors and privacy; older users worry about empathy and loss of control. A single, unified message will not resolve both. Communications and onboarding should address the specific risk frame relevant to each segment, rather than relying on a generalized articulation of the technology's benefits.

09 Plan for the gap between adoption and value

Customers will most readily adopt AI in low-risk, low-emotion scenarios. The greatest value of AI lies in harder moments such as financial difficulty, bereavement and retirement planning. Design entry points where users feel safe, and design pathways from those entry points into the moments where AI can genuinely change outcomes. The objective is not to accelerate adoption of the harder scenarios, which tends to produce resistance, but to make the route from one to the other intelligible, so that users who have built confidence in low-stakes interactions know how to extend that confidence when the higher-stakes moments arrive.

From capability to context

Taken together, these implications point to a single conclusion. Successful AI is not defined by technological capability alone, but by the precision with which that capability is positioned within the right context. For organizations, this means AI strategies must be grounded in a deep understanding of the specific use case, the user's situation, and the role AI is being asked to play in that moment.

This rewards restraint, sequencing and a willingness to leave space where space is what users require. It also rewards specificity. The more precisely an organization can describe the moment it is designing for, the more likely it is that the AI it deploys into that moment will be welcome there.

> **The next generation of AI-powered financial services will not be defined by the organizations that automate the most. It will be defined by those that understand, with precision, when to step forward and when to step aside.**

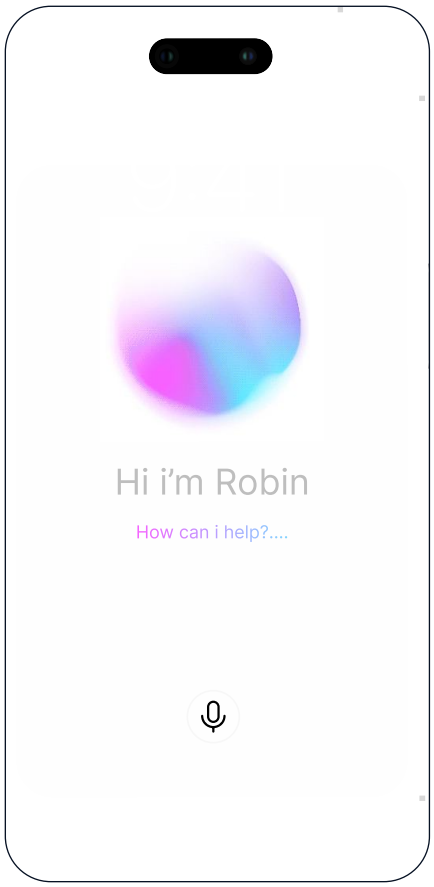

Section 4

Cross-market comparison: the United Kingdom and the Netherlands

Same research questions, same broad methodology, two different cultural contexts. What does reading the UK and Dutch findings alongside each other reveal? The shared dynamics are real, but the way they are expressed diverges in places that matter for design and strategy. The differences are not contradictions; they are different ways of resolving the same underlying tension between support and control.

Two markets, one underlying logic

Both studies are grounded in the same fundamental question: how do consumers engage with AI-supported services in situations that are complex, personal and often emotionally charged? In both the UK and the Netherlands, the answer is not binary. Users do not simply accept or reject AI. They evaluate its role in relation to the situation in which it is applied. In both markets, AI is broadly accepted when it brings clarity and reduces complexity, and that acceptance falls away as autonomy rises, particularly when AI is expected to act or decide on the user's behalf.

That is the shared logic. The differences lie in how the logic is expressed.

The balance between support and control

Across both markets, acceptance is governed by a single underlying tension: the balance between support and control. Consumers welcome AI when it deepens their understanding and helps them navigate complexity. They resist relinquishing control over decisions that carry personal or financial consequences.

The difference lies in where users place the emphasis. UK consumers focus more strongly on the experience of the interaction, evaluating whether they feel understood, recognized and supported. Dutch consumers focus more strongly on the outcome, evaluating clarity, efficiency and a direct understanding of what to do next. These are not opposing values but different reading orders.

> "AI is much more efficient. And you have to ask yourself whether a human advisor would really give you the same things you get from AI, because AI simply knows so much.
>
> **Dutch participant, female, 59**

Empathy: central in the UK, situational in the Netherlands

The most distinctive contrast between the two markets concerns the role of empathy. This contrast should be interpreted carefully. The point is not that Dutch consumers are unfeeling, or that British consumers are more in need of reassurance. Both populations expect to be treated as adults navigating complicated situations. They differ in the form of respect they expect that treatment to take. In the Dutch register, respect is more often expressed through clarity, brevity and the assumption that the user can handle the facts. In the UK register, respect is more often expressed through attentiveness, acknowledgement and sensitivity to the emotional weight of the moment. Neither is more sophisticated than the other; they are different defaults that any AI service operating across both markets must be able to switch between.

> In the UK, **empathy** is the lens. In the Netherlands, **control** is the lens.

Trust: emotional space versus rational assurance

Trust matters in both markets, but is interpreted differently. In the UK, AI is partly framed as a potential "neutral space", a non-judgmental environment in which users can engage with sensitive topics without shame. This emotional neutrality contributes meaningfully to a sense of safety, and helps explain why participants in the UK study sometimes preferred AI to a human in moments of embarrassment or vulnerability.

In the Netherlands, trust in AI is approached primarily as a rational judgment rather than an emotional response. Users emphasize concrete criteria such as data security, transparency and accountability when deciding whether an AI system can be relied upon. Trust is therefore evaluated rather than felt. Users ask whether they understand how the system works, whether it behaves predictably, and whether responsibility is clearly assigned.

Taken together, these differences show that while trust matters in both markets, it operates through different logics.

> In the UK, **trust** is shaped by emotional safety and a non judgmental space. In the Netherlands, **trust** is established through rational evaluation.

For organizations operating across both markets, the same product will need to demonstrate different things first. In the UK, the early signal that earns engagement is often emotional fluency. In the Netherlands, the early signal is often procedural clarity. Getting the order wrong is a more reliable way to lose users than any deficit in capability.

Human versus AI: bias and pragmatism

In both markets, a clear preference for human interaction persists in complex or high-stakes situations. The UK research describes this preference as a learned bias, shaped by past experiences with limited chatbots and likely to evolve as AI becomes more capable.

In the Netherlands, the same preference is observable but is accompanied by a stronger streak of pragmatism. Dutch users prefer humans in principle and use AI in practice when it offers tangible benefits in speed, convenience or efficiency. The result is a hybrid model in which the choice between humans and AI is driven less by principle and more by what works in the moment.

This produces a position that is unusually stable, less a transitional state on the way to full automation than a settled division of labor in which humans and AI each occupy roles the other does not perform well. Strategies that assume Dutch consumers are simply waiting to be persuaded out of their preference for humans are likely to underestimate how rational, and how durable, that preference actually is.

Context as the defining factor

Taken together, this comparison makes one point clear. While the principles underlying AI acceptance may be broadly consistent across markets, their realization is always shaped by context. Expectations of AI are not formed in the abstract, they emerge from locally grounded norms about communication, decision-making and trust. What matters is not whether these expectations differ, but how they give concrete form to shared concerns about agency, responsibility and meaning.

For organizations, the implications are substantial. AI acceptance cannot be treated as a universal phenomenon. Any meaningful implementation must be understood within a specific context: the country, the sector, the use case and the situation of the user. Ignoring these contextual layers risks misalignment between technological capabilities and user expectations, and tends to produce services that perform poorly in every market simultaneously.

The deeper point is that AI strategies do not stand outside culture, they sit within it. The technology is global; the meaning of the technology is not. Organizations that take the time to understand which meaning is operating in which market will be better placed to design services that earn the kind of trust this report has been describing throughout: trust that is given conditionally, monitored carefully, and extended only gradually to the moments that matter most.

> **The role AI is allowed to play is never determined by the technology itself, but by the meaning users assign to it within a given context.**

Appendix

Research methodology

This report is based on a mixed methods research approach combining qualitative depth with quantitative validation. The methodology was designed to explore how Dutch consumers engage with AI supported financial services, particularly in moments that are complex, high stakes or emotionally significant.

Qualitative research

The qualitative component consisted of 32 one to one in depth interviews with Dutch consumers. Each interview lasted approximately 75 minutes and was conducted remotely.

Participants were selected to reflect a broad range of life situations, financial confidence levels and attitudes towards technology. All participants had recently experienced at least one financially complex or emotionally significant situation (such as complaint handling, financial difficulty, home buying or life event related decisions), ensuring that discussions were grounded in lived experience rather than abstract opinion.

The interviews followed a semi structured format, allowing for consistency across conversations while leaving room for individual experiences, hesitations and contradictions to surface. Visual stimuli, including prototype scenario cards, were used to help participants react to concrete examples of AI supported services operating at different levels of autonomy.

Quantitative research

The qualitative findings were validated through a large scale quantitative study conducted among 2,000+ Dutch consumers, aged 18 and above. The sample was structured to provide a good reflection of the Dutch population across key demographic variables.

The survey focused on four core areas directly aligned with the qualitative themes:

- Baseline trust in a financial AI assistant to perform the correct actions in financially complex situations
- Perceived understanding and preference when comparing advice from a human employee versus a financial AI assistant in complex situations

- Acceptance of AI roles across levels of autonomy, distinguishing between AI explaining information, providing advice, taking action, or making final decisions on the user's behalf
- Barriers to adoption, explored through a ranked assessment of concerns such as fear of errors, lack of human empathy, privacy and data concerns, loss of control, and the perceived "black box" nature of AI

In addition, respondents were asked to rank financially complex scenarios, including complaint handling, buying a home, retirement planning, inheritance and financial difficulties, to indicate where they would be most willing to first engage with an AI assistant.

Read alongside the qualitative interviews, the quantitative data provides population level validation of the patterns observed in depth. Together, the two methods reveal not only whether Dutch consumers accept AI, but under which conditions, in which roles, and in which moments that acceptance holds.

Key Definitions

Generative AI: Generative AI (GenAI) refers to a category of artificial intelligence systems designed to generate new content, such as text, images, audio, video or code, by learning patterns from large datasets and applying those patterns to produce new outputs. Unlike traditional analytical systems, generative AI does not merely retrieve or summarize existing information but synthesizes new material based on its learned representations.

Agentic Generative AI: Agentic Generative AI refers to AI systems that, within a clearly defined scope, are authorized to take actions or make decisions in pursuit of specific goals. These systems move beyond explanation or recommendation by executing tasks on the user's behalf, such as progressing a process, initiating actions, or coordinating next steps.

In the context of this research, agentic AI is not understood as unbounded autonomy. Its acceptability depends on explicit limits around when it may act, what it may act upon, and how users can intervene, override, or escalate to human support. Learning or adaptation over time is acceptable only insofar as it remains transparent, controllable, and accountable within these boundaries.

AX (Artificial Intelligence Experience): Artificial Intelligence Experience (AX) describes the overall quality and character of a person's interaction with AI enabled systems. Similar to Customer Experience (CX), AX captures both functional and experiential dimensions, including clarity, usability and perceived value.

In this report, AX places particular emphasis on trust, user control, transparency and the handoff between AI and human support. A positive AX is not determined solely by technical performance, but by whether users feel able to understand what the system is doing, retain agency over important decisions, and rely on clear accountability when outcomes matter.

Prototype scenario card stimulus

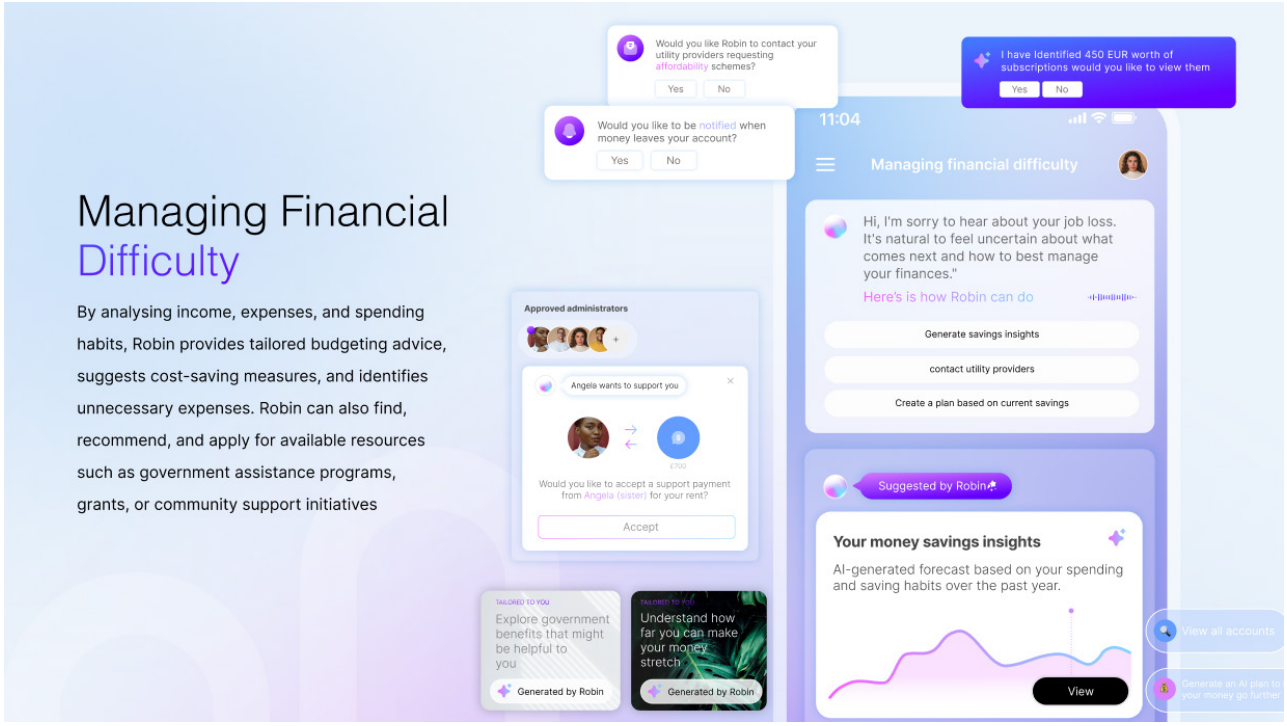

Scenario: Financial difficulty
After losing your job, you face a sudden shift in your financial situation. With ongoing financial responsibilities, you want to assess your position and understand your next step

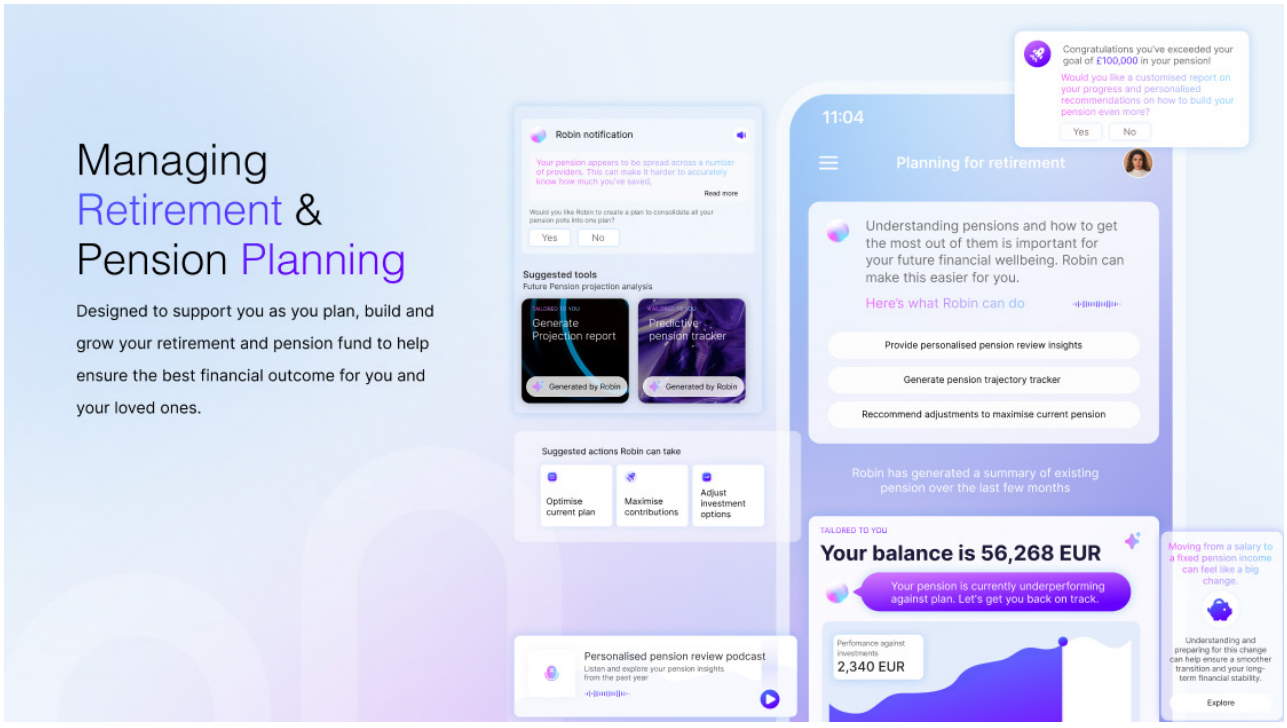

Scenario: Planning for retirement
You realize that it's time to think about your pension. While retirement can feel distant, you're thinking through if you're doing enough and what more you could be doing to prepare.

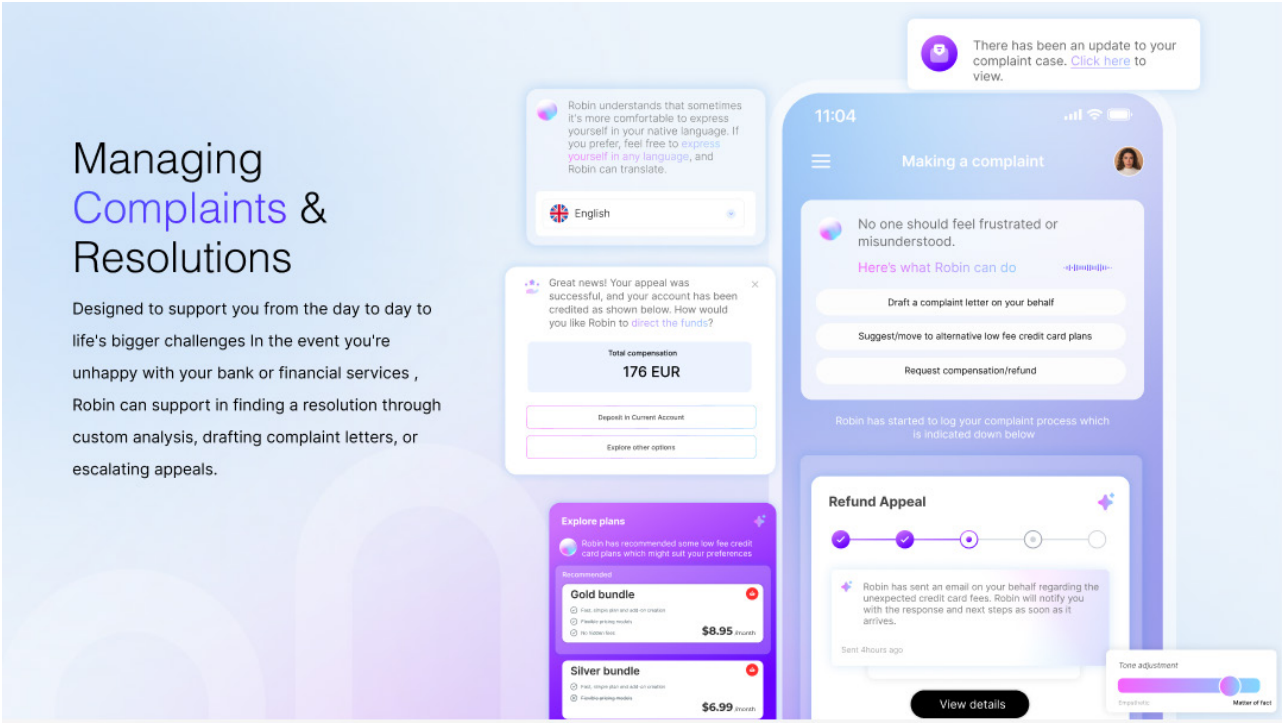

Scenario: Making a complaint
You've noticed unexpected fees on your bank account, leaving you frustrated and annoyed. You want to file a complaint to ensure your concerns are heard and the issue is resolved.

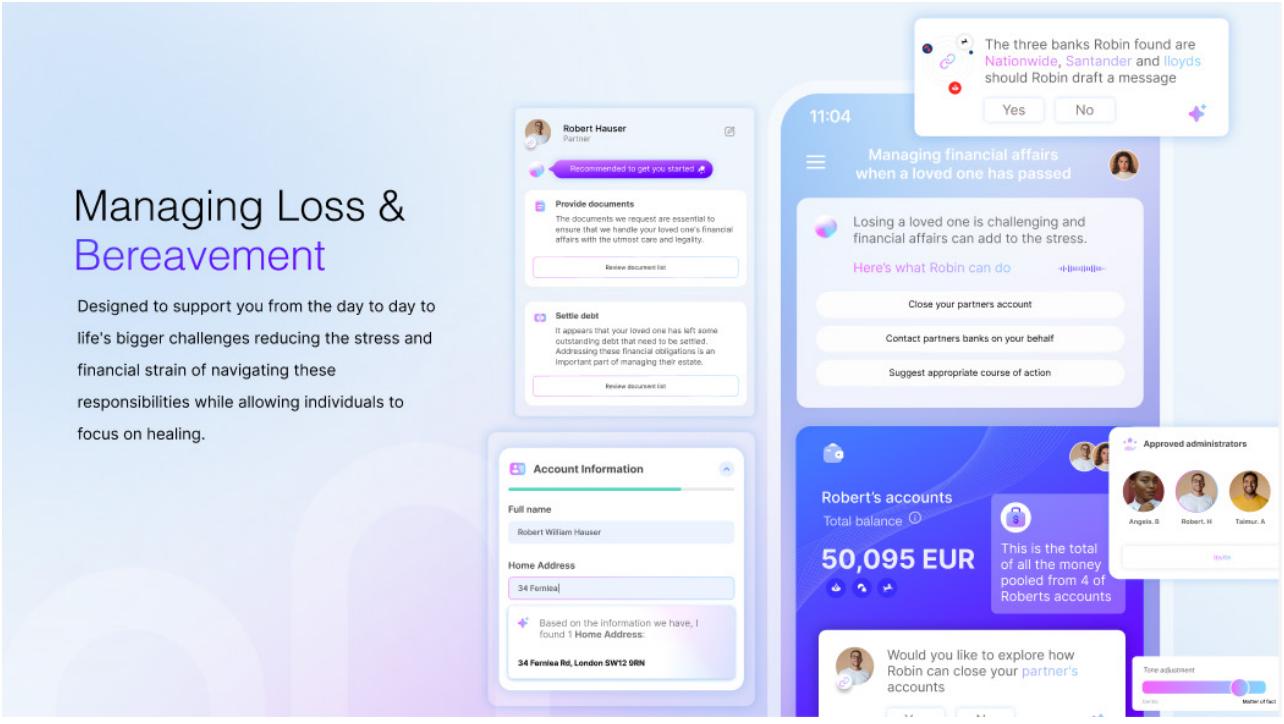

Scenario: Bereavement
After your partner's passing, you have stepped into the role of executor. This means handling the difficult task of closing their bank accounts and settling their affairs. It's a process that requires careful attention, and you are focused on navigating this step by step.

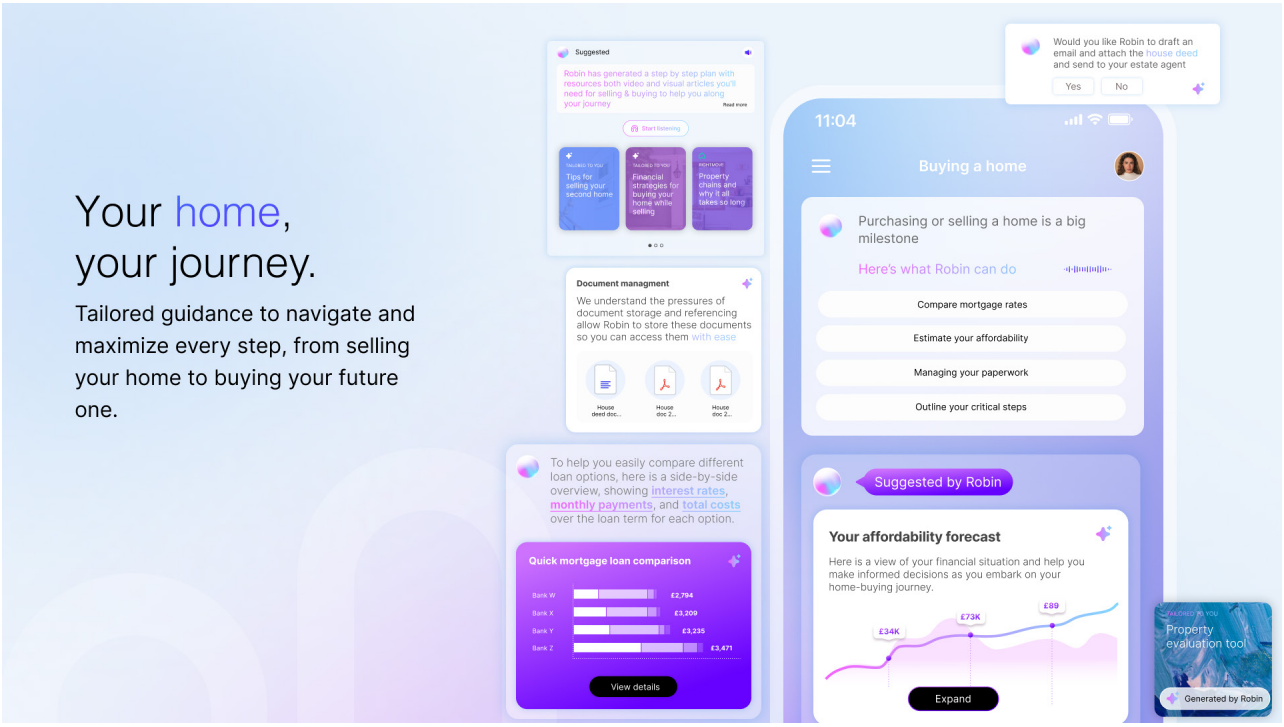

Scenario: Buying a home

You are currently looking to sell your current home and buy a new one which feels like double the workload. You need to understand what's involved and assess the financial implications.

Thank you

If you would like a briefing of this research or a discussion on how your teams can apply these insights, please contact: Patrick Ruijs at patrick.ruijs@studio.ey.com

Meet the Human Signals team

Patrick Ruijs
Partner
EY Studio+ | NL

Boaz Burkunk
Design Lead
EY Studio+ | NL

Claire Galjart
Senior Service Designer
EY Studio+ | NL

Anouk Mouthaan
Researcher
EY Studio+ | NL

Peter Neufeld
Founding Editor of Human Signals
Global Head of Financial Services, EY Studio+
EY Studio+ | UK&I

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets. Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multi-disciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

About EY Studio+
EY Studio+ creates transformative experiences that move people and shape markets. Our solutions work at the intersection of creative design and commercial insight, powered by cutting-edge technology — to drive real business value in a captivating way. We reimagine customer centricity and shape new possibilities at scale, drawing on EY's full spectrum of services, sector insights and diverse ecosystem partners. EY Studio+ consulting services related to front office transformation are an offering by a number of EY member firms across the globe.
For more information, please visit studio.ey.com.

ED None

ey.com