

The Psychological Speed Gap Framework

A human-centred model for responsible AI adoption



The Psychological Speed Gap Framework

Signature idea

AI evolves in months. Trust evolves in years. The Psychological Speed Gap sits in between.

Enhanced framework documentation
Version 1.0

Document purpose

This document provides an enhanced, fully usable version of the Psychological Speed Gap Framework. It converts the original framework into a structured guide that can support strategy, policy, service design, AI governance, customer experience, workforce development and responsible technology adoption.

The document is written for leaders, product teams, risk and compliance teams, customer operations, design teams, data science teams and any organisation seeking to deploy AI in a way that remains trustworthy, inclusive and human-centred.

Core proposition

The central question is not only whether AI can technically perform a task. It is whether people and institutions can realistically understand, govern, trust and navigate the change.

Executive summary

Artificial intelligence is moving at a pace that is difficult for institutions and individuals to match. New models, automation capabilities, analytics and intelligent agents can emerge in months, while organisations need time to govern, implement, test, train and embed them. People need even longer to build understanding, trust and confidence.

The Psychological Speed Gap describes the widening distance between these different speeds of change. It is the space between technological capability, institutional readiness and human psychological adaptation. When the gap widens, services may become harder to navigate, trust may weaken, support demand may increase and vulnerability may become more likely during digital interactions.

The framework is built around five connected layers:

- The Psychological Speed Gap, the core concept that names the distance between AI speed and human or institutional adaptation.
- The Three Curves Model, which visualises AI Capability, Institutional Adaptation and Human Psychological Adaptation moving at different speeds.
- SAFE-AI, which explains the forces behind the gap: Speed of Technology, Human Adaptation, Institutional Friction and Ethical Response.
- The Human Adaptation Layer, which identifies psychological drivers and situational amplifiers that shape behaviour, trust and confidence.
- HARMONIC, which provides a practical set of principles for closing the gap through responsible, inclusive and confidence-building design.

This enhanced version also adds practical guidance for interpreting behaviour, designing services, governing AI, measuring risk and embedding the framework into organisational decision-making.

Framework at a glance

Layer	Purpose	Practical question
Psychological Speed Gap	Defines the widening distance between AI speed and the pace at which institutions and people can adapt.	Where is technology moving faster than readiness, understanding or trust?
Three Curves	Shows the relationship between AI Capability, Institutional Adaptation and Human Psychological Adaptation.	Which curve is moving fastest, and where is the gap widening?
SAFE-AI	Explains the forces that create the gap.	What is driving the pressure: technology speed, human adaptation, institutional friction or ethical response?
Human Adaptation Layer	Explains behaviour through psychological drivers and situational amplifiers.	What might behaviour be telling us about cognitive load, trust, control, risk or confidence?
HARMONIC	Provides responsible design and governance principles.	How do we close the gap without removing human control, support or dignity?

Figure 2. Framework Architecture

A layered model moving from diagnosis to explanation, behavioural insight and responsible response.




Figure 1. Framework architecture, moving from diagnosis to response.

Contents

- 1. The core concept**
- 2. The Three Curves Model**
- 3. Workforce Development Capacity**
- 4. SAFE-AI, the forces behind the gap**
- 5. The Human Adaptation Layer**
- 6. HARMONIC, principles for responsible response**
- 7. Behavioural interpretation guidance**
- 8. Design and governance guidelines**
- 9. Applying the framework across the AI lifecycle**
- 10. Measurement and maturity assessment**
- 11. Practical tools and checklists**
- 12. Example use cases**
- 13. Glossary**
- 14. Closing summary**

1. The core concept

1.1 Definition

The Psychological Speed Gap is the widening distance between the pace of technological change and the slower pace at which institutions and people can realistically adapt to it.

AI capability may develop rapidly, but organisations must still work through regulation, governance, infrastructure, risk management, operating model change and workforce development. People must also build understanding, confidence and trust through lived experience. These forms of adaptation are slower, uneven and context-dependent.

The gap appears when the rate of technological acceleration exceeds the rate of institutional readiness and human psychological adaptation.

Plain language definition

The Psychological Speed Gap is what happens when technology changes faster than people and organisations can confidently understand, govern and use it.

1.2 Why the gap matters

The gap matters because AI is not experienced as an abstract technical capability. It is experienced through journeys, decisions, messages, prompts, recommendations, security checks, support interactions and outcomes. When those experiences move faster than a person can understand or trust, the technology may become technically impressive but psychologically difficult to navigate.

This is especially important in financial services and other regulated sectors, where technology intersects with money, trust, identity, fraud risk, vulnerability and personal wellbeing.

When the gap widens, organisations may see:

- customers misunderstanding automated or digital processes
- people becoming less confident in institutional decisions
- increased pressure on contact centres and support teams
- higher risk of exclusion for people with low digital confidence
- more exposure to fraud, coercion or misleading external influence
- over-reliance on automated outputs where human judgement is still needed
- unintended harm caused by journeys that are efficient but not navigable

1.3 What the framework is, and what it is not

The framework is	The framework is not
A human-centred model for understanding the relationship between AI speed, institutional readiness and human adaptation.	A rejection of AI or a claim that technology should not evolve.
A guide for designing AI-enabled services that remain understandable, inclusive and supportable.	A narrow technical architecture model.
A behavioural model that helps interpret hesitation, confusion, avoidance, over-trust and loss of confidence.	A diagnostic tool for labelling people as vulnerable or incapable.
A governance lens for identifying when organisations need to slow down, support, explain or escalate.	A substitute for regulatory, legal, ethical or clinical judgement.

2. The Three Curves Model

The Three Curves Model visualises the Psychological Speed Gap as the space that opens between three different speeds of change: AI Capability, Institutional Adaptation and Human Psychological Adaptation.




Figure 2. The Three Curves Model, showing where the Psychological Speed Gap widens.

2.1 Curve 1, AI Capability

The AI Capability Curve represents the rapid acceleration of artificial intelligence. It includes the development of more powerful models, generative AI systems, behavioural analytics, intelligent automation and agents capable of performing multi-step tasks.

This curve rises steeply because technical capability can improve through research cycles, data access, infrastructure investment, competitive pressure and rapid product iteration. In many contexts, the capability curve is now measured in months rather than years.

2.2 Curve 2, Institutional Adaptation

The Institutional Adaptation Curve represents the pace at which organisations can safely adopt, govern and embed AI into real services. It rises more slowly because institutions must manage complexity and protect customers.

Institutional adaptation is shaped by:

- regulation and supervisory expectations
- risk management and assurance processes
- data protection, model governance and ethical review
- legacy technology and operating constraints
- customer communications and support design
- workforce development, training and confidence
- procurement, supplier management and accountability

This curve should not be treated as a failure of speed. In regulated environments, some friction is protective. The question is whether friction is purposeful, transparent and human-centred, or whether it becomes a hidden barrier that prevents meaningful adoption.

2.3 Curve 3, Human Psychological Adaptation

The Human Psychological Adaptation Curve represents the pace at which individuals can understand, trust and integrate new systems into everyday behaviour. It rises more slowly and less evenly because people adapt through experience, confidence, familiarity and context.

Human adaptation is influenced by digital confidence, age, health, stress, financial pressure, life events, accessibility needs, confidence with institutions and previous experience of technology.

This curve is important because AI adoption is not successful simply because a system is available. Adoption becomes meaningful when people can navigate it, recover from mistakes, seek help, understand what is happening and feel that their autonomy is preserved.

2.4 Where the gap appears

The Psychological Speed Gap appears where AI capability accelerates faster than institutional readiness and human adaptation. The gap is not a single fixed distance. It changes by product, channel, customer group, situation and level of support.

A service can look mature from a technical perspective while still being psychologically difficult for users to understand. Equally, a workforce may have access to powerful AI tools while lacking the protected time, practice and confidence to use them responsibly.

3. Workforce Development Capacity

3.1 Why workforce development belongs inside the framework

Workforce Development Capacity is a critical factor shaping the Institutional Adaptation Curve. Organisations do not adopt AI through policy and technology alone. They adopt it through the practical ability of their people to understand new tools, apply judgement, support customers and challenge outputs where needed.

Where learning time is squeezed by delivery pressure, adoption can become superficial. Awareness may increase, but deeper skill, confidence and judgement may remain concentrated among a small group of early adopters.

Key point

Workforce Development Capacity is not simply a training issue. It is a structural factor that determines whether AI adoption becomes meaningful, inclusive and sustainable.

3.2 The workforce catch

Many organisations face a practical catch. AI could help people work more effectively, but people need time to learn how to use it well. If there is no protected learning time, the workforce may be expected to adopt tools while also maintaining existing workloads, absorbing organisational change and responding to rising service demand.

This can produce a hidden capability gap inside the institution. The organisation may appear to be progressing because tools exist, pilots are running and policies are being written. However, confidence may remain uneven across teams, and the people closest to customer friction may have limited opportunity to shape how AI is implemented.

3.3 What good workforce development needs

- Protected time to practise with realistic tasks, not only awareness sessions.
- Clear boundaries on what AI should and should not be used for.
- Role-specific guidance for frontline teams, operational teams, managers, analysts and senior leaders.
- Safe spaces for challenge, uncertainty and learning from mistakes.
- Practical examples that connect AI use to customer outcomes, inclusion, risk and vulnerability.
- Feedback loops from employees who see customer confusion, escalation or harm in real journeys.
- Confidence-building support for people who are not early adopters.

3.4 Risk of shallow adoption

Signal	What it may indicate	Organisational response
AI tools are available but unevenly used.	Confidence and skill are concentrated among a small group.	Provide role-based practice and peer learning, not just generic tool access.
Teams use AI for low-risk drafting but avoid deeper use.	People may not understand boundaries, accountability or quality checks.	Clarify governance, review standards and escalation routes.
Employees accept AI output without challenge.	Over-trust may be replacing professional judgement.	Reinforce critical evaluation, explainability and human accountability.
Frontline teams report customer confusion but are not included in design.	Customer insight is being disconnected from technical implementation.	Bring frontline evidence into journey design, risk review and policy iteration.

4. SAFE-AI, the forces behind the gap

SAFE-AI explains why the Psychological Speed Gap occurs. It identifies four interacting forces: Speed of Technology, Human Adaptation, Institutional Friction and Ethical Response.

Figure 3. SAFE-AI Forces

Four forces explain why the Psychological Speed Gap appears and why it needs active governance.




Figure 3. SAFE-AI, four forces shaping the Psychological Speed Gap.

4.1 S, Speed of Technology

AI capability evolves quickly. New functions may be introduced before organisations have fully understood the behavioural, ethical, operational and customer implications of the last wave of change.

The design challenge is not to slow all innovation. It is to understand where speed creates risk because people, processes or controls cannot adapt at the same pace.

4.2 A, Human Adaptation

People adapt through repeated experience, explanation, successful outcomes and a sense of control. Human adaptation is therefore slower than technical deployment.

Customers, employees and decision-makers may all experience adaptation pressure. For customers, this may appear as hesitation, channel switching or reliance on support. For employees, it may appear as uncertainty, over-cautiousness, over-trust or uneven adoption.

4.3 F, Institutional Friction

Institutional friction includes regulation, governance, risk controls, ethical review, infrastructure, supplier management, operational complexity and workforce development. Some friction is essential because it protects people from harm and helps maintain trust.

The governance question is whether friction is designed as a purposeful safety mechanism or allowed to become a confusing obstacle that slows learning without increasing protection.

4.4 E, Ethical Response

Ethical Response is the organisational choice layer. It determines whether AI is deployed in ways that improve support, inclusion, fairness, fraud prevention and confidence, or whether it creates bias, opacity, exclusion, over-automation and loss of autonomy.


Ethical response must be operational, not decorative. It should be visible in design decisions, support pathways, oversight mechanisms, measurement, workforce training and customer outcomes.

5. The Human Adaptation Layer

The Human Adaptation Layer explains why the Psychological Speed Gap is not only technical or organisational. It is behavioural. People need to form mental models, decide whether systems can be trusted, assess risk, feel in control and believe they can complete tasks successfully.

Figure 4. Human Adaptation Layer

Behaviour changes when psychological drivers are intensified by real-life circumstances.



When amplification is high, hesitation, avoidance, repeated confirmation or impulsive action may signal context, not incapability.

Figure 4. Human Adaptation Layer, showing psychological drivers and situational amplifiers.

5.1 Psychological drivers

Five psychological drivers sit beneath the Human Psychological Adaptation Curve. They help explain why adaptation is gradual, uneven and sensitive to system design.

Driver	Definition	Risk when ignored	Design response
Cognitive Load	The mental effort required to process information and make decisions.	People may rush, avoid, disengage or rely on shortcuts.	Reduce complexity, chunk information and allow slower decision-making.
Trust Formation	The belief that a system is reliable, fair and predictable.	People may reject the system or over-rely on it without review.	Explain decisions, provide consistency and make oversight visible.
Perceived Control	The belief that a person can influence or manage what happens.	People may feel trapped, anxious or dependent on others.	Offer clear choices, recovery routes, human support and reversibility.
Risk Perception	A subjective interpretation of danger, uncertainty or consequence.	People may avoid services or become vulnerable to reassurance from others.	Make risk information clear, timely and proportionate.
Self-efficacy and Confidence	The belief that a person can complete a task successfully.	People may withdraw from digital services or avoid complex decisions.	Build confidence through clear feedback, successful steps and accessible help.

5.2 Situational amplifiers

Situational Amplifiers are real-life circumstances that intensify psychological pressure. They do not make a person permanently vulnerable. Instead, they may temporarily reduce a person's ability to process information, assess risk, maintain confidence or challenge external influence.

Amplifier	How it affects adaptation	Possible behavioural signal
Emotional Stress	Reduces processing capacity and increases emotional response to uncertainty.	Rushed decisions, avoidance, reassurance seeking.
Life Events and Disruption	Interrupts routines and forces unfamiliar decisions during strain.	Difficulty completing tasks that would usually feel manageable.
Time Pressure	Pushes people towards rapid heuristics and default choices.	Skipping information, impulsive responses, accepting prompts without review.
Financial Anxiety	Narrows attention and increases perceived consequence of errors.	Delaying, seeking reassurance, channel switching.
Health and Cognitive Factors	May affect memory, attention, comprehension or decision speed.	Repeated errors, difficulty following multi-step journeys.
Social Influence and External Pressure	Introduces urgency, authority cues or conflicting instructions.	Following third-party instructions, deferring decisions, unusual transaction behaviour.

5.3 Behaviour should be interpreted in context

A key strength of this framework is that it discourages simplistic interpretation. Hesitation does not automatically mean incapability. Confidence does not always mean understanding. Fast completion does not always mean informed consent. Repeated confirmation does not always mean poor service literacy. Behaviour is a signal, not a conclusion.

Behavioural principle

Interpret behaviour as evidence of possible friction, context or risk. Do not use behaviour alone to label a person, remove autonomy or make assumptions about capability.

6. HARMONIC, principles for responsible response

HARMONIC translates the framework into practical design and governance principles. It helps organisations close the Psychological Speed Gap by aligning AI adoption with human pace, adaptation capacity, risk recognition, oversight, outcomes, navigability, inclusion and confidence.

Figure 5. HARMONIC Principles

Eight principles for closing the gap through human-centred AI design and governance.



HARMONIC is the practical response to the Psychological Speed Gap.

Figure 5. HARMONIC principles, corrected and structured as eight distinct design commitments.

H, Human Pace: AI should be introduced at a pace that allows people to understand, test and trust new systems. Fast deployment without explanation or support can increase confusion and reduce confidence.

A, Adaptation Capacity: People adapt differently. Services should account for different levels of digital confidence, financial knowledge, cognitive load, emotional resilience and accessibility need.

R, Risk Recognition: Systems should recognise situations where customers may be confused, stressed, under pressure or at risk of fraud, coercion or harm. Recognition should trigger supportive intervention, not punitive treatment.

M, Meaningful Oversight: Human oversight must remain visible and actionable. It should include review, escalation, accountability and human judgement in complex or high-impact situations.

O, Outcomes Focus: AI success should be measured by customer outcomes as well as efficiency. Important measures include fairness, clarity, accessibility, confidence, recovery and trust.

N, Navigability: AI-enabled services should be understandable and recoverable. People should know what is happening, how to get help and how to correct mistakes.

I, Inclusive Design: Services should work for people with different abilities, confidence levels, circumstances, languages, devices and support needs.

C, Confidence Building: Successful AI adoption should leave people feeling secure, respected, supported and more confident, not merely processed more quickly.

6.1 HARMONIC as a design test

HARMONIC can be used as a design review before an AI-enabled service goes live. It can also be used during operational monitoring when support data, complaints, fraud signals or customer feedback suggest that the gap is widening.

Question	Evidence to look for
Have we paced the change so people can understand it?	Communications, user testing, phased rollout, support readiness.
Have we designed for different adaptation capacities?	Accessibility testing, inclusion review, vulnerable customer journey testing.
Can the system recognise risk without over-labelling people?	Risk signals, escalation criteria, safeguards against bias.
Is oversight meaningful rather than symbolic?	Review rights, human escalation, audit trails, accountable ownership.
Are outcomes measured beyond efficiency?	Customer confidence, complaints, errors, recovery, fairness and support data.
Can people navigate, pause, recover and get help?	Clear steps, undo routes, accessible support and channel choice.
Has inclusive design shaped the service from the start?	Diverse testing, frontline insight, accessibility standards and user research.
Does the journey build confidence?	User feedback, repeat use, reduced anxiety, improved understanding.

7. Behavioural interpretation guidance

The framework can be used to interpret behavioural signals in AI-enabled journeys. The aim is not to diagnose the individual. The aim is to understand where the service may be creating friction, where context may be amplifying risk and where support may need to be offered.

Observed behaviour	Possible interpretation	Supportive response	Caution
Pausing repeatedly or spending a long time on a screen.	Cognitive load, uncertainty, low confidence or fear of error.	Simplify instructions, offer examples, allow save and return, make help visible.	Do not assume the person lacks capacity or is being difficult.
Repeated confirmation questions.	Low perceived control, high consequence, unclear next steps.	Confirm what will happen, explain reversibility and offer human support.	Do not treat reassurance seeking as consent.
Rapid completion without reading key information.	Time pressure, over-trust, stress or external influence.	Introduce friction at high-risk points and make review moments clear.	Do not assume speed equals understanding.
Switching from digital to phone or branch support.	Trust uncertainty, accessibility need, complexity or context pressure.	Preserve channel choice and learn from friction points.	Do not treat channel switching as failure.
Abandoning a journey.	Overload, loss of confidence, unclear value or fear of consequence.	Offer recovery, reminders and a simpler route back.	Do not remove access because the person did not complete first time.
Following unusual third-party instructions.	Possible coercion, scam pressure or misplaced reassurance.	Trigger safety checks, pause high-risk actions and provide clear warnings.	Do not disclose sensitive signals to potentially coercive third parties.
Strong confidence in an AI-generated answer.	Possible over-reliance on automation.	Show uncertainty, explain limits and require review for high-impact decisions.	Do not let confidence replace verification.

Interpretation rule

The same behaviour can have different causes. Design should create safe routes for clarification, support and recovery before harm occurs.

8. Design and governance guidelines

The Psychological Speed Gap Framework can be turned into a practical design and governance standard. The following guidelines help organisations move from concept to implementation.

8.1 Design guidelines

1. Design for slower understanding at important moments. Do not assume the customer or employee has the same context as the system designer.
2. Make AI involvement clear where it affects decisions, recommendations, prioritisation or communication.
3. Use plain language explanations that describe what happened, why it matters and what the person can do next.
4. Provide visible human support, especially where decisions affect money, access, identity, fraud risk or wellbeing.
5. Build recovery into the journey. People should be able to correct mistakes, pause, save, escalate or change channel.
6. Use friction carefully. Add protective friction where risk is high, but avoid unnecessary friction that increases stress without improving safety.
7. Avoid over-personalisation that feels intrusive or removes user control.
8. Test journeys with people who have different levels of digital confidence, cognitive load, accessibility need and life circumstances.

8.2 Governance guidelines

9. Assess AI initiatives against all three curves: capability, institutional readiness and human psychological adaptation.
10. Include workforce development as an adoption control, not as a post-launch communication activity.
11. Review customer outcomes alongside efficiency, automation rate and cost reduction.
12. Ensure human oversight is specific, resourced and empowered to act.
13. Monitor behavioural signals for evidence of widening gaps, such as abandonment, escalation, complaints and repeated support contact.
14. Use frontline and support team evidence as part of model and journey governance.
15. Maintain clear accountability for automated or AI-assisted outcomes.
16. Create feedback loops that allow policy, design, training and technical controls to improve together.

8.3 Ethical safeguards

Responsible AI design should protect autonomy, dignity and access. Safeguards should be visible in the service, not only in policy documentation.

Safeguard	Purpose	Example
Transparency	Helps people understand what is happening.	Explain when automation has influenced an outcome or support route.
Choice	Preserves autonomy and avoids forced digital dependence.	Offer alternative channels for complex or high-risk journeys.
Recovery	Prevents small errors becoming harmful outcomes.	Allow corrections, reversals, pauses and human review.
Escalation	Ensures human judgement is available when complexity rises.	Route vulnerable, disputed or high-impact decisions to trained staff.
Fairness review	Reduces bias and exclusion.	Test outcomes across groups, circumstances and digital confidence levels.
Workforce confidence	Ensures staff can use and challenge AI appropriately.	Give teams time, examples and decision boundaries.

9. Applying the framework across the AI lifecycle

The framework is most useful when it is applied across the full lifecycle of AI adoption, from early idea to live monitoring.

Figure 6. Responsible Response Cycle

A practical operating cycle for teams designing or managing AI-enabled services.




Figure 6. Responsible Response Cycle for AI-enabled services.

Lifecycle stage	Psychological Speed Gap question	Practical actions
Idea and discovery	What human problem are we trying to solve, and could AI change the behaviour or trust dynamics?	Map users, contexts, vulnerabilities, support needs and potential unintended consequences.
Design	Where might cognitive load, trust, control or risk perception become difficult?	Prototype explanations, support routes, recovery options and human escalation.
Build and test	Can different users understand and navigate the journey under realistic pressure?	Test with diverse users, frontline teams and scenarios involving stress, urgency and uncertainty.
Governance review	Are controls, oversight and accountability strong enough for the level of impact?	Review HARMONIC evidence, risk controls, fairness, transparency and workforce readiness.
Launch	Have people, channels and support teams been prepared for the change?	Use phased rollout, clear communications, monitoring and staff guidance.
Live operation	Is the gap widening in real behaviour?	Monitor abandonment, escalation, complaints, confidence, support demand and harm indicators.
Learning and iteration	What is customer and workforce evidence telling us?	Update journeys, scripts, training, governance and technical controls based on evidence.

10. Measurement and maturity assessment

The Psychological Speed Gap cannot be measured through one metric alone. It should be assessed through a combination of customer experience, behavioural data, workforce readiness, outcome quality and governance evidence.

10.1 Suggested indicators

Indicator area	Example measures	What it may reveal
Customer understanding	Comprehension testing, explanation usefulness, support questions.	Whether people understand what the system is doing and what they should do next.
Confidence and trust	Feedback, repeat use, complaint themes, channel switching.	Whether AI-enabled journeys increase or reduce confidence.
Navigability	Completion rate, abandonment, time on step, error recovery, save and return.	Whether the journey is usable under real conditions.
Support demand	Contact rate, escalation volume, repeated contacts, assisted completion.	Where the system is creating friction or uncertainty.
Risk and vulnerability	Fraud interventions, coercion signals, vulnerable customer referrals, complaint severity.	Whether the journey recognises heightened risk and responds appropriately.
Workforce readiness	Training depth, confidence to challenge AI, adoption patterns, quality review.	Whether institutional adaptation is deep enough to sustain responsible use.
Governance quality	Audit trails, decision review, fairness testing, issue resolution.	Whether oversight is meaningful and evidence-based.

10.2 Maturity model

Level	Maturity state	Description
1	Reactive	AI is adopted mainly through technical capability. Human adaptation and customer trust are considered after issues appear.
2	Controlled	Basic governance exists, but behavioural insight, inclusion and workforce development are inconsistent.
3	Supported	Customer support, accessibility, explanations and workforce guidance are designed into key journeys.
4	Integrated	The framework is embedded into product design, risk review, service operations and customer outcome measurement.
5	Adaptive	The organisation continuously senses, interprets and responds to widening gaps, using customer and workforce evidence to improve AI responsibly.

10.3 Practical maturity questions

- Can we explain where the Psychological Speed Gap is most likely to appear in our AI roadmap?
- Do we know which customer journeys carry the greatest risk of confusion, exclusion or over-trust?
- Have we tested AI-enabled journeys under stress, urgency or low-confidence conditions?
- Do staff have enough time and guidance to use AI responsibly?
- Can customers reach a human when the system becomes difficult, sensitive or high-impact?
- Are we measuring confidence, recovery and fairness, not just speed and automation?
- Can we show how customer evidence changes design, governance and training?

11. Practical tools and checklists

11.1 Pre-launch Psychological Speed Gap assessment

Assessment question	Yes, No or Evidence
Have we identified where AI capability may move faster than customer understanding?	
Have we assessed whether the institution has the workforce capacity to support the change?	
Have we mapped psychological drivers, including cognitive load, trust, control, risk and confidence?	
Have we considered situational amplifiers such as stress, time pressure, financial anxiety and external influence?	
Have we designed clear explanations for automated or AI-assisted outcomes?	
Can people pause, recover, correct mistakes or access human support?	
Have we tested the journey with users of different digital confidence levels?	
Have we defined escalation triggers and human oversight responsibilities?	
Have we agreed outcome measures beyond efficiency and speed?	
Have we built a live feedback loop from customers, staff and risk events?	

11.2 HARMONIC design checklist

Principle	Checklist prompt
Human Pace	Is the change introduced with enough explanation, support and time for people to adapt?
Adaptation Capacity	Does the journey work for people with different confidence levels, circumstances and abilities?
Risk Recognition	Can the system identify heightened risk, confusion or possible external pressure?
Meaningful Oversight	Is there a named human review process with authority to intervene?
Outcomes Focus	Are we measuring fairness, clarity, accessibility, confidence and recovery?
Navigability	Can people understand where they are, what happens next and how to get help?
Inclusive Design	Have diverse users, frontline teams and accessibility needs shaped the service?
Confidence Building	Does the experience leave people feeling more secure, respected and capable?

11.3 Red flag questions

- Are we assuming customers understand AI-enabled decisions because the process is technically accurate?
- Are we using automation to reduce contact without understanding why people need contact?
- Are we removing friction where friction is actually protective?
- Are we adding friction that increases stress without reducing risk?
- Are we treating vulnerability as a fixed customer attribute rather than a dynamic interaction between person, context and system?
- Are we asking frontline staff to support AI adoption without giving them meaningful learning time?
- Are we measuring success through speed while ignoring trust, comprehension and recovery?

12. Example use cases

The following examples show how the framework can be applied. They are deliberately generic and should be adapted to the specific organisation, product, risk profile and customer population.

12.1 AI-assisted fraud warning

An organisation introduces AI to identify unusual payment behaviour and present customers with targeted warnings. The AI Capability Curve rises quickly because detection improves. However, customer trust and comprehension may lag behind if the warning is unclear, too technical or appears at a stressful moment.

Framework response:

- Use Human Pace by providing short, plain language warnings with a visible option to slow down.
- Use Risk Recognition to detect urgency, unusual behaviour and possible external coaching.
- Use Meaningful Oversight by allowing escalation to a trained human support route.
- Use Navigability by clearly explaining what the customer can do next and how to recover.
- Measure outcomes through avoided fraud, customer understanding and post-warning confidence.

12.2 AI-generated customer communication

A team uses generative AI to draft customer letters, emails or chat responses. The benefit is faster communication, but the risk is that wording may be technically correct while psychologically unhelpful, especially for customers under stress.

Framework response:

- Assess cognitive load and emotional tone, not only accuracy.
- Include human review for high-impact or sensitive communications.
- Train staff to edit AI outputs for clarity, empathy, legal accuracy and customer context.
- Monitor complaints, repeat contacts and confusion after communication changes.
- Build staff confidence to challenge AI drafts rather than accept them automatically.

12.3 AI support tool for frontline colleagues

A frontline team receives an AI assistant to suggest next best actions. The institutional benefit may be consistency and speed. However, adoption may remain shallow if colleagues do not have protected practice time or if they are uncertain about when to rely on the tool.

Framework response:

- Treat workforce development as a control, not a communications add-on.
- Give staff role-specific examples and safe practice scenarios.
- Clarify accountability and when human judgement must override the tool.
- Collect frontline evidence about customer friction and use it to improve design.
- Measure colleague confidence, challenge behaviour and quality outcomes.

13. Glossary

Term	Meaning
Psychological Speed Gap	The widening distance between AI speed and the slower pace of institutional and human adaptation.
AI Capability Curve	The rate at which AI models, tools and automation capabilities improve.
Institutional Adaptation Curve	The rate at which organisations can govern, implement, support and embed AI responsibly.
Human Psychological Adaptation Curve	The rate at which people build understanding, trust, confidence and new behaviours around technology.
Workforce Development Capacity	The organisation's real ability to give people time, support and practice to use AI meaningfully and responsibly.
SAFE-AI	A model explaining the forces behind the gap: Speed of Technology, Human Adaptation, Institutional Friction and Ethical Response.
Psychological Drivers	Core mechanisms that shape adaptation: cognitive load, trust, control, risk perception and confidence.
Situational Amplifiers	Real-life pressures that intensify the drivers, such as stress, life events, time pressure and external influence.
HARMONIC	A responsible response model: Human Pace, Adaptation Capacity, Risk Recognition, Meaningful Oversight, Outcomes Focus, Navigability, Inclusive Design and Confidence Building.
Navigability	The ability of a person to understand a journey, know what happens next, get help and recover from mistakes.

14. Closing summary

The Psychological Speed Gap Framework offers a human-centred way to understand one of the defining challenges of the AI era. AI capability can accelerate quickly, but institutional readiness and human psychological adaptation move more slowly. When these speeds diverge, the gap becomes a site of confusion, reduced trust, exclusion, support pressure and potential harm.

The Three Curves Model helps organisations see the gap. SAFE-AI explains the forces that create it. The Human Adaptation Layer explains why people respond differently depending on cognitive load, trust, control, perceived risk, confidence and life context. HARMONIC provides a practical response by guiding AI design towards human pace, inclusion, navigability, oversight and confidence-building.

The framework does not argue against innovation. It argues for innovation that remains humanly navigable. The organisations most likely to succeed with AI will not be those that move fastest at every point. They will be those that understand when to move quickly, when to slow down, when to explain, when to support and when to bring human judgement back into the process.

Final thought

The real test of responsible AI is not only what systems can do. It is whether people can understand, trust and navigate them without losing confidence, control or dignity.

Appendix A. One-page framework summary

This page can be used as a quick reference for workshops, governance discussions or design reviews.

Component	One-line use
Psychological Speed Gap	Names the core problem: AI moves faster than people and institutions can adapt.
Three Curves	Shows where capability, readiness and confidence diverge.
SAFE-AI	Explains the forces causing the gap.
Human Adaptation Layer	Explains behaviour through psychological drivers and contextual amplifiers.
HARMONIC	Turns insight into design and governance principles.
Responsible Response Cycle	Creates a live operating rhythm: Sense, Interpret, Support, Escalate and Learn.

Workshop prompt: Choose one AI-enabled journey. Map the three curves, identify SAFE-AI pressures, list likely psychological drivers and situational amplifiers, then review the journey against HARMONIC.

Appendix B. Workshop exercise

17. Select a live or planned AI use case.
18. Describe the customer or colleague journey in plain language.
19. Identify where AI capability is ahead of institutional readiness.
20. Identify where customer or colleague adaptation may be slow or uneven.
21. Map the psychological drivers most likely to be activated.
22. Map the situational amplifiers most likely to intensify pressure.
23. Apply each HARMONIC principle and identify gaps.
24. Define three design changes, three governance actions and three measurement indicators.
25. Agree an owner for each action and a review point after launch.