

4:47

— A LEJACIE PAPER

Safe to Say, Safe to Reach, Safe to Use

An Emerging Framework for Closing the Paper-to-Person
Gap in Regulated Expertise so it can reach, matter to and
safely help real people

Is it safe to say?
Is it safe to reach?
Is it safe to use?

Chami Rupasinghe

Admitted Australian lawyer (non-practising)
Founder, Lejacie

Draft v0.1

30 September 2026
lejacie.com · chami@lejacie.com

Contents

Version and how to cite	4
Abstract	5
1. Introduction: problem, consequences, purpose and roadmap	6
1.1 The Paper-to-Person Gap	6
1.2 The actual problem	7
1.3 The conditions are not exceptional	12
1.4 The professional dilemma	13
1.5 What the gap costs	14
1.6 What a person needs	14
1.7 How the framework fits together	15
2. Definitions, audience and scope	16
2.1 Core terminology	16
2.2 Audience and scope	21
3. Research synthesis, conceptual foundations and proposed contribution	31
3.1 Research fields informing the framework	31
3.2 How the fields connect	34
3.3 Understanding is necessary, but not the endpoint	35
3.4 Reach involves encounter, relevance and relationships	35
3.5 Use depends on the surrounding pathway	36
3.6 Authority must remain intelligible and contestable	37
3.7 Knowledge travels in both directions	38
3.8 What is different about the proposed framework	39
3.9 Application sectors	39
3.10 Scope of the contribution	40
3.11 Status of additional sources	41
4. Framework development and evidence status	42
4.1 How the framework was developed	42
4.2 The framework is a synthesis, not a systematic review	44
4.3 What the existing literature establishes	44
4.4 What this paper proposes	45
4.5 What has not yet been established	46
4.6 Next stage: evaluation, co-design and revision	47
5. Three conditions and governing principles	49
Governing principles	62
6. Practical method	66
How the method works	69
0. Orient	71
1. Capture	73

2. Anchor	74
3. Bound	76
4. Reach	77
5. Translate	78
6. Connect	80
7. Verify	82
8. Test	83
9. Maintain	84
The minimum safe transfer	88
7. The 4:47pm Test	90
8. Worked example	98
9. Discussion, limitations and evaluation agenda	113
Limitations	113
Evaluation agenda	119
10. Conclusion and key points	122
Conclusion	122
Key points	122
11. References	125

Version and how to cite

Current version: Draft v0.1

Status: Proposed practice framework for discussion, testing and empirical evaluation

Date: 30 September 2026

Author: Chami Rupasinghe, Admitted Australian lawyer (non-practising), Founder, Lejacie

Title: Safe to Say, Safe to Reach, Safe to Use: An Emerging Framework for Closing the Paper-to-Person Gap in Regulated Expertise so it can reach, matter to and safely help real people

This document is an evolving working paper. It presents a proposed framework, practical method and stress test for examining how consequential expertise travels from its source to the people expected to understand, judge or act on it.

The framework is not a substitute for legal, clinical, financial, professional or other specialist advice; formal decision-making; professional standards; regulatory compliance; accessibility assessment; service design; or AI governance and assurance. It is intended to help professionals and institutions identify how those responsibilities connect across a real-world pathway.

The version number should change when the framework, method, definitions, standards, tests or claims change materially. Minor corrections, formatting changes and reference updates may be recorded as point releases, for example, v0.1.1.

Suggested citation

Rupasinghe, C. (2026). Safe to Say, Safe to Reach, Safe to Use: An Emerging Framework for Closing the Paper-to-Person Gap in Regulated Expertise so it can reach, matter to and safely help real people (Draft v0.1). Lejacie.

Short-form citation

Safe to Say, Safe to Reach, Safe to Use (Draft v0.1, 2026).

Suggested attribution

Adapted from Safe to Say, Safe to Reach, Safe to Use: An Emerging Framework for Closing the Paper-to-Person Gap in Regulated Expertise so it can reach, matter to and safely help real people (Draft v0.1, 2026), by Chami Rupasinghe, Lejacie.

Quoting and reuse

Please cite the version used. Because the framework is under development, readers should not assume that later versions retain the same wording, method, examples, boundaries or evidence claims.

Abstract

Law, policy, research and professional expertise can be correct at their source and still fail the person they were meant to help. When expertise leaves the expert, institution or system that holds it, the conclusion may travel without the source, conditions, uncertainty, authority, relevance, hidden questions, practical pathway and route back to qualified help that a person needs to use it safely. This is the Paper-to-Person Gap: the distance between regulated expertise being correct enough to issue and making a safe, useful difference in real life.

This paper introduces Safe to Say, Safe to Reach, Safe to Use, a proposed framework for addressing that gap. The framework includes a three-part standard, a practical ten-step method and the 4:47pm Test for stress-testing whether expertise works under constrained real-world conditions. Safe to Say asks whether expertise is accurate, properly sourced, qualified, bounded and within authority. Safe to Reach asks whether it can reach the people who need it and give them a truthful reason to pay attention, without manufacturing fear, urgency, certainty or pressure. Safe to Use asks whether enough of the expert's reasoning, hidden questions, limits, pathway, handover and recovery route has travelled for a person to understand what matters, judge what they can rely on, take a feasible next step, seek help, question an error and recover when the first route fails.

The framework is intended for professionals and institutions responsible for consequential expertise: expertise that can affect rights, safety, health, finances, access to services or future choices. Its contribution is to integrate three responsibilities that are often handled separately: preserving what regulated expertise is entitled to say; ensuring it reaches the people who need it in a form that matters; and making it usable enough to support real-world judgement and action. Reach and use may never exceed what the source and carrier are entitled to support.

Safe to Say, Safe to Reach, Safe to Use is a proposed framework that requires empirical evaluation across regulated settings. Its claim is that institutions should be able to test whether consequential expertise is safe to say, safe to reach and safe to use, and identify when any remaining failure belongs to the system rather than the person.

Introduction: problem, consequences, purpose and roadmap

1.1 The Paper-to-Person Gap

Imagine receiving a letter from a court, government department, insurer, employer, school, lawyer, hospital or regulator. It says something important has happened: a decision, deadline, right, risk, entitlement, obligation or process may affect you.

It says:

You may be eligible. Your circumstances may affect the outcome. We cannot guarantee the result. You should seek independent advice. Time limits apply.

The letter may be completely correct.

That is the problem.

Across law, policy, regulation, research and professional practice, consequential expertise may be accurate, complete, published or formally issued and still fail in the life of the person it was meant to help.

A person may not be able to:

- Find it before it is too late.
- See that it may matter to them.
- Recognise that a right, risk, choice, deadline or route to help is involved.
- Understand what is happening in practical terms.
- Work out what they can rely on and what they should not assume.
- Identify who decides and what remains uncertain.
- See the questions they need answered before they know to ask them.
- Take a safe, realistic next step.
- Follow through when time, stress, competing responsibilities or a practical barrier intervenes.
- Reach help.
- Correct, challenge, review, complain, escalate or recover if the first route fails.

The gap opens when expertise is organised according to the logic of the profession, while the person is left to turn it into the logic of their own situation, judgement and action.

An expert can often see automatically which facts matter, what exceptions may apply, how urgent a matter is, who decides, how much confidence an answer deserves, what questions need asking and what to do if the first route fails.

A person outside that field cannot silently supply that reasoning.

People also do not receive consequential expertise in ideal conditions. They may be tired, frightened, angry, grieving, rushed, using a phone, reading in a second language, short of money, missing documents, concerned about privacy or unable to reach a service. The professional helping them may have limited time, incomplete information or limited authority.

The problem is not that people are careless or incapable. The problem is that institutions often treat publication, legal sign-off, a referral or a document being sent as evidence that expertise has transferred.

It has not.

Where the information fails	What the person experiences
It cannot be found, accessed or reached in time	"I did not know this existed until it was too late."
It gives no reason to attend	"I did not realise this mattered until it was too late."
It feels generic or institutional	"I do not know whether this is about me."
It gives facts without a practical picture	"I read it, but I still do not know what is happening."
It assumes knowledge the person does not have	"I did not know I needed to ask that, bring that document or take that step."
It hides the limits of the answer	"I do not know what I can rely on."
It gives a next step that is not feasible	"I meant to deal with it, but I could not make it happen."
It ends with a disclaimer or referral	"I still do not know what to do, or where to go if that fails."
It offers no way to correct or challenge an error	"The system is wrong, but I do not know how to question it."

1.2 The actual problem

The problem is not:

"People do not understand complicated information."

The problem is not even:

"Expert information is poorly explained."

The deeper problem is that consequential expertise is often designed, approved and delivered as though accuracy at the source is enough.

It is not enough.

A statement can be legally correct, evidence-based, professionally responsible, technically accessible and formally approved, and still be unusable to the person whose rights, safety, health, finances, access to services or future choices may depend on it.

The person may receive the conclusion without the reasoning that makes it intelligible. They may receive the rule without knowing whether it applies. They may receive a warning without knowing whether the risk is immediate, likely, serious or relevant to them. They may receive a referral without knowing what the next service can actually do. They may receive an online form without the documents, time, privacy, money, language, confidence, connection or support required to complete it.

The failure is therefore not simply a failure of information.

It is a failure of **transfer**.

Expertise has not successfully travelled merely because it has been issued, published, uploaded, mailed, translated, explained once or referred onward. It has travelled safely only when the person who needs it can realistically:

- encounter it in time;
- recognise why it may matter;
- understand the relevant part of it;
- identify what is known, uncertain, conditional or outside the source's authority;
- judge what they can appropriately rely on;
- recognise the questions that matter in their circumstances;
- take a feasible next step;
- reach qualified help where the matter requires individual assessment, advice or decision-making;
- question, correct, review or challenge an error; and
- recover when the intended pathway does not work.

This is why the Paper-to-Person Gap is not solved by making a document shorter, friendlier or more visually appealing, although each may help.

The gap appears whenever the burden of interpretation, navigation, action and recovery is quietly transferred from the institution that holds expertise to the person least equipped or least resourced to carry it.

The transfer error

Institutions often use completion of an internal task as evidence that the external task has been completed.

Institutional completion	What may still be unresolved for the person
The policy is approved	"How does this affect me?"
The legal content is correct	"What can I actually rely on?"
The notice has been sent	"Did I see it, understand it and receive it in time?"
The webpage is published	"Can I find it when I need it?"
The material is in plain language	"What should I do in my circumstances?"
The form is online	"Can I complete it on my phone, with my documents and in private?"
The referral has been given	"Will that service help me, and what happens if it cannot?"
The decision includes review rights	"Can I realistically ask for review before the deadline?"
The AI system gives an answer	"Is this general information, a recommendation or something I can safely rely on?"
The service process exists	"Can I get through it while managing work, care, money, safety, disability, language or stress?"

None of these institutional actions is unimportant. Each may be necessary. But none is sufficient evidence that a person has received what they need to make a safe judgement or take action.

The framework therefore shifts the unit of analysis from the document, message, tool or isolated interaction to the full pathway through which expertise must travel.

Source → carrier → person → judgement → action → help, review or recovery

The source may be a statute, policy, professional standard, research finding, clinical guideline, official decision, risk assessment or expert judgement.

The carrier may be a lawyer, practitioner, frontline worker, court, regulator, government service, website, form, interpreter, community organisation, public campaign, automated notice or AI-supported system.

The person may be an individual, family, patient, client, consumer, worker, student, community member, small organisation or another professional who must act on the expertise.

At every point, something can be lost, distorted, overclaimed, hidden, delayed, made inaccessible or left without a workable route forward.

The missing work

The work that often goes missing is not simply "explaining better."

It is the work of carrying forward the elements that enable a person to use expertise safely:

What must travel	Why it matters
Source	A person needs to know where the information, rule, finding or decision comes from and whether it is official, current and applicable.
Authority	A person needs to know who is entitled to explain, advise, assess, decide, review or make a binding determination.
Meaning	A person needs a practical account of what is happening, rather than a conclusion alone.
Conditions	A person needs to see the criteria, facts, exceptions and dependencies that may affect whether the expertise applies.
Uncertainty	A person needs to know what is not known, what may change and what cannot be concluded from the available information.
Relevance	A person needs a truthful reason to recognise why the expertise may matter to them.
Hidden questions	A person needs help identifying the facts, documents, choices or questions that an expert would know to consider.
Boundary	A person needs to distinguish general information from personalised advice, assessment, prediction or decision-making.
Action	A person needs a next step that is specific enough to be feasible in their circumstances.
Handover	A person needs to know who can take over when individual judgement, representation, assessment or decision-making is required.
Recovery	A person needs a route to correct an error, seek review, complain, escalate, re-enter the pathway or obtain alternative help.

The framework does not require every resource or carrier to provide all of these elements in full. A brief public notice, for example, cannot contain the reasoning of a complete legal assessment or clinical consultation.

It does require each carrier to be clear about what it can safely carry, what it cannot determine, what conditions or limits matter, and how a person can move to the next appropriate source of help.

The accountability problem

The Paper-to-Person Gap persists partly because the relevant work is distributed.

One team owns accuracy. Another owns publication. Another owns accessibility. Another owns service delivery. Another owns digital systems. Another owns complaints. Another owns AI. Another owns risk. The person experiences the combined system, but the organisation may assess each part separately.

This creates predictable gaps:

- A legally accurate resource may be difficult to find.
- An accessible webpage may point to a service with no capacity.

- A clear form may be impossible to complete without documents the person does not have.
- A timely warning may create fear without showing a realistic protective action.
- A well-designed referral may fail if the next service does not accept the person.
- An AI assistant may retrieve accurate information but blur the line between explanation and personal advice.
- A review process may exist formally but be inaccessible within the relevant deadline.
- A disclaimer may protect a professional boundary while leaving the person with no usable pathway.

The framework treats these not as isolated communication failures, user failures or frontline failures. They are often failures of system design, ownership, coordination and accountability.

Its question is therefore not merely:

“Is this information accurate?”

It is:

“What would have to be true for a person, in realistic conditions, to use this expertise safely, and who is responsible for making those conditions possible?”

From publication to practical effect

A system has not completed its public-facing responsibility when it has made expertise available. It has completed more of that responsibility when the expertise can support a person’s informed judgement and realistic action without:

- overstating what the source supports;
- concealing conditions, uncertainty or authority limits;
- manufacturing fear, urgency, certainty or pressure;
- asking a person to infer expert reasoning they do not have;
- shifting system burdens onto people without acknowledging them;
- presenting an unavailable pathway as a solution; or
- using a disclaimer, referral or automated response as a substitute for a workable handover.

This does not mean that every person will obtain the outcome they want, that every system barrier can be removed or that every form of expertise can be simplified without loss.

It means that institutions should be able to identify where the pathway breaks, distinguish a genuine limit from a preventable design failure, and avoid treating a person’s inability to proceed as evidence that they did not try hard enough.

The next sections set out the proposed framework for doing that: **Safe to Say, Safe to Reach** and **Safe to Use**.

1.3 The conditions are not exceptional

The Paper-to-Person Gap does not affect everyone in the same way. But in the public-serving systems where people seek legal, health, financial, social and regulatory help, the conditions that make expertise hard to use are often not unusual cases at the edge of the system. They are the ordinary conditions in which people arrive.

People may be trying to use important expertise while they are distressed, short of money, caring for children or other family members, experiencing violence, coercion or unsafe monitoring, using English as an additional language, living with disability, chronic illness or mental ill-health, without stable phone or internet access, without privacy, unfamiliar with the system, geographically isolated or unable to pay for specialist advice.

This is not a list of exceptional user needs to accommodate once the “standard” resource is finished. It is the reality the resource, pathway and professional transfer should be designed around from the beginning.

In Victoria, community legal centres supported 54,475 people through the Community Legal Services Program in 2024–25. Of those clients, 61% were experiencing financial disadvantage, 54% were experiencing family violence, 27% were single parents, 27% had disability and/or mental illness, 18% were from culturally and linguistically diverse backgrounds, and 8% were homeless or at risk of homelessness.¹

Victoria Legal Aid similarly reports that many of its clients experience cultural barriers, discrimination, disability, mental-health issues or social and geographic isolation. In 2024–25, one in four clients had no income, 36% disclosed a disability or mental-health issue, one in four lived in regional Victoria, and 6% required an interpreter.²

The design implication is straightforward:

Do not build for the calm, confident, well-resourced person with time, privacy, complete documents, stable internet and ready access to advice, then call everyone else a special case.

Build for the conditions in which people actually need help.

This means assuming, and testing for, realities such as:

- The person may read on a phone, between other responsibilities.
- They may be scared, exhausted, ashamed, angry, grieving or distrustful.
- They may not have the document the system expects.
- They may need an interpreter, accessible format, advocate, support person or trusted intermediary.
- They may be unable to call during business hours or wait on hold.
- They may not be able to safely receive mail, notifications or calls.
- They may not know the system’s terms, categories, deadlines or usual questions.
- They may not have money for individual advice.
- They may have good reason, based on lived or historical experience, not to trust the institution.

Safe to Say, Safe to Reach, Safe to Use does not treat those conditions as proof that a person lacks capability, motivation or engagement.

It treats them as design conditions.

If the expertise, delivery route or pathway does not work under those conditions, the failure belongs at least partly to the system that designed it, not solely to the person who could not use it.

1.4 The professional dilemma

Professionals responsible for law, policy, research and high-stakes guidance often have to help a person without having all the facts, holding the authority to decide the outcome or being entitled to give individual advice.

They know that one missing condition, overlooked exception or overconfident statement can cause real harm. They may not be able to promise a result, tell a person which option to take, guarantee that a service will help or apply a general rule to that person's particular facts.

But the person still needs more than "it depends."

They need to know what this may mean, what facts matter, what they can safely rely on, what they can do now, who decides and where to go if the answer needs a qualified human.

That is the dilemma: the professional must not go too far, while the person cannot be left with too little.

Without a method, professionals often protect the boundary through caveats, disclaimers and referrals. Those may be necessary. But if they are not paired with practical meaning, a real handover and a path forward, caution becomes withdrawal.

Safe to Say, Safe to Reach, Safe to Use creates a third option: preserve the boundary, explain everything that can safely be explained and build the pathway at the point where expertise must be handed over.

1.5 What the gap costs

For people	For professionals	For institutions and systems
Missed rights, deadlines, reviews, payments, remedies or services	Repeat explanation, triage, repair work and moral distress	Avoidable complaints, delay, escalation and unequal outcomes
Unsafe or delayed action	Time spent compensating for unclear resources and broken handovers	Resources that are technically complete but fail in practice
Over-reliance on incomplete or unsuitable information	Pressure to give individual answers without authority or enough facts	Records that mistake system failure for “did not proceed” or “failed to respond”
Reliance on confident but unreliable AI, scam, influencer or misinformation sources	Defensive retreat into caveats, referrals and disclaimers	Loss of warranted trust and a growing advantage for false certainty
Carrying the work of navigating between services, explaining the same story and recovering from failed pathways	Burnout caused by repeatedly repairing preventable transfer failures	Reforms that exist on paper but do not change lived reality

1.6 What a person needs

Consequential expertise must support three connected human activities: **understanding, judgement and action**. These activities can overlap and recur. They are not a fixed sequence.

Activity	The person's question	What the resource and pathway must support
Understand	“What is happening, and why might it matter to me?”	Practical meaning, relevant context, roles, conditions and questions the person may not know to ask
Judge	“What can I rely on, and what should I not assume?”	Visible evidence, uncertainty, scope, authority and distinctions between general explanation, individual advice and formal determination
Act	“What can I realistically do?”	Rights, choices, feasible steps, reachable help, contestability and recovery when a route fails

Safe to Reach addresses whether the person encounters the expertise and recognises its relevance. Safe to Say preserves the meaning, evidence and authority on which understanding and judgement depend. Safe to Use supports these activities within the person's actual circumstances.

Action may include seeking help, preparing a question, declining an option, pausing or challenging a decision. Successful use is not equivalent to agreement or compliance.

1.7 How the framework fits together

Safe to Say, Safe to Reach, Safe to Use is a framework for making regulated expertise accurate enough to stand behind, human enough to matter and practical enough to change what a real person can do.

Not:

- safer transfer only;
- better communication;
- better information;
- plain language;
- user-centred content;
- persuasion;
- compliance;
- knowledge translation alone.

But:

A framework for making regulated expertise work for the people who have to live with it.

That holds all three pieces:

- it must remain regulated and defensible;
- it must reach human attention and matter;
- it must be usable in action.

Layer	Name	What it is	The question it answers	What someone does with it
1. The problem	The Paper-to-Person Gap	The name for the failure: expertise can be correct at its source but still fail the person who needs it	"What is going wrong here?"	Diagnose what has been lost between the source and the person
2. The proposed standard	Safe to Say · Safe to Reach · Safe to Use	The three conditions that responsible regulated expertise must meet	"What would good look like?"	Check whether the expertise is trustworthy, reaches the person and can be used safely
3. The method	Orient → Capture → Anchor → Bound → Reach → Translate → Connect → Verify → Test → Maintain	The practical process for designing or rebuilding the resource, interaction or pathway	"How do we make it good?"	Work step by step from source material through to real-world use
4. The stress test	The 4:47pm Test	A reality check for difficult ordinary conditions	"Would this work when life is messy?"	Test whether it works for a tired, pressured person and a busy professional

Definitions, audience and scope

2.1 Core terminology

Expertise

In this paper, expertise means specialised knowledge, evidence, interpretation, assessment, judgement, guidance or decision-making that a person or institution holds because of training, experience, authority, research, professional role or access to information.

Expertise may be expressed through:

- Advice, recommendations or professional explanations.
- Formal decisions, notices, orders or determinations.
- Laws, policies, rules, standards and guidance.
- Research findings, evidence summaries, guidelines and practice resources.
- Risk assessments, warnings, eligibility decisions and compliance material.
- Forms, letters, portals, websites, scripts and public-information resources.
- Referral pathways, service navigation and handover processes.
- AI-supported search, retrieval, explanation, navigation or decision-support systems.

The term includes both the underlying content and the practical activity of carrying that content to another person.

A legal rule, for example, is not the same thing as a lawyer's advice about how the rule applies to a particular person. A clinical guideline is not the same as an individual diagnosis or treatment recommendation. A government webpage is not the same as a binding decision. The framework depends on keeping these differences visible.

Regulated expertise

Regulated expertise is expertise shaped by formal rules about what may be said, who may say it, how claims must be supported, what limits must be made visible, what decisions may be made or how people may reasonably rely on it.

The source of regulation may include:

Source of regulation or control	Examples of what it governs
Law and legislation	Rights, obligations, powers, eligibility, deadlines, remedies, offences, procedural requirements and review rights
Professional regulation	Licensing, scope of practice, ethical duties, competence, supervision, confidentiality, informed consent and disciplinary accountability
Institutional authority	Who may issue a decision, certify a fact, determine eligibility, investigate, approve, refuse, review or enforce
Evidence and research standards	Research quality, evidentiary support, uncertainty, clinical guidelines, peer review, reporting and claims about effectiveness
Administrative and procedural rules	Notice, reasons, hearings, applications, appeals, complaints, records, service eligibility and decision review
Privacy, safety and safeguarding duties	Information collection, use, sharing, consent, risk management, mandatory reporting and protective action
Financial, consumer and fiduciary duties	Disclosure, suitability, fair dealing, conflicts, product governance, consumer protection and reliance obligations
Accessibility and equality duties	Reasonable adjustments, non-discrimination, accessible communication, language access and equitable service delivery
Technology and AI governance	Reliability, transparency, human oversight, data governance, monitoring, accountability, security and escalation

Regulated expertise is broader than professional advice. It includes official information, professional explanations, research translation, public guidance, procedural material and systems that influence what a person understands or does.

Examples include:

- A court notice explaining a hearing date, order, deadline or review right.
- A government letter about benefits, eligibility, immigration, housing, taxation, licensing or compliance.
- A legal-information resource explaining a process, without advising on an individual matter.
- A clinical discharge instruction, consent explanation, public-health warning or referral pathway.
- A financial hardship notice, insurance claim decision, product disclosure or complaints process.
- A regulator's public guidance, enforcement notice or professional standard.
- A research summary translated into policy, practice or public information.
- A school, university or employer policy that affects access, discipline, safety or support.
- A public-facing chatbot that explains an official process or directs users to services.
- A frontline worker's explanation of what a person can do next within a service system.

Regulation does not necessarily make expertise trustworthy, useful or safe. It establishes duties, boundaries and accountability structures that shape the conditions under which expertise is produced and carried.

What counts as regulated expertise

In this paper, **regulated expertise** means knowledge, guidance, assessment, judgement or information that is shaped by formal duties about what may be said, who may say it, how it must be evidenced, what limits must be made visible or how it may be relied on.

It includes expertise that is governed by one or more of the following:

Source of regulation or control	What it can govern
Law and legislation	Rights, obligations, powers, eligibility, offences, remedies, deadlines, review rights and formal decision-making
Regulation and licensing	Who may provide services, give advice, assess risk, make decisions or represent another person
Professional standards and codes	Competence, ethics, confidentiality, record-keeping, supervision, informed consent and scope of practice
Evidence standards	How claims must be supported, how uncertainty must be represented and what conclusions can properly be drawn
Institutional authority	Who can approve, refuse, determine, review, certify, investigate or enforce
Privacy, safety and safeguarding duties	What information can be collected, shared, disclosed or acted on and how risks must be managed
Administrative and procedural rules	Notice requirements, hearing processes, review rights, complaint routes, deadlines and service eligibility
Contractual or fiduciary duties	Disclosure, fair dealing, suitability, conflicts, consumer protection and reliance obligations
AI and technology governance	Reliability, transparency, human oversight, monitoring, data governance, accountability and escalation

Regulated expertise is not limited to “advice.”

It includes:

- a court or tribunal notice;
- a government eligibility decision;
- a hospital discharge instruction;
- a public-health warning;
- a financial-services explanation;
- an insurance claim decision;
- a child-safety protocol;

- a professional standard;
- a research guideline;
- a regulator FAQ;
- a family-violence court pathway;
- a university research translation;
- a practitioner’s explanation;
- a referral pathway; and
- an AI-generated answer that appears to carry professional authority.

The key question is not:

“Is this document legally regulated?”

The key question is:

Could a person reasonably use, rely on or act on this expertise in a way that affects their rights, safety, health, finances, access to services or future choices?

If yes, the framework may be relevant.

Consequential expertise

Consequential expertise is expertise that can materially affect a person’s rights, safety, health, finances, access to services or future choices if it is misunderstood, unavailable, inaccurate, over-relied on, acted on too late or left without a workable route to help.

Consequential expertise is defined by its likely practical effect, not by the prestige of its source or the formality of its presentation.

It may include:

- Information about a legal deadline, review right, court order or obligation.
- Guidance about symptoms, treatment, medication, consent, care or safety.
- A decision about income support, housing, insurance, education, employment or migration status.
- A warning about a product, public-health risk, natural hazard, environmental exposure or financial harm.
- A resource explaining how to seek protection, report harm, make a complaint or obtain urgent assistance.
- Information about eligibility for a service, funding, compensation, accommodation or support.
- A professional or AI-supported explanation that may influence a person’s next decision.

An apparently ordinary document can be consequential because of the setting in which it is received. A short message about an appointment may determine whether a person attends a hearing, keeps a benefit, receives treatment, avoids a penalty or obtains help before a deadline. Conversely, a formally regulated resource may be relatively low consequence in a particular context.

The framework therefore asks teams to assess consequence in context.

The relationship between the terms

Regulated expertise and consequential expertise overlap, but they are not identical.

	Regulated expertise	Consequential expertise
Primary question	What rules, authority, standards or duties shape this expertise?	What may happen to a person if this expertise is unavailable, misunderstood, wrong, over-relied on or unusable?
Focus	Source, authority, professional role, evidence, procedure and accountability	Real-world stakes, dependence, vulnerability, timing, practical options and foreseeable harm
Can exist without the other?	Yes. Some formally governed expertise may have limited practical effect in a particular use case	Yes. Some high-stakes information may be produced outside a formally regulated profession or institution
Typical risk	Expertise is carried beyond what the source, evidence or carrier's authority can support	A person cannot use, act on, challenge or recover from expertise that materially affects them
Framework concern	Is it safe to say?	Is it safe to reach and safe to use?

The terms can be expressed as two different lenses:

Regulated expertise = expertise shaped by formal duties and authority
Consequential expertise = expertise with material real-world effects

The framework is most clearly directed to the overlap:

Regulated $\cap$ Consequential

This is the zone where a failure in how expertise is communicated, carried, applied, accessed or handed over may affect a person's rights, safety, health, finances, access to services or future choices.

For example:

Example	Regulated?	Consequential?	Why the framework may matter
A court notice giving a hearing date and consequences of non-attendance	Yes	Yes	The notice must be authoritative, reach the person in time, be understandable and provide a practicable route to support or review
A hospital discharge instruction	Yes	Yes	The instruction must accurately reflect clinical advice, be usable under stress and provide a route for follow-up if symptoms worsen
A financial-services hardship pathway	Yes	Yes	The information must accurately state options, be understandable, provide feasible contact routes and not imply outcomes that depend on assessment
A public research summary about a low-risk historical topic	Often	Not necessarily	Accuracy and transparency still matter, but the full framework may be applied proportionately
A social-media post on a serious but unregulated health claim	Not necessarily	Potentially	The framework may help assess reach, reliance, uncertainty and escalation, even if the source is outside formal regulation
A workplace policy about a minor administrative preference	Often	Usually low	The framework may be relevant, but a lighter application is likely sufficient
A public-facing AI assistant explaining a government eligibility process	Yes, through the system it represents	Often yes	The system must distinguish official information from individual determination, show limits and provide access to human help or review

The framework therefore does not require teams to decide that every piece of information is either fully regulated and high stakes or entirely outside scope. It supports proportional use.

A short, low-consequence resource may need a light-touch check: accuracy, clear wording, appropriate source identification and a contact route.

A high-stakes decision letter, safety warning, legal process, clinical instruction, public-benefit pathway or AI-supported triage system may require deeper work across content, reach, accessibility, authority, service capacity, handovers, contestability and recovery.

2.2 Audience and scope

Who uses the framework

Safe to Say, Safe to Reach, Safe to Use is for professionals and institutions responsible for expertise that other people must understand, judge or act on.

It is particularly relevant where getting the expertise wrong, relying on it too far, failing to act or being unable to reach help could materially affect a person's rights, safety, health, finances, access to services or future choices.

The framework can be used by:

Who	What they may use it for
Lawyers, legal teams and legal-assistance services	Public legal information, client resources, notices, referral pathways, legal education, triage tools, self-help material and AI-supported legal-information services
Courts, tribunals and dispute-resolution services	Hearing notices, orders, court forms, process information, interpreter pathways, review information, self-represented litigant resources and support-service handovers
Government departments and public agencies	Eligibility decisions, review and appeal notices, public guidance, benefits information, compliance material, service pathways, forms, implementation resources and public-facing AI tools
Regulators and professional bodies	Regulatory guidance, enforcement notices, complaints material, reporting requirements, standards, disciplinary information and professional-practice resources
Health, social-service and community organisations	Consent information, discharge instructions, referral pathways, safety information, care navigation, public guidance and decision-support resources
Financial services, insurance and consumer services	General information, claims material, hardship pathways, dispute information, product explanations, complaints routes and escalation resources
Researchers, universities and evidence bodies	Research summaries, evidence translations, practice guides, public reports, implementation resources and policy-facing material
Frontline teams and program leaders	Scripts, triage tools, handovers, service pathways, training, escalation protocols and resources used in time-pressured interactions
Service designers, legal designers, content designers and accessibility practitioners	Resource design, form redesign, pathway mapping, user testing, accessibility work and implementation projects
AI product, governance, risk and assurance teams	AI-supported public information, navigation tools, triage systems, conversational interfaces, retrieval tools, escalation design, carrier boundaries and monitoring plans

The framework may also be useful to funders, commissioners, educators, professional associations, community organisations and people with lived experience who are helping design, test or govern how expertise reaches the public.

It is not limited to one profession or sector. Its relevance depends on the **consequence of the expertise**, not the job title of the person producing it.

Intended audience

The framework is intended for people and institutions responsible for producing, approving, carrying, translating, designing, implementing or governing consequential expertise.

This includes people who create content, but it is not limited to them. The Paper-to-Person Gap is rarely solved by writers or designers alone. It requires cooperation among the people who own the source, control the channel, understand the intended audience, provide the service, hold the relevant authority and receive people when the first pathway fails.

Source owners and decision-makers

These roles determine, approve, interpret or hold responsibility for the underlying expertise.

They may include:

- Legislatures, policy owners, regulators and statutory authorities.
- Courts, tribunals, decision-makers and registry staff.
- Professional bodies, licensing authorities and standards setters.
- Senior lawyers, clinicians, financial professionals, researchers and subject-matter experts.
- Government departments, public agencies and local authorities.
- Universities, research institutes, evidence bodies and guideline developers.
- Boards, executives, risk owners and accountable officers.
- Product owners and model owners responsible for AI-supported systems.

Their role is not only to ensure that the source is correct. It is also to decide what can safely be carried beyond the source, which conditions and limits must travel with it, when individual assessment is required and what help, review or correction routes need to exist.

Carriers and intermediaries

Carriers are people, systems or organisations through which expertise moves from its source toward the person who may use it.

They may include:

- Lawyers, legal-information providers, paralegals and community legal workers.
- Health professionals, care navigators, social workers and community-service staff.
- Financial counsellors, claims officers, caseworkers and customer-service teams.
- Public servants, contact-centre staff, inspectors, regulators and compliance teams.
- Teachers, university staff, workplace advisers and human-resources practitioners.
- Translators, interpreters, accessibility specialists and communications teams.
- Community organisations, peer workers, advocates and trusted local intermediaries.
- Service designers, content designers, digital teams and implementation leads.
- Search tools, websites, online forms, portals, automated notices and AI-supported assistants.

A carrier may be authorised to provide individual advice or make a decision. Or they may be authorised only to explain general information, guide a person through a process, assist with access or make a referral.

The framework treats that distinction as central. The aim is not for every carrier to do everything. The aim is for each carrier to know what it can safely carry, what it cannot determine and how to help a person cross the boundary to qualified support.

System and pathway owners

System and pathway owners shape whether a person can act on the expertise once they receive it.

They may include:

- Service managers and operational leaders.
- Intake, triage, scheduling and referral teams.
- Digital-service owners and portal administrators.
- Complaints, review and escalation teams.
- Accessibility, language-access, privacy and safeguarding teams.
- Data, technology, cybersecurity and AI-governance teams.
- Procurement, commissioning, funding and partnership leads.
- Quality-assurance, evaluation and continuous-improvement teams.

These roles may not create the substantive content. However, they may determine whether the content is available in time, works on a phone, can be interpreted, can be acted upon, reaches the right service, receives a response or can be corrected when it is wrong.

People with lived experience

People who use, depend on or are affected by consequential expertise are not merely recipients of a finished product. They are essential participants in identifying where the Paper-to-Person Gap appears.

They may include:

- People directly affected by a decision, condition, rule or service pathway.
- Families, carers, support people and advocates.
- Community representatives and organisations.
- People with disability, chronic illness, language barriers, limited digital access or insecure housing.
- People experiencing financial hardship, violence, discrimination, legal problems, migration uncertainty or geographic isolation.
- Frontline professionals who see how people experience systems in practice.

The framework should be tested with intended users and people who encounter the relevant system under realistic conditions. This is not because all responsibility should shift to “user feedback.” It is because institutions cannot reliably infer what a person can understand, access, do or recover from without involving people who experience the pathway.

Governance, assurance and evaluation roles

The framework may also be used by people responsible for asking whether a system is working as intended.

They may include:

- Governance boards and senior accountable officers.
- Internal audit, risk, legal, privacy and compliance teams.
- Regulators, ombuds services, inspectors and quality bodies.
- Funders, commissioners and procurement teams.
- Researchers, evaluators and implementation partners.
- AI assurance, model-risk and responsible-technology teams.
- Consumer, patient, client and community advisory groups.

For these users, the framework offers a way to ask a cross-functional question that can otherwise fall between teams:

Is the expertise not only correct at its source, but safe to say, safe to reach and safe to use for the people whose decisions and outcomes it may affect?

What they apply it to

The framework can be applied to any resource, interaction, pathway or system through which regulated expertise has to travel from its source to another person.

This includes:

What is being worked on	Examples
Formal documents	Decision notices, orders, letters, contracts, compliance notices, discharge summaries, consent forms, policies and procedural directions
Public resources	Websites, FAQs, guides, fact sheets, explainer pages, portals, videos, translated material, chat tools and public campaigns
Professional resources	Practice guides, standards, protocols, scripts, checklists, training materials, decision aids, implementation tools and referral guides
Forms and processes	Applications, review requests, complaints forms, consent processes, reporting pathways, eligibility assessments and evidence-gathering steps
Service pathways	Referrals, warm handovers, triage processes, court-support pathways, escalation routes, review and appeal pathways, and interagency coordination
Live interactions	Front-counter explanations, telephone triage, workshops, consultations, community education, court support and service navigation
Research translation	Research summaries, reports, evidence briefings, policy recommendations, practice resources, guidelines and implementation materials
AI-supported systems	Public-facing chatbots, legal-information tools, clinical-information tools, navigation assistants, retrieval systems, decision-support interfaces and automated notices
System reform work	New legislation, policy implementation, service redesign, regulatory reform, digital transformation and new professional-practice models

The method may be applied to a single document or interaction. But it is most powerful when applied to the whole pathway around it:

Source → resource or carrier → person → judgement → action → help, review or recovery

A letter may be clear, but the pathway can still fail if the phone number does not work, the referral is not accepted, the next service has no capacity, the person cannot use the online form or no one owns the recovery route.

The framework therefore asks teams to apply it to the **resource and the pathway together** wherever possible.

Applications

The framework can be applied to a single item, a repeated interaction, a service pathway or a whole system.

Documents and communications

Examples include:

- Letters, notices, decisions, orders and reasons.
- Forms, applications, consent documents and declarations.
- Websites, portals, FAQs, fact sheets and explanatory videos.

- Public warnings, campaigns, media releases and social-media content.
- Policies, procedures, professional standards and practice guides.
- Research summaries, evidence briefings and implementation resources.
- Emails, SMS reminders, scripts, call-centre prompts and chatbot responses.
- Product disclosures, claim letters, hardship information and complaints guidance.

For a single item, the framework can test whether the content is accurate and bounded, whether it reaches the right people in time and whether it gives a person enough support to understand the practical meaning, act safely and obtain help where needed.

Decisions, processes and pathways

Examples include:

- Court, tribunal and administrative review processes.
- Applications for benefits, housing, education, migration, licences or funding.
- Clinical care, discharge, consent, referral and follow-up pathways.
- Financial hardship, insurance claims, consumer complaints and dispute-resolution processes.
- Child protection, family violence, safeguarding and crisis-response pathways.
- Regulatory reporting, investigation, enforcement and remediation processes.
- University, workplace, school and professional-discipline processes.
- Emergency warnings, public-health communications and disaster-recovery support.
- Public-service navigation and cross-agency referral pathways.

For a pathway, the framework examines the movement between stages. It asks whether the person understands what each stage means, knows who has authority, can meet practical requirements, reaches the next service successfully and has a recovery route if a handover fails.

Digital and AI-supported systems

The framework can be applied to:

- Search and retrieval systems.
- Information websites and interactive eligibility tools.
- Online forms and self-service portals.
- Automated letters, reminders and decision notifications.
- Chatbots, virtual assistants and AI-supported information tools.
- AI-supported triage, prioritisation, risk scoring and decision-support systems.
- Knowledge-management systems used by frontline staff.
- Translation, summarisation and content-generation tools used to carry expert information.

The framework is particularly relevant where a digital or AI-supported system appears authoritative, personalised or decisive to the person using it.

It asks:

- What source material does the system rely on?

- Is the content current, traceable and properly qualified?
- What is the system authorised to explain, recommend, assess or decide?
- Can a person distinguish general information from individual advice or a formal decision?
- Does the system make uncertainty and limits visible?
- Does it provide a realistic route to a qualified human where required?
- Can a person correct an error, contest an outcome or obtain review?
- Does it work for people with limited digital access, accessibility needs, language needs or safety concerns?
- Is there a process for monitoring harmful patterns, updating content and responding when the system fails?

The carrier matters

The same underlying expertise can be carried differently depending on who is delivering it.

Carrier	What they may commonly do	What may require a boundary or handover
Authorised professional	Explain, apply expertise to facts, advise within their scope and make recommendations where authorised	Formal decisions or specialised advice outside their role
Court, regulator or decision-maker	Issue a formal determination, explain official process and identify review rights	Independent advice, advocacy or safety planning
Frontline worker or community practitioner	Explain general process, name public pathways, provide generic examples, support access and make referrals	Applying rules to individual facts, predicting outcomes or recommending a personal course
Public website, form or guide	Explain general rules, conditions, pathways, deadlines and support routes	Assessing an individual position or determining what they should do
AI-supported system	Retrieve approved information, explain general material, guide users through public pathways and escalate	Personalised legal, clinical, financial or other regulated advice unless specifically designed, authorised and governed to do so
Peer worker, advocate or support person	Share lived experience, explain practical pathways, help a person prepare questions and support access	Making professional determinations or presenting personal experience as authoritative advice

The carrier's authority sets the ceiling for how far the expertise may be personalised.

The framework does not treat this as a reason to stop being useful.

It asks:

What can this carrier safely carry, and what must be handed over to someone with the authority to take the next part of the journey?

What the framework does not cover alone

The framework is not a substitute for:

- individual legal, clinical, financial, migration or other specialist advice;
- formal decision-making;
- risk assessment or safety planning for a specific person;
- therapy, crisis response or emergency intervention;
- regulatory compliance advice;
- audit, assurance or certification;
- a full accessibility assessment;
- a complete service-design process;
- a full AI governance system; or
- reform of an unjust rule, inadequate service or under-resourced system.

It can identify where these are needed.

It can also identify when a resource or pathway is asking a person to carry risks that belong to the institution, professional or system.

But it cannot make a missing service exist, make an unfair rule fair, guarantee a person's safety or create meaningful choice where none exists.

Where those conditions prevent use, the framework treats them as **system findings**, not failures of the person.

Boundaries and exclusions

The framework is a method for examining how expertise is carried. It is not itself a source of substantive expertise or a replacement for the duties that govern a particular field.

It does not:

- Give legal, clinical, financial, migration, tax, therapeutic or other specialist advice.
- Determine an individual's rights, diagnosis, eligibility, risk or best course of action.
- Replace professional judgement, informed consent, risk assessment or safety planning.
- Create authority where a carrier does not have authority.
- Replace legal compliance, professional regulation, ethical review, accessibility assessment, privacy assessment, clinical governance, service design or AI assurance.
- Guarantee that a pathway has sufficient funding, workforce capacity, service availability or lawful options.
- Resolve an unjust law, policy, decision rule or institutional practice merely by explaining it more clearly.
- Eliminate the need for advocacy, representation, structural reform or accountability when the underlying system is unfair.

The framework can, however, make these limits visible.

It can identify when an apparently communicative problem is actually a service-capacity problem, an authority problem, an accessibility problem, a policy problem, an evidence problem or an injustice in the underlying system.

It can also identify when an institution has placed unreasonable interpretive, practical or emotional labour on the person expected to use its expertise.

Proportionate use

The framework is intended to be applied proportionately.

Not every resource requires the same level of research, design, testing, governance or documentation. The depth of application should reflect the potential consequence of failure, the degree of uncertainty, the complexity of the pathway, the vulnerability or dependence of intended users, the authority claimed by the carrier and the availability of recovery routes.

A proportionate assessment may consider:

Consideration	Questions
Potential impact	Could failure affect rights, safety, health, finances, access to services or future choices?
Time sensitivity	Is there a deadline, acute risk, rapidly changing situation or narrow opportunity to act?
Complexity	Does the person need to understand conditions, exceptions, evidence, multiple steps or interactions between systems?
Authority and reliance	Will people reasonably treat the source, carrier or tool as official, expert, binding or personally applicable?
User circumstances	Are people likely to be stressed, unwell, time-poor, digitally excluded, unsafe, unfamiliar with the system or using a second language?
Pathway dependence	Does safe action depend on another service, referral, appointment, document, payment, interpreter, advocate or decision-maker?
Recoverability	Can a person correct an error, obtain review, appeal, complain, escalate or access another route if the first one fails?
Scale and repeatability	Could a flaw affect many people, recur automatically or be amplified through digital or AI-supported systems?

The framework is therefore not a demand for maximum process in every case. It is a way to match the level of care, testing and accountability to the consequence of getting the transfer wrong.

Research synthesis, conceptual foundations and proposed contribution

Safe to Say, Safe to Reach, Safe to Use is grounded in research and practice concerned with how consequential expertise becomes understandable, trustworthy, reachable and usable beyond the expert, profession or institution that holds it.

The framework does not treat responsible expertise transfer as a wholly new problem. Its component concerns are already addressed across legal design, access to justice, plain language, information behaviour, human-centred design, implementation science, health literacy, risk communication, procedural justice, safety science, accessibility, professional responsibility and AI governance.

Its proposed contribution is to connect selected responsibilities from these fields within one method for examining how consequential expertise travels across sources, carriers, people, services, institutional boundaries and recovery routes.

The synthesis is organised by **research field**. A separate sector map identifies settings in which the framework may be applied. This distinction matters: a field supplies conceptual, empirical or practical grounding; an application sector identifies where the framework may be useful. The inclusion of a sector does not establish that the framework has been tested or validated in that sector.

The synthesis is guided by one practical question:

What must travel with consequential expertise for a person to recognise its relevance, understand what matters, judge what they can appropriately rely on and take a feasible next step, without losing the conditions, uncertainty, authority and boundaries that make the expertise trustworthy?

3.1 Research fields informing the framework

The following table identifies the sources and fields that inform the framework. It does not suggest that every source supports every element of the framework, or that the fields are separate and exclusive. Several sources and concepts are relevant across more than one field.

Sources marked as **additional grounding** or **further grounding required** should be treated as proposed sources for examination and verification unless they have been fully reviewed and incorporated into the final synthesis.

Research field	Sources for the synthesis	What the framework draws from it
Legal design and people-centred justice	Hagan (2018; 2019); Chung and Kim (2022); Malik (2026); OECD (2021)	Participation, human-centred legal resources, justice journeys, practical access and designing legal systems around people's needs and experiences
Access to justice, legal capability and expertise	OECD and Open Society Foundations (2019); Pleasence and Balmer (2019); Balmer et al. (2024); Sandefur (2015)	Problem recognition, legal confidence, access to suitable help and the reasoning experts supply that non-experts may not possess
Plain language, information design and cognition	ISO 24495-1:2023; Assy (2011); Stoll et al. (2022); Sweller, van Merriënboer and Paas (2019)	Findability, comprehension, structure, cognitive demands, working-memory limits and the limits of simplifying consequential expertise
Information behaviour and everyday information seeking (additional grounding)	Wilson (1999; 2000); Savolainen (1995); Dervin (1998); Kuhlthau (1991); Lloyd (2010)	Recognising information needs, searching without expert vocabulary, sense-making, uncertainty, incidental encounter and use of informal intermediaries
Human-centred and service design	ISO 9241-210:2019; Design Council, Double Diamond; Hagan (2018; 2019)	Context research, problem definition, prototyping, user testing, journey mapping and examining a resource within its surrounding service pathway
Knowledge translation and implementation science	Graham et al. (2006); Field et al. (2014); Aarons, Hurlburt and Horwitz (2011); Moullin et al. (2019); Damschroder et al. (2022); Proctor et al. (2011)	Contextual adaptation, barriers and enablers, implementation feasibility, ownership, evaluation, monitoring and sustainment
Administrative burden	Moynihan, Herd and Harvey (2015); Halling and Baekgaard (2024)	Learning, compliance and psychological costs imposed by institutional processes; distinguishing communication problems from structural and operational barriers
Health literacy and checking understanding	Sørensen et al. (2012); Nutbeam and Lloyd (2021); AHRQ (2024); Ha Dinh et al. (2016); Talevski et al. (2020)	Access, understanding, appraisal and use; teach-back; and checking whether meaning has transferred rather than assuming that information delivery succeeded
Risk and uncertainty communication	Trevena et al. (2013); Bonner et al. (2021); van der Bles et al. (2019; 2020); Kerr et al. (2022)	Balanced risk information, probability, uncertainty, evidence limits, calibrated claims and warranted trust
Community engagement and public-interest communication (additional grounding)	WHO emergency risk-communication guidance (2017); WHO risk communication and community engagement competency framework (2021); Reynolds and Seeger (2005); Covello (2003)	Listening, trusted relationships, relevant channels, community feedback, proportionate urgency and communication that supports informed protective action

Research field	Sources for the synthesis	What the framework draws from it
Behaviour change and intervention evidence	Michie, van Stralen and West (2011); Mertens et al. (2022); Maier et al. (2022)	Capability, opportunity and motivation; barriers to follow-through; and caution about assuming that behavioural interventions reliably produce intended effects
Shared decision-making and decision support (additional grounding)	NICE decision-support standards; Stacey et al., Cochrane decision-aids review	Balanced options, benefits, harms, uncertainty and personal values where a meaningful choice exists
Procedural and administrative justice (additional grounding)	Tyler (2006); Lind and Tyler (1988); Blader and Tyler (2003); Mashaw (1983); Adler (2010)	Voice, respectful treatment, intelligible reasons, authority, institutional decision-making, contestability and meaningful correction or review
Human factors, safety and handovers (additional grounding)	AHRQ TeamSTEPPS Handoff; Starmer et al. (2014); Manser and Foster (2011); Hollnagel (2014); Resilient Health Care, Volume 2	Constrained conditions, workload, ordinary variation, transfer of responsibility, receiver acknowledgement, failed routes, adaptation and recovery
Trauma-informed practice and cultural safety (additional grounding)	SAMHSA (2014); Ramsden (2002); Curtis et al. (2019); Dudgeon, Milroy and Walker (2014); relevant Aboriginal and Torres Strait Islander-led sources	Safety, control, disclosure, power, historical experience, cultural context and avoiding interactions that cause further harm
Epistemic justice and co-production (additional grounding)	Fricker (2007); Dotson (2011); Beresford (2002); Ocloo et al. (2021); Oliver, Kothari and Mays (2019)	Whose knowledge counts, credibility, lived expertise, participation, institutional blind spots and the burdens or risks of co-production
Accessibility, language access and digital inclusion (additional grounding)	WCAG 2.2; Convention on the Rights of Persons with Disabilities, Articles 9 and 21; Karliner et al. (2007); Flores (2005); relevant local language-access guidance	Accessible formats, disability access, assistive technology, interpreting, multilingual communication, channel choice, device constraints and alternatives to digital pathways
Professional responsibility and information governance (further grounding required)	Relevant professional codes, legislation and jurisdiction-specific guidance; ISO 15489-1:2016	Authority, competence, confidentiality, scope, source provenance, version control, currency, records, correction and accountability
AI governance and appropriate reliance	NIST AI Risk Management Framework 1.0 (2023); Buçinca, Malaya and Gajos (2021); Vasconcelos et al. (2023); Schemmer et al. (2023); Dahl et al. (2024); Magesh et al. (2025)	Reliability, unsupported claims, human oversight, actual reliance, verification, escalation, monitoring and contestability

These fields overlap. For example, a person's ability to understand a decision may involve plain language, health literacy, legal capability, accessibility, procedural justice and trauma-informed practice at the same time. The table identifies how each field informs the proposed framework; it does not assign exclusive ownership of any concern to one discipline.

3.2 How the fields connect

The synthesis does not suggest that expertise travels through a simple sequence from source to person. The responsibilities are interdependent.

Authority, evidence and professional responsibility establish what can responsibly be carried. Information behaviour, communication and community-engagement research explain how expertise becomes findable, relevant and meaningful in the conditions of everyday life. Human-centred design, service design, knowledge translation and implementation research examine how expertise connects to feasible action and sustained use. Justice, safety, cultural, accessibility and governance perspectives establish how participation, boundaries, accountability, contestability and recovery should be protected.

The relationship can be expressed as follows:

Authority and evidence → truthful encounter and interpretation → practical action, help, contestability and recovery

The framework's three conditions hold these responsibilities together:

Framework condition	Fields most directly informing it	Central concern
Safe to Say	Professional responsibility; information governance; risk and uncertainty communication; procedural justice; AI governance	Whether expertise is accurate, traceable, qualified, bounded, current and within the authority of the source and carrier
Safe to Reach	Information behaviour; community engagement; public-interest communication; risk communication; accessibility; digital inclusion	Whether people can encounter expertise in time, through safe and usable channels, and recognise a truthful reason why it may matter
Safe to Use	Plain language; cognition; legal capability; health literacy; shared decision-making; service design; implementation science; handover practice; trauma-informed practice	Whether people can understand practical meaning, identify limits, take a feasible next step, obtain qualified help, challenge error and recover if a pathway fails

The same field may inform more than one condition. Accessibility, for example, affects whether a person can reach a resource and whether they can use the action pathway it describes. Procedural justice informs authority, explanation and review. Information behaviour informs both reach and the hidden questions needed for use. Human factors informs practical use but also tests whether expertise remains safely bounded under pressure or system failure.

The framework's proposed contribution is not that these fields have failed to recognise such connections. It is to make the connections operational in one shared standard, method and stress test for consequential expertise.

3.3 Understanding is necessary, but not the endpoint

Plain-language, information-design and cognitive research inform the framework's attention to wording, structure, findability and avoidable demands on the reader. Legal-capability, health-literacy and professional-expertise research extend that concern to the questions, reasoning and practical orientation a person needs in order to use information.

A person may understand every word in the statement:

"Your circumstances may affect the outcome."

Yet still be unable to identify which circumstances matter, what document to find, who makes the decision, what evidence is needed or what question to ask next.

The framework therefore treats understanding as more than sentence-level comprehension. It concerns a person's ability to form a practical picture of what is happening, identify what the expertise establishes and does not establish, recognise what remains uncertain and determine whether qualified assessment or advice is needed.

This distinction matters because clear language can still leave expert reasoning hidden. An expert may automatically recognise a relevant exception, missing document, deadline, procedural risk or question. A person unfamiliar with the system may not know that the issue exists, let alone know how to ask about it.

Teach-back and related approaches inform the framework's testing logic. Information being supplied is not evidence that its meaning has transferred. A more useful question is whether a person can explain, in their own words:

- What the information means in practical terms.
- What may or may not apply in their circumstances.
- What they can rely on.
- What remains uncertain or requires further assessment.
- What they would do next.

The purpose is not to test or blame the person. It is to test whether the resource, interaction or pathway has made the necessary meaning available.

3.4 Reach involves encounter, relevance and relationships

Information-behaviour research directs attention to the conditions that occur before comprehension. People may not recognise that they have an information need, know the institutional terms for their problem or begin with a question an expert would regard as complete. They may search in ordinary language, encounter information incidentally, avoid searching because the issue is stressful or confusing, or seek help first through family, peers, community workers, frontline staff, trusted intermediaries or AI-supported tools.

Safe to Reach therefore concerns more than whether information has been published, sent or made technically searchable. It asks whether the people who may need the expertise can realistically encounter it in time, recognise why it may matter and access it through channels they can safely and practically use.

Risk communication, community engagement and public-interest communication broaden this account. They direct attention to trust, uncertainty, listening, channel choice, community feedback and the competing explanations that may shape how people interpret an issue. Community engagement is particularly important because communication is not only the distribution of institutional knowledge. It involves understanding how people interpret a situation, what they fear, what they need to know, which messengers they trust and what practical barriers they face.

Within the framework, reach has two connected dimensions:

Dimension	Question
Practical reach	Can the intended person encounter and access the expertise in time, through a channel they can use?
Human reach	Can the person recognise why it may matter, assess whether it is relevant and identify a reason to attend to it?

Safe to Reach places an explicit ethical constraint on both. Relevance, urgency and engagement must remain proportionate to what the facts support. Attention is not itself success. The purpose is informed judgement and feasible action, not agreement, disclosure, institutional convenience or compliance obtained through fear, pressure or manufactured urgency.

3.5 Use depends on the surrounding pathway

Knowledge translation, implementation science and service design inform the framework's attention to context, operational arrangements, barriers, implementation and sustained use. Administrative-burden research adds a complementary question: what must people learn, provide, endure, coordinate or risk before a right, protection or service becomes practically available?

These foundations support examining the resource and the pathway together.

A correct explanation may direct a person to an unavailable service. An understandable form may require documents they cannot obtain. A website may describe an accessible process while requiring a long phone call during working hours. A professional may accurately state that they cannot advise on individual circumstances without making clear who can help, how to reach them or what to do if the first route fails.

Each team or organisation may complete its own task while leaving the person responsible for connecting the entire journey.

The framework therefore considers the pathway through which expertise must travel:

Source → carrier → person → judgement → action → help, review or recovery

This is not a claim that every carrier must resolve every problem. It is a requirement to make boundaries, dependencies and handovers visible.

Handover practice contributes an important distinction. A referral tells a person about another possible service. A handover concerns whether relevant information, authority and responsibility have actually passed to a receiving party that can accept the next stage of work. In clinical safety practice, handovers involve transfer of information together with authority and responsibility, with receiver acknowledgement, clarification and an opportunity for questions.³

The framework adapts this insight cautiously. It does not assume that clinical handover protocols or their outcome evidence transfer unchanged to legal, regulatory or public-service settings. It does, however, ask comparable questions:

- What information must travel to the next carrier?
- What remains uncertain?
- Who is authorised to take responsibility next?
- Can the receiving service accept the person?
- Does the person understand what will happen?
- What is the alternative if the transfer fails?

3.6 Authority must remain intelligible and contestable

Professional responsibility informs what a carrier may explain, advise, assess, recommend or determine. It establishes the importance of competence, scope, confidentiality, evidence, authority, accountability and correction.

Procedural and administrative justice broaden the issue from correct communication to the exercise of consequential authority.

A person may need to understand who has made, or will make, a decision; what criteria or evidence matter; what reasons support an outcome; what information they can provide; and how they can ask for correction, review, complaint or reconsideration. They may also need to be treated respectfully and be able to participate in a process without first mastering its professional language.

Clear wording cannot make an unfair process fair. Nor can an accessible interface compensate for a decision that is unlawful, discriminatory or unsupported by evidence. But an unintelligible, inaccessible or non-contestable pathway can prevent a person from participating in, responding to or challenging consequential institutional action.

Within the framework, authority should therefore be visible rather than assumed:

- What is the source of this expertise?
- Who is entitled to explain it?
- Who is entitled to give individual advice, make an assessment or issue a decision?
- What can the person appropriately rely on?
- What can they question, correct or challenge?
- What route exists if they believe the institution has made an error?

Shared decision-making and decision-support practice inform a related distinction. Some situations involve reasonable options, uncertainty and trade-offs that depend partly on a person's values and circumstances. Others involve professional assessment, institutional determination or legal obligation.

The framework should not manufacture choice where no lawful or practical choice exists. It should support meaningful choice where it does exist, and make authority, reasons, conditions, response rights and review routes visible where it does not.

3.7 Knowledge travels in both directions

The framework must not treat institutional expertise as complete and lived experience as merely feedback on presentation.

People affected by consequential systems hold knowledge about their circumstances, safety, cultural context, resources, priorities and the practical effects of institutional requirements. They may know why a route is inaccessible, why a deadline is unrealistic, why contact is unsafe, why a form cannot be completed or why an apparently sensible referral will fail.

Epistemic-justice and co-production perspectives sharpen this point. They ask whose testimony is treated as credible, whose vocabulary defines the problem, whose questions shape the resource and whether institutional categories make a person's experience difficult to recognise or explain.

These concerns matter at the beginning of the method, not only after a resource is designed. They affect:

- What is captured as the problem.
- Which source is treated as authoritative.
- What assumptions are made about a person's circumstances.
- Which questions become visible.
- How a service pathway is tested.
- What counts as evidence of success or failure.
- Who has authority to interpret a breakdown.

Trauma-informed practice and cultural safety further shape how the framework approaches disclosure, contact, participation and testing. Their relevance concerns power, control, safety, historical experience and the possibility of harm, not simply a warmer tone or more empathetic wording.

For Australian applications, cultural safety should be grounded in Aboriginal and Torres Strait Islander-led scholarship, community authority and context-specific expertise. It should not be reduced to generic cultural competence, translation or inclusion. The framework should ask not only whether people can access a pathway, but whether contact with the pathway is safe, respectful and capable of recognising their knowledge without requiring unsafe disclosure or institutional conformity.

3.8 What is different about the proposed framework

The proposed distinction lies in the organisation of responsibilities, rather than in the invention of its component ideas.

Safe to Say, Safe to Reach, Safe to Use offers the following proposed features.

Proposed feature	What it does
A shared three-condition standard	Examines accuracy and authority, truthful reach and practical use together. Success in one condition cannot compensate for a material failure in another.
An explicit account of what must travel	Treats the conclusion alone as insufficient where conditions, uncertainty, authority, relevance, hidden questions, reliance limits or a practical pathway are needed for safe use.
A carrier-specific method	Examines how the same expertise changes when carried by an authorised professional, institution, frontline worker, public resource, community intermediary or AI-supported system.
Connected boundaries and handovers	Treats the limit of general information as a trigger for examining qualified help, access conditions, responsibility and fallback, not merely for adding a disclaimer.
A shared real-conditions stress test	Uses the 4:47pm Test to examine whether the person, professional and system can still function under plausible constraints, including missed contact, limited time, missing information, stress and failure of the expected route.
A route from design findings to institutional accountability	Identifies when a problem cannot be repaired through wording or interface design and instead requires action by a service owner, policy owner, funder, decision-maker, governance function or another responsible institution.
A whole-pathway account of use	Treats practical use as extending beyond comprehension to action, qualified help, contestability and recovery.

These features may already appear in particular methods, sectors or applications. Their distinctiveness and practical value as an integrated framework remain questions for comparative literature review, co-design, application and empirical evaluation.

The proposed contribution is a repeatable way to hold connected responsibilities together across the whole pathway:

Preserve what expertise can responsibly support; enable it to reach the people who may need it; and make its meaning, limits and routes forward usable in real life.

3.9 Application sectors

The framework may be applied wherever expertise can materially affect rights, safety, health, finances, access to services or future choices.

These are potential application sectors. They are not sectors in which the framework has already been validated, and the inclusion of a sector does not mean that its existing legal, professional, ethical, accessibility or governance requirements can be replaced by the framework.

Application sector	Resources and pathways the framework could examine
Law and justice	Legal information, client resources, court and tribunal notices, forms, orders, advice boundaries, self-help resources, legal-assistance pathways, review processes and referral systems
Government and public administration	Eligibility guidance, decision notices, benefits applications, obligations, public-service portals, review and appeal pathways, complaints processes and cross-agency service navigation
Regulation and professional practice	Standards, compliance guidance, enforcement notices, licensing information, disciplinary processes, complaints systems, public warnings and professional-practice resources
Health and care	Consent information, discharge instructions, treatment or screening decision support, referrals, care transitions, public-health communications, care navigation and follow-up pathways
Community and social services	Family-violence support, housing, disability, child safety, homelessness, migration, income support, community navigation and interagency handovers
Finance, insurance and consumer services	Product information, disclosure documents, claims decisions, hardship assistance, complaints processes, dispute-resolution pathways, fraud warnings and financial-counselling referrals
Research and education	Evidence summaries, research translation, practice guidance, professional learning, institutional policies, student-support resources and research-to-practice pathways
Digital and AI-supported services	Public chatbots, navigation tools, triage interfaces, automated notices, online forms, retrieval tools and decision-support systems operating across the other sectors

The framework may be applied to an individual resource, such as a letter or web page. It may also be applied to a live interaction, a referral process, a service journey, an online portal, an AI-supported system or a broader policy-implementation pathway.

The level of application should be proportionate. A low-consequence resource may require a light-touch review of source, clarity, channel and contact route. A high-stakes decision, safety warning, clinical instruction, legal deadline, hardship pathway or AI-supported service may require deeper work on authority, access, personalisation boundaries, service capacity, handovers, contestability and recovery.

3.10 Scope of the contribution

The framework's purpose is not to collapse these research fields into one new discipline.

Its purpose is to provide a shared way for institutions to ask whether consequential expertise remains safe through the full path from source to person:

Source → carrier → person → judgement → action → help, review or recovery

At each stage, expertise can fail in a different way.

It may be inaccurate, unsupported or beyond authority. It may be technically correct but impossible to find. It may be found but not recognised as relevant. It may be understandable but not actionable. It may point to a service that cannot help. It may offer a referral without a real handover. It may rely on a digital channel that a person cannot safely use. It may be delivered in a way that requires unsafe disclosure, discounts lived knowledge or prevents a person from challenging an error.

The framework proposes that these failures should not be treated as separate communication, service, accessibility, operations or governance issues once they affect the same person's ability to use consequential expertise.

They are connected responsibilities.

The framework therefore asks institutions to preserve what expertise can responsibly support; enable it to reach the people who may need it; and make its meaning, limits and routes forward usable in real life.

3.11 Status of additional sources

Several sources listed in the table are newly identified grounding rather than confirmed components of the original synthesis. They should remain subject to review, verification and appropriate contextual use.

This is particularly important for:

- Sources whose relevance depends on a specific jurisdiction or professional setting.
- Clinical evidence proposed for adaptation to non-clinical settings.
- Cultural-safety sources, which should be led by relevant communities and applied with appropriate local expertise.
- Accessibility and language-access sources, which may require jurisdiction-specific legal and service standards.
- Professional responsibility and information-governance sources, which depend on the applicable profession, legislation and institutional role.
- Emerging literature that has not yet been verified.

Schuster and Scheu (2026), on communication of scientific uncertainty and trust in science, should remain outside the confirmed synthesis until its publication details, scope and findings have been verified.

Framework development and evidence status

Safe to Say, Safe to Reach, Safe to Use was developed as an interdisciplinary synthesis around a practical problem: consequential expertise can be accurate, approved and professionally responsible at its source, yet still fail the person expected to understand, judge or act on it.

The framework was not derived from a single discipline or presented as a new substitute for established fields. It developed by bringing together concerns that are often addressed separately: preserving the integrity and authority of expert knowledge; enabling people to encounter it in time and recognise why it may matter; and supporting practical understanding, informed judgement, action, qualified help and recovery when the expected route fails.

The initial synthesis drew from legal design and people-centred justice; access to justice, legal capability and professional expertise; plain language, information design and cognition; human-centred design; knowledge translation and implementation science; administrative burden; health literacy and teach-back; risk and uncertainty communication; behaviour change; and AI governance and appropriate reliance.

As the framework developed, further literature was identified as relevant to questions that its initial formulation raised: information behaviour and everyday information seeking; community engagement and public-interest communication; procedural and administrative justice; shared decision-making and decision support; human factors, safety science and handover practice; trauma-informed practice; cultural safety; epistemic justice and co-production; accessibility, language access and digital inclusion; and information governance.

These fields do not supply a single settled model. Nor do they each map neatly onto one part of the framework. They provide concepts, evidence, methods and cautions that inform different aspects of how expertise is carried, received, interpreted, acted upon, challenged and maintained.

The framework's proposition is therefore an integrative one:

Consequential expertise should be examined not only for whether it is accurate enough to issue, but also for whether it is safe to say, safe to reach and safe to use in the real conditions of the people whose rights, safety, health, finances, access to services or future choices may be affected.

4.1 How the framework was developed

The synthesis was organised around recurring responsibilities rather than around disciplinary boundaries.

Across the literature, different fields raised versions of the same practical questions:

- What is the source of the knowledge, rule, evidence, decision or recommendation?
- What does the source support, and what remains uncertain, conditional or outside scope?
- Who has authority to explain, advise, assess, recommend or decide?
- How can a person recognise that expertise may matter to them before it is too late?
- How can they find it without already knowing the system's language or structure?
- What practical meaning, hidden questions and relevant conditions must travel with the conclusion?
- What can a person appropriately rely on, and what requires individual assessment or qualified advice?
- What does a feasible next step look like under ordinary constraints?
- What happens when another service, professional, digital system or decision-maker must take over?
- How can a person correct an error, seek review, complain, escalate or recover when the intended route fails?
- Who owns the barrier when wording, publication or referral cannot repair it?

These questions became the framework's organising responsibilities.

The three conditions (Safe to Say, Safe to Reach and Safe to Use) were developed as a way to hold them together.

Condition	Responsibility brought together
Safe to Say	Source integrity, evidence, qualification, uncertainty, authority, professional boundaries, appropriate reliance, currency and correction
Safe to Reach	Timely encounter, discoverability, accessible channels, truthful relevance, proportionate urgency, trusted pathways, community context and safe contact
Safe to Use	Practical meaning, hidden questions, feasible action, informed judgement, authority boundaries, qualified handover, contestability and recovery

The ten-step method was developed to translate these conditions into practical work. It asks teams to identify the consequential decision or action; locate source and authority; define the claim and carrier boundary; identify intended users and real conditions; establish the moment of relevance; map the pathway and handovers; carry the minimum safe meaning; make action and escalation feasible; stress-test under constrained conditions; and assign accountability for maintenance and improvement.

The 4:47pm Test developed from a recurring concern across human-centred design, health literacy, administrative burden, trauma-informed practice, accessibility, service design and safety science: people do not encounter consequential expertise in ideal conditions.

They may be tired, rushed, distressed, on a phone, in a second language, without the documents they need, unable to contact the first service, concerned about privacy or safety, or unsure whether they have understood the issue correctly. The test asks whether expertise and its pathway remain usable when these ordinary conditions are present.

4.2 The framework is a synthesis, not a systematic review

This paper is a conceptual and practice synthesis. It does not claim to be a systematic review, meta-analysis, evidence standard or comprehensive account of every relevant discipline.

The source set was assembled to inform a framework for responsible expertise transfer across regulated and consequential settings. It was guided by relevance to the Paper-to-Person Gap rather than by a pre-registered search protocol, exhaustive database search or formal quality appraisal of all available evidence.

That means several limits should be made explicit.

First, the framework reflects the literature and practice traditions identified during its development. Other relevant traditions may not yet have been located or incorporated.

Second, the source set includes different kinds of authority: empirical studies, systematic reviews, conceptual scholarship, standards, professional guidance, implementation frameworks, policy documents and practice tools. These sources do not make the same kind of claim and should not be treated as equivalent evidence.

Third, the framework uses concepts developed in particular sectors, especially law, health, public administration, design, communication and AI governance. Their transfer across settings should be examined rather than assumed.

Fourth, the categories used in this paper are interpretive. The decision to organise responsibilities as Say, Reach and Use; to describe a Paper-to-Person Gap; to identify a minimum safe transfer; and to use a 4:47pm Test is the paper's own proposal. These are not labels that emerge unchanged from a single established body of literature.

The framework should therefore be read as a transparent synthesis and an invitation to test, refine or reject its organisation where evidence and experience warrant.

4.3 What the existing literature establishes

The literature provides substantial support for the importance of many individual responsibilities incorporated into the framework.

Plain-language, information-design and cognitive research support attention to clarity, structure, cognitive demand and comprehension. Health-literacy research supports the importance of people being able to access, understand, appraise and use information and services. Teach-back demonstrates the value of checking understanding rather than treating delivery as proof of comprehension.

Information-behaviour research supports the proposition that people do not always begin with a clear information need or the language necessary to search for an answer. They may encounter information incidentally, rely on social networks and intermediaries, avoid searching because of stress or uncertainty, or seek help through channels that institutions do not treat as official.

Legal design, people-centred justice and access-to-justice research support attention to the person's journey through legal systems, the practical capability required to recognise and respond to legal problems, and the gap between formal rights and usable access.

Knowledge translation, implementation science and service design support attention to context, barriers, adaptation, adoption, handovers, operational arrangements, monitoring and sustained use. Administrative-burden research supports attention to the learning, compliance and psychological costs that institutions impose on people in their efforts to access rights, benefits, services or remedies.

Risk communication, uncertainty communication and public-interest communication support calibrated communication of evidence, risk and uncertainty. Community-engagement approaches support listening, feedback, trusted relationships and the practical importance of communication routes that work in local context.

Professional responsibility, information governance and AI governance support attention to competence, scope, source traceability, authority, oversight, correction, accountability and the risk that people may rely on a source or system more than it can safely support.

Procedural and administrative justice support the importance of intelligible authority, reasons, voice, respectful treatment, correction and meaningful review. Shared decision-making supports the fair presentation of options, benefits, harms, uncertainty and personal values where a person has a meaningful choice. Human-factors and safety literature support attention to workload, variation, transitions, failure conditions and recovery. Trauma-informed practice, cultural safety and epistemic-justice scholarship support attention to power, safety, disclosure, credibility, historical distrust, cultural context and the limits of institutional knowledge.

Together, these bodies of work establish that the concerns taken up by this framework are real and consequential.

They do not, however, establish that the three conditions are the only or best way to organise those concerns; that the ten-step method is complete; that the 4:47pm Test is reliable; or that the proposed integration improves outcomes.

4.4 What this paper proposes

This paper makes six proposed contributions.

First, it names the Paper-to-Person Gap: the distance between expertise being correct enough to issue and becoming safe and useful in the life of the person expected to use it.

Second, it proposes that three conditions must be considered together:

- **Safe to Say:** whether expertise is accurate, traceable, qualified, bounded and within the authority of the source and carrier.
- **Safe to Reach:** whether it can reach the people who may need it in time, through usable channels, with a truthful reason to recognise why it may matter.
- **Safe to Use:** whether enough meaning, limits, hidden questions, next steps, handovers and recovery routes have travelled with the expertise for a person to understand, judge and act safely.

Third, it proposes that these conditions are jointly required. Accuracy cannot compensate for inaccessibility. Reach cannot compensate for an unsupported claim. Clarity cannot compensate for an unusable or unsafe pathway. A resource should not be treated as successful merely because

it has been approved, published, sent, translated, referred onward or made available online.

Fourth, it proposes a minimum safe transfer. The conclusion alone may not be enough. Depending on the setting, a person may also need the source, authority, conditions, uncertainty, relevance, hidden questions, reliance limits, next action, route to qualified help and recovery option.

Fifth, it proposes a carrier-specific method. The same expertise may be carried by an authorised professional, public authority, frontline worker, website, interpreter, community organisation, form, portal or AI-supported system. Each carrier has different authority, limitations and responsibilities. The framework asks what that carrier can safely explain, what it cannot determine and how a person should be connected to the next authorised source of help.

Sixth, it proposes the 4:47pm Test as a shared practical stress test. The test asks whether a person can still identify what matters, what they can rely on, what they can do next and how they can recover if they encounter the expertise under constrained conditions and the first expected route does not work.

These are propositions for use and evaluation. They should not be treated as settled findings.

4.5 What has not yet been established

The framework has not yet been empirically validated.

This paper does not establish that applying Safe to Say, Safe to Reach, Safe to Use will improve comprehension, trust, access, appropriate reliance, task completion, help-seeking, safety, equity, institutional accountability or substantive outcomes.

It does not establish that the three conditions are exhaustive, mutually distinct or reliably interpreted by different teams.

It does not establish that the ten-step method can be applied consistently across law, health, public administration, financial services, research translation, social services or AI-supported systems.

It does not establish that the 4:47pm Test predicts real-world failure, identifies problems missed by other methods or produces better outcomes when used in review and design.

It does not establish that a resource or pathway is fair, lawful, accessible, clinically safe, professionally compliant or ethically acceptable merely because it appears to meet the framework's requirements.

It does not establish that a person will obtain the outcome they seek, that a referred service will have capacity, that a review mechanism will provide a remedy or that a system's underlying rules are just.

These limits are particularly important where the framework is used in high-stakes settings. A well-designed communication cannot compensate for an unlawful rule, inaccessible service, inadequate workforce, unaffordable requirement, discriminatory practice, unavailable interpreter, unsafe disclosure requirement or absent route to specialist help.

The worked examples in this paper are illustrative. They show how the framework can organise questions and identify potential pathway failures. They are not intervention studies, outcome evaluations or evidence that the proposed method has improved people's experience or results.

4.6 Next stage: evaluation, co-design and revision

The next stage is to test whether the framework is useful, feasible and responsible in practice.

Evaluation should begin with real, consequential examples: a decision letter, review pathway, discharge instruction, public warning, benefit application, financial-hardship process, regulatory guidance, referral system or AI-supported information tool.

The first question is conceptual:

Does the framework reveal a problem that existing content, legal, accessibility, service-design, governance or operational review processes would not otherwise identify together?

The second question is practical:

Can the people responsible for the pathway use the method without creating unworkable delay, duplicated approval, excessive documentation or new burdens for frontline staff and intended users?

The third question is outcome-focused:

Does applying the framework improve a person's ability to recognise relevance, understand what matters, identify limits, take a feasible next step, reach qualified help, question an error and recover when the first route fails?

Evaluation should involve intended users, people with lived experience, frontline professionals, source owners, decision-makers and pathway owners. It should include people likely to face access barriers, including people with disability, language and interpreting needs, limited digital access, geographic isolation, financial hardship, safety concerns or prior negative experience of institutions.

Relevant evidence may include:

- Source, authority and currency reviews.
- User research, cognitive interviews and task-based testing.
- Teach-back or scenario-based assessment of practical understanding.
- Accessibility, language-access and channel testing.
- Analysis of search, navigation, completion, abandonment and repeat-contact patterns.

- Referral, handover, waiting-time, review, complaints, correction and escalation data.
- Frontline observations of recurring confusion, workarounds and failed pathways.
- Qualitative accounts of burden, safety, trust, dignity, agency and practical barriers.
- Comparison with existing practice where ethically and operationally feasible.
- Monitoring of unintended consequences, including delay, over-disclosure, excessive caution, reduced autonomy, new exclusion or inappropriate reliance.

Evaluation should also document adaptation. A framework intended for varied settings cannot be applied identically everywhere. But changes should be recorded clearly enough to distinguish the framework's core responsibilities from local implementation choices.

The framework should remain open to revision. Its value will depend not on the appeal of its language, but on whether it helps institutions identify preventable transfer failures, assign responsibility for them and improve the real conditions in which people must understand, judge and act on consequential expertise.

Three conditions and governing principles

The framework proposes a standard for consequential information: information that can materially affect a person's rights, obligations, safety, health, finances, access to services or future choices.

Could the person find it, see why it mattered, understand what was happening, see the questions they needed answered, tell what they could safely rely on and take a realistic next step when it mattered?

A resource fails if it is easy to read but drops an important condition, hides uncertainty, overstates evidence or blurs a professional boundary.

It also fails if it is accurate and easy to find but leaves the person unable to understand what matters, judge its limits, see a hidden question, take a feasible step, reach help or recover if the pathway fails.

What safe means

Safe does not mean risk-free, guaranteed to produce the correct outcome or a substitute for individual advice, professional judgement or a formal decision.

Safe means designed to support understanding, warranted reliance and feasible action while making material limits, uncertainty, authority and escalation needs visible.

Why all three conditions are needed

The standard has three conditions: Safe to Say, Safe to Reach and Safe to Use.

If you have only...	You may get...
Compliance / regulation	Accurate, qualified material that nobody understands, notices or uses
Marketing / cognition	Attention and action achieved through overconfidence, emotional pressure or a promise the expertise cannot support
Human-centred or legal design	A clear, useful experience that accidentally drops an exception, overstates the evidence or crosses into unauthorised advice

Each condition constrains the other two.

- Compliance without Reach can produce expertise that is technically correct but invisible, cold or impossible to enter.
- Reach without Say can become false certainty, emotional pressure or manipulation.

- Human-centred design without Say can accidentally delete an exception, overstate evidence or turn general expertise into unauthorised individual advice.
- Say and Reach without Use can leave a person interested, informed and still unable to do anything safely.

How the three conditions work together

Safe to Say, Safe to Reach, Safe to Use is a proposed framework for carrying consequential expertise from its source to the people who may need to understand, judge or act on it.

It is not a linear communications model. The three conditions must hold together. Expertise that is easy to find but inaccurate is unsafe. Expertise that is accurate but impossible to locate, understand or act on is insufficient. Expertise that is clear and engaging but exceeds the authority of its source or carrier is unsafe.

The framework therefore asks three connected questions:

Safe to Say + Safe to Reach + Safe to Use

Condition	Core question	Failure if absent
Safe to Say	Is this expertise accurate, properly sourced, qualified, bounded and within the authority of the source and carrier?	People may be misled, over-rely on content, receive unauthorised advice or act on claims that cannot be supported.
Safe to Reach	Can the expertise reach the people who may need it, in time and through usable channels, with a truthful reason to recognise why it may matter?	Expertise remains technically correct but unseen, ignored, received too late or encountered without sufficient relevance or trust.
Safe to Use	Has enough meaning, context, limitation, pathway and recovery route travelled with the expertise for a person to make a safe judgement and take a feasible next step?	People may understand fragments but remain unable to act, seek qualified help, question an error or recover when the intended route fails.

The framework is designed for shared use. A source owner may be responsible for accuracy and authority; a communications or community team may lead reach; a service, design or frontline team may lead usability; and a governance or quality function may oversee the whole pathway. No single team has to perform every task. However, someone must be accountable for whether the three conditions hold together.

Safe to Say

Can we stand behind it?

Safe to Say is the compliance, professional responsibility and authority part of the standard.

It asks whether the expertise is accurate, properly sourced, qualified, bounded and within the authority of the person, institution or system carrying it.

A law, policy, research finding, professional judgement or formal decision cannot safely travel if it loses the things that make it trustworthy:

- the source;
- the evidence;
- the condition;
- the exception;
- the deadline;
- the uncertainty;
- the authority structure;
- the ethical or privacy requirement; or
- the boundary between general explanation, individual advice and formal determination.

Safe to Say means that what reaches the person is still something the relevant source owner, professional or institution can stand behind.

It does not mean:

- say as little as possible;
- write in legalese;
- hide behind disclaimers;
- preserve every sentence exactly as written;
- remove all practical explanation because the answer depends on facts; or
- send a person away with “seek advice” and no usable pathway.

It means preserving what carries meaning while changing the form freely where the meaning survives.

Safe to Say does not mean “send everything that could lawfully be said.” What travels must be appropriate for its purpose, audience, delivery route, carrier and level of stakes. It must not omit material information in a way that misleads or strands a person, but it must also respect privacy, safety, relevance, procedural fairness and professional limits.

What Safe to Say tests

Question	What it checks
What is the authoritative source?	Law, policy, evidence, professional standard, formal decision or identified judgement
What exactly does it establish?	The claim that can safely be carried
What could change the answer?	Conditions, facts, thresholds, exceptions, missing information and time limits
What remains uncertain?	Evidence limits, disagreement, unknown facts or a decision not yet made
Who has authority?	Who can explain, advise, assess, decide, review or formally determine
How far can this carrier go?	What a lawyer, frontline worker, website, caseworker, educator or AI system may properly say
What must the resource not imply?	Eligibility, certainty, predicted outcome, individual advice or formal determination

The central boundary

Safe to Say has two duties that must be held together:

1. Do not say more than the expertise can support.
2. Do not say so little that the person is misled by omission or left without a usable path.

For example, this may be safe but insufficient:

“Your circumstances may affect the outcome. Seek independent advice.”

It may correctly signal that a resource cannot assess the person's position. But it does not tell them:

- what circumstances may matter;
- what they can use the resource for;
- what they should not assume it decides;
- when individual advice is needed;
- who can provide it;
- how to contact them; or
- what to do if that help is unavailable.

A disclaimer is not a boundary plan. The boundary must be designed into the pathway.

What Safe to Say produces

A source and boundary record that identifies:

- the relevant source material;

- material claims and their authority;
- conditions, exceptions, uncertainty and limits;
- the carrier's authority;
- where personalisation must stop;
- what may be translated freely;
- what must be preserved or reviewed; and
- where the resource must hand over to a qualified person or formal decision-maker.

What the condition concerns

Safe to Say concerns the integrity and authority of the expertise being carried.

It asks whether the content is accurate enough to issue, traceable to an appropriate source, current, appropriately qualified and expressed within the authority of the person or system carrying it.

The condition applies to more than factual accuracy. A statement can be factually correct while still being unsafe if it implies a conclusion that the evidence does not support, omits a decisive qualification, presents general information as personal advice or is delivered by a carrier without authority to make the claim.

Essential requirements

The following requirements are core. They should be met in every application, although the depth of evidence and documentation should be proportionate to consequence.

Requirement	Test
Traceable source	Can the material be linked to an authoritative source, evidence base, formal decision, approved policy or identified professional judgement?
Accuracy and currency	Is it factually accurate for the context, and has it been checked against current law, policy, evidence, guidance or service arrangements?
Appropriate qualification	Does it preserve relevant conditions, exceptions, uncertainty, time limits, eligibility criteria and limits of evidence?
Scope and authority	Is the source or carrier authorised and competent to make this statement, provide this explanation, offer this recommendation or make this decision?
Appropriate reliance	Is it clear what a person may use the content for, and what requires individual assessment, advice, representation or decision-making?
Review and correction	Is there an owner, review point and correction process if the source changes, the content is found to be wrong or the system produces harmful error?

What Safe to Say does not require

Safe to Say does not require every carrier to reproduce the full source, provide individual advice or resolve every uncertainty.

A public webpage may accurately explain general eligibility criteria without determining whether a particular person qualifies. A frontline worker may explain a process without predicting its outcome. An AI-supported assistant may retrieve approved information without being authorised to assess an individual's rights, diagnosis or risk.

The requirement is that the boundary remains visible and that the person is not invited to rely on the carrier beyond what it can safely support.

Safe to Say review questions

- What is the source of this expertise?
- Who owns its accuracy and currency?
- What does the source establish, and what does it not establish?
- What conditions, exceptions, uncertainty or evidence limits must remain visible?
- Who is authorised to explain, apply, assess, recommend or decide?
- Is the carrier giving general information, individual advice, a recommendation, an assessment or a formal decision?
- Could a reasonable person interpret this as more certain, personalised, binding or authoritative than it is?
- What should happen if the content is wrong, outdated, misapplied or contradicted by a person's circumstances?

Safe to Say protects people from being misled.

Safe to Reach

Will it reach the people who need it and help them recognise why it matters?

Safe to Reach is the attention, relevance and human-cognition part of the standard.

It asks whether responsible expertise can reach the people who need it, in a form that gives them a truthful reason to pay attention and stay with it long enough to use it.

This is where the framework learns from the legitimate craft of marketing, audience insight, behavioural science, storytelling, public-interest practice and misinformation research.

Those fields understand something regulated professions often leave to chance:

A person cannot use expertise they do not notice, cannot locate, do not recognise as relevant or do not care enough to stay with.

Safe to Reach does not turn regulated expertise into marketing. It puts attention and engagement under a stricter ethical rule.

The care it creates must be care the facts genuinely warrant.

The aim is to reveal a true stake that careful expertise may have buried, not to manufacture a feeling the truth does not support.

The aim is not to make someone afraid, compliant, outraged or emotionally attached to an institution. It is to help them recognise a real right, risk, choice, deadline, consequence or opportunity before it is too late.

Two kinds of reach

Kind of reach	The question	What it requires
Practical reach	Can the person find, access and receive this in time?	The right delivery route, timing, language, format, accessibility, privacy and channel
Human reach	Can they see why it may matter to them and stay with it?	Relevance, a recognisable situation, real human stakes, proportionate urgency, a coherent account and a visible next move

What Safe to Reach tests

Question	What it checks
Will the right person encounter this?	Findability, delivery route, timing and access
Can they use the channel?	Phone, digital access, language, privacy, disability access, safety, cost and time
Will they recognise themselves?	Relevance cues, examples, categories and situations
Can they see the real stake?	Rights, risks, consequences, choices, opportunity or deadline
Is the urgency real?	Proportionate timing without manufactured panic
Does it acknowledge the human reality?	Fear, shame, grief, anger, distrust, overload or anticipated consequences without exploiting them
Does it have an accountable voice?	Who is speaking, why they know and who is responsible
Does it give a reason to continue?	A visible next move, not a vague instruction or institutional preamble

The ethical governor

Safe to Reach must never use:

- false urgency;
- fear, shame or humiliation;
- scapegoating, grievance or identity pressure;
- hidden persuasion;
- false certainty;
- selective omission of material information;

- a personal story that implies an unwarranted outcome;
- apparent personalisation that exceeds the carrier's authority; or
- an action prompt designed to secure compliance rather than support informed judgement.

The rule is simple:

The moment the feeling outruns the facts, Safe to Reach has failed Safe to Say.

Safe to Reach does not mean using expertise to make someone agree, comply, disclose personal information, accept a decision or take a particular option. It means giving them enough orientation, dignity, story and truthful human context to make their own informed judgement.

What Safe to Reach produces

A reach plan that identifies:

- the people the expertise needs to reach;
- where and when they are likely to encounter it;
- practical access barriers;
- the first thing they need to see;
- the truthful human stake;
- the right level of urgency;
- relevant examples, messengers or entry points;
- the visible next move; and
- ethical guardrails against pressure or manipulation.

What the condition concerns

Safe to Reach concerns whether expertise can arrive in time, through a route people can use, with enough relevance and trust for them to notice and engage with it.

It recognises that publishing accurate information does not make it available in practice. A resource may exist but be buried in a website, delivered through an unsuitable channel, unavailable in a person's language, inaccessible on a phone, sent after a deadline or framed so generically that a person does not recognise its relevance.

Safe to Reach does not mean maximising attention at any cost. It requires truthful relevance. A person should be given a warranted reason to pay attention without manufacturing fear, urgency, certainty or pressure.

Essential requirements

Requirement	Test
Defined intended users	Has the institution identified who may need the expertise, including people likely to be excluded by standard channels?
Timely availability	Can people encounter it before the relevant deadline, decision, risk, appointment, expiry or point of action?
Usable channel	Is it available through channels people can realistically access, such as phone, mobile, print, in-person, interpreter-supported or community routes where appropriate?
Discoverability	Can a person find it without already knowing the institution's terminology, structure or internal process?
Truthful relevance	Does it explain why the information may matter, for whom and when, without implying that it necessarily applies to every recipient?
Proportionate attention	Is the level of urgency, warning, certainty and emotional framing proportionate to the evidence, stakes and available action?
Trusted and safe delivery	Is the source or messenger understandable and credible to intended users, and can people receive or access the information safely and privately where necessary?

What Safe to Reach does not require

Safe to Reach does not require universal awareness, personalised outreach in every instance or persuasive messaging designed to force action.

It does require a reasoned decision about who may need the expertise, how they are likely to encounter it and what barriers may prevent them from doing so. It also requires institutions to avoid treating a resource as “available” merely because it is online or technically published.

In some contexts, urgent or emotionally serious communication is appropriate. A safety warning, legal deadline, public-health instruction or imminent service change may require direct language and rapid contact. The condition is not “avoid fear.” It is “do not manufacture fear or urgency.” The framing must remain truthful, proportionate and connected to an available action or help route.

Safe to Reach review questions

- Who may need this expertise, and who may be missed by the usual channel?
- When does it need to reach them to remain useful?
- How will they know that it may matter to them?
- Can they find it without knowing specialist language or internal service names?
- Does the channel work on a phone, with low bandwidth, through an interpreter or without private digital access?
- Is the communication safe for people facing surveillance, coercion, family violence, stigma or privacy risk?
- Does the urgency match the evidence and the actual window for action?

- Is there a realistic next step, service or support route after the communication gains the person's attention?
- What evidence will show whether the intended people encountered it in time?

Safe to Reach protects people from being missed.

Safe to Use

Can they use it when it matters?

Safe to Use is the legal-design, human-centred-design and real-life pathway part of the standard.

It asks whether enough of the expert's reasoning, questions and pathway has travelled for a person to safely understand, judge, act, seek help, question an error and recover if the first route fails.

This is where the framework addresses the problem experts often cannot see unaided:

A person does not know what they do not know to ask.

An expert may automatically know what facts matter, what document to find, what exception may apply, who decides, what the usual next step is, what a service will ask for and what to do if that service cannot help.

A person who is unfamiliar with the system does not.

Safe to Use makes that invisible expertise practical.

It is designed for the conditions in which people often need help, not for an ideal reader with time, money, confidence, privacy, complete documents, stable internet and ready access to advice.

What Safe to Use carries

Human-use requirement	The person's question	What the expertise must provide
Practical meaning	"What is happening?"	A clear account of the situation, roles, sequence, decision and practical consequence
Hidden questions	"What else might matter?"	The questions, documents, deadlines, conditions and decision points the person may not know to ask about
Appropriate reliance	"What can I rely on?"	What is established, conditional, uncertain, incomplete, outside scope or for someone else to decide
Authority	"Who decides?"	The difference between a source, explainer, adviser, assessor, decision-maker and reviewer
Agency and action	"What can I do now?"	Rights, choices, risks, consequences and a feasible next step
Follow-through	"Can I actually manage this?"	A practical plan: when, where and how to act; what to have ready; likely obstacles and a backup
Help and handover	"Who can help me?"	A named service, contact route, access details, hours, eligibility and what to say
Contestability and recovery	"What if this is wrong or does not work?"	Review, correction, complaint, appeal, escalation, independent help, fallback and recovery routes

Safe to Use does not mean

Safe to Use does not mean:

- guaranteeing an outcome;
- pretending a difficult system is easy;
- telling a person which option to choose;
- turning general explanation into individual advice;
- assuming a person can act merely because they understand;
- treating a person's failure to act as lack of motivation; or
- using a warm tone to disguise an unavailable service or unjust process.

Where the person cannot act because of safety, money, time, capacity, access, service failure or absence of a meaningful choice, the method identifies a system finding. It does not treat the person as the failure.

What Safe to Use produces

A usable resource and pathway containing:

- practical explanation;
- hidden questions and relevant conditions;
- reliance cues and visible limits;

- rights, options, action and follow-through support;
- named help;
- a real handover;
- contestability;
- fallback; and
- recovery.

What the condition concerns

Safe to Use concerns whether expertise has travelled with enough practical meaning for a person to understand the relevant situation, judge what they can rely on and move safely through the next part of the pathway.

It is not identical to readability. A person may be able to read every sentence while still not understanding what is happening, what matters in their circumstances, which questions they need to ask, what information they need to gather, where they should go or what to do when the first route fails.

Safe to Use requires the resource, interaction or system to support realistic judgement and action without crossing into unauthorised personalised advice or decision-making.

Essential requirements

Requirement	Test
Practical meaning	Does the person have a usable explanation of what is happening, why it may matter and what the expertise does and does not establish?
Visible conditions and limits	Are key eligibility criteria, exceptions, uncertainty, risks, deadlines and limits of the answer visible at the point where they matter?
Hidden-question support	Does the material help people identify facts, documents, choices or questions that an expert would know to consider?
Feasible next step	Does it identify an action that can realistically be taken, including what is needed, by when, through which channel and at what cost or access requirement?
Authority boundary	Can the person tell what the carrier can explain, what requires individual assessment or advice and who has authority to decide?
Qualified handover	Where the boundary is reached, is there a practical route to an authorised person or service that can take the next step?
Contestability and recovery	Can the person correct an error, ask for reasons, seek review, complain, escalate, re-enter the pathway or obtain alternative help if the first route fails?

What Safe to Use does not require

Safe to Use does not mean that every person must receive a tailored answer, that every pathway must guarantee an outcome or that a general-information resource must become a substitute for professional advice.

A self-help resource may appropriately stop at explaining the general process, identifying common questions and directing the person to authorised assistance. A frontline worker may support a person to prepare for a conversation without advising what they should decide. A service may identify a review route without being able to guarantee that review will change the outcome.

The requirement is that the system does not conceal these limits or make a person carry the full burden of inferring them.

Safe to Use review questions

- Can a person explain, in practical terms, what is happening and why it may matter?
- Can they identify what they know, what remains uncertain and what depends on their circumstances?
- Are critical conditions, deadlines and exceptions visible before the person acts?
- Does the resource surface the questions, facts and documents that may matter?
- Is the next step specific, feasible and proportionate to the person's likely circumstances?
- Can the person tell whether they need advice, assessment, representation, interpretation, advocacy or another form of support?
- Does the handover connect them to a service that has relevant authority, capacity and access arrangements?
- What happens if the person cannot use the stated channel, misses the deadline, cannot gather the documents or is turned away?
- Can they obtain reasons, challenge an error, seek review, complain or escalate?

Safe to Use protects people from being abandoned.

The distinction in one view

	Safe to Say	Safe to Reach	Safe to Use
Primary knowledge base	Compliance, regulation, professional standards, evidence and ethics	Marketing craft, human cognition, attention, desire, storytelling and public-interest practice	Legal design, human-centred design, service design, accessibility and implementation
Core job	Preserve what can responsibly travel	Make sure it reaches the person and matters	Make it workable in a person's real life
Main question	Can we stand behind it?	Will it reach them and help them recognise why it matters?	Can they use it when it matters?
Main problem	"Are we allowed to say this, and is it true?"	"Will the person notice this, enter it and care?"	"Can they actually understand, judge and act?"
What has to travel	Law, evidence, conditions, exceptions, uncertainty, authority and boundaries	Relevance, human stakes, a recognisable story, trust cues and a truthful reason to stay with it	Practical meaning, hidden questions, reliance cues, action, help, contestability and recovery
Main risk	Misleading, non-compliant or unauthorised expertise	Manipulation, false urgency, irrelevance or being missed	Inappropriate reliance, confusion, abandonment or pathway failure
What it protects against	Being misled	Being missed	Being abandoned
What success looks like	Truth, authority and limits survive	The person notices and recognises why it may matter	The person can understand, judge, act, get help, question and recover

Governing principles

The governing principles are the framework's non-negotiable commitments: they guide decisions where the ten steps, delivery channels or local implementation details vary. They ensure the framework remains grounded in authority, truthfulness, practical usefulness, fairness and accountability.

1. Preserve source integrity

Expertise must remain faithful to what its source is entitled to support.

Do not simplify by removing material conditions, uncertainty, exceptions, time limits, evidentiary limits or qualifications. Do not present a general rule as a personal conclusion, or an emerging finding as settled fact.

Test: Would the source owner recognise this as an accurate and properly bounded representation of the underlying law, evidence, decision, policy or professional judgement?

2. Match claim to authority

No carrier should say, imply, recommend, assess or decide more than their role, competence and authority permit.

The framework does not require every person, website or AI system to provide individual advice. It requires each carrier to make its boundary visible, explain what it can safely explain and offer a real pathway to authorised help where individual judgement is required.

Test: What is this carrier entitled to do, and what must it hand over?

3. Make relevance truthful

People need a clear reason to attend to consequential expertise. That reason must be warranted by the facts, evidence, risk and person's likely circumstances.

The framework permits direct warnings, serious language and genuine urgency where they are justified. It rejects manufactured fear, urgency, certainty, pressure or emotional manipulation as substitutes for relevance.

Test: Does the communication explain why this may matter without overstating what is known, likely or required?

4. Design for real conditions

Expertise must work under the conditions in which people actually encounter it, not only in ideal conditions imagined by the institution.

People may be stressed, tired, time-poor, using a phone, reading in a second language, lacking documents, concerned for their safety, living with disability, without private internet or unable to contact the first service they try. The framework treats these conditions as design inputs, not user shortcomings.

Test: Could a person use this safely at 4:47pm, under ordinary constraints and after the first route has failed?

5. Support judgement, not compliance

The purpose is to enable informed, safe judgement and action, not to secure attention, completion or behavioural compliance at any cost.

A successful system does not simply persuade a person to act. It helps them understand what matters, what they can rely on, what remains uncertain and which action is feasible in their circumstances. Where the person has a genuine choice, the framework should preserve that choice.

Test: Does this support the person's understanding and agency, or does it merely drive a preferred institutional outcome?

6. Make boundaries useful

A professional, institutional or technical boundary should not become a dead end.

It may be necessary to say, “We cannot advise on your individual circumstances,” “This tool cannot determine your eligibility,” or “This information does not replace clinical assessment.” But that statement should be accompanied, where possible, by an explanation of what the resource can do, what further help is needed, how to reach it and what to do if the first route fails.

Test: Does the boundary protect against inappropriate reliance while leaving the person with a usable route forward?

7. Treat access as pathway access

Access is more than being able to open a document or website. People need a realistic way to understand the material, complete the relevant step, reach help, use an interpreter or accessible format, ask for review and recover from failure.

A resource is not fully accessible if the route it describes is inaccessible in practice.

Test: Can intended users access and use the whole pathway, not merely the first resource?

8. Build in contestability and recovery

People must be able to question, correct, review, complain, escalate or seek alternative help when information, decisions, services or digital systems fail.

The framework treats error correction and recovery as normal parts of consequential systems, not as exceptional afterthoughts. A pathway with no realistic recovery route shifts too much risk onto the person.

Test: If this is wrong, unavailable, misunderstood or unsuccessfully acted on, what can the person do next?

9. Share accountability across the pathway

No single team necessarily controls the whole journey, but the institution must not allow distributed responsibility to become unowned failure.

Source owners, legal and professional-governance teams, communications, accessibility, service delivery, digital, AI, complaints and operational teams should have clear, connected responsibilities. The person experiences one system, not separate internal functions.

Test: Who owns accuracy, reach, accessibility, action, handover, correction and recovery, and can those owners work together?

10. Apply proportionately and learn

The level of review, testing, documentation and governance should reflect the consequence of getting the transfer wrong.

High-stakes decisions, deadlines, safety warnings, clinical instructions, financial hardship pathways and AI-supported systems that appear authoritative require deeper testing. Lower-consequence resources may need a lighter application. In every case, the framework should be tested, monitored and revised in response to evidence of failure.

Test: Is the level of care proportionate to the likely harm of misunderstanding, inaction, delay, over-reliance or pathway failure?

Principle summary

Principle	What it protects
Preserve source integrity	Accuracy, conditions and uncertainty
Match claim to authority	Scope, competence and appropriate reliance
Make relevance truthful	Attention without manipulation
Design for real conditions	Usability under pressure and constraint
Support judgement, not compliance	Agency and informed action
Make boundaries useful	Safe escalation rather than abandonment
Treat access as pathway access	Practical inclusion, not mere publication
Build in contestability and recovery	Correction, review and resilience
Share accountability across the pathway	Joined-up institutional responsibility
Apply proportionately and learn	Feasibility, improvement and evidence-based revision

Practical method

What the method addresses

The Paper-to-Person Gap opens when an institution provides information but leaves the person to supply the reasoning, orientation, coordination and recovery work required to use it.

A notice may say that time limits apply without making the relevant deadline easy to identify. Guidance may explain that circumstances affect eligibility without helping a person recognise which circumstances matter. A referral may advise someone to seek help without establishing whether suitable assistance is reachable, affordable, authorised to help or available before the relevant deadline. A digital tool may provide accurate general information without making clear what it cannot determine or how a person can obtain qualified human support.

In each case, the source may be correct. The failure occurs in the transfer.

The person is expected to undertake work that an expert or institution may perform automatically: identify relevant facts, infer conditions and exceptions, distinguish explanation from individual advice or formal determination, prepare the right questions, locate a suitable service, coordinate multiple steps and recognise when further judgement is required.

Safe to Say, Safe to Reach, Safe to Use proposes a method for making this missing work visible, testing whether it has been carried by the resource or pathway, and assigning responsibility where it has not.

Transfer failure	What the method asks teams to do
Meaning, evidence, conditions or authority are lost, obscured or overstated	Identify the authoritative source and the claim it supports; capture the reasoning that matters; preserve material qualifications, uncertainty and conditions; and define what the carrier is and is not authorised to explain, advise, assess or decide.
Accurate expertise cannot be found, arrives too late or does not register as relevant	Identify who may need it, when it needs to reach them and which channels they can realistically use; design for discoverability and truthful relevance; and ensure that urgency, certainty and emotional framing remain proportionate to the facts.
A person can understand the resource but cannot use the pathway around it	Connect explanation to a feasible next step, including practical requirements, deadlines, access needs, qualified help, handover, review and recovery if the first route fails.
The resource implies more than the source or carrier can support	Distinguish general information from individual advice, assessment, recommendation and formal determination; make reliance limits visible; and design a practicable transition to an authorised source of help.
Separate approvals leave an end-to-end failure undetected	Review substance, delivery, accessibility, service capacity, escalation and recovery together; test realistic use; and assign connected responsibilities for maintenance and repair.
A structural barrier prevents action	Record the barrier as a system finding; identify the responsible owner; distinguish a communication problem from a service, policy, funding, authority, technology or institutional problem; and escalate the issue where local repair is insufficient.
The first contact, referral or digital route fails	Identify alternatives, clarify who remains responsible, provide a route to correction or escalation and make recovery visible rather than treating the first failed route as the end of the person's options.

The method is not designed to make every institution responsible for solving every problem a person faces. It is designed to prevent institutions from presenting a technically correct resource, disclaimer or referral as if it completes the work of helping.

Helping without exceeding authority

The method addresses a practical professional dilemma: how to provide useful help without overstating evidence, making an unsupported assessment or exceeding the authority of the person, service or system carrying the expertise.

A carrier may be able to explain a general rule, process, risk, deadline, evidence source or service pathway. They may not be authorised to determine eligibility, predict a legal outcome, diagnose a condition, recommend a financial course of action or make a binding decision.

The framework does not remove that boundary. It makes the boundary operational.

It asks the professional or institution to establish:

- What can responsibly be explained from the available source.
- What conditions, uncertainty and limitations must remain visible.

- What a person can appropriately rely on.
- What requires individual advice, assessment, representation, decision-making or another qualified form of help.
- Which person, service or institution has authority to provide that next stage.
- Whether the route to that help is realistic, safe and accessible.
- What the person can do if the expected route is unavailable, unsuitable or unsuccessful.

This turns a boundary from a dead end into a designed transition.

For example, “We cannot advise you about your individual circumstances” may be a necessary and accurate statement. On its own, however, it can leave a person with no practical way forward. A more responsible transfer may explain the relevant general process, identify the facts or documents likely to matter, state what the resource cannot determine, provide a realistic route to qualified advice and identify an alternative or escalation option if the first route does not work.

The aim is not to authorise every carrier to answer an individual question. It is to make the distinction between general explanation, individual advice, assessment and formal determination visible, and to ask whether the person has a practicable route between them.

Connecting institutional responsibilities

Accuracy, communication, accessibility, service delivery, professional responsibility, digital systems, complaints and review may each be managed by separate functions. Each review may be sound within its own scope while the person’s overall pathway remains unusable.

A legal team may confirm that a notice is accurate. A communications team may make it clear. An accessibility team may ensure that the webpage works with assistive technology. An operations team may confirm that a referral process exists. A digital team may ensure that an online form is live.

Yet the person may still be unable to identify the deadline, gather the required evidence, safely use the channel, reach the referred service, obtain an interpreter, receive an answer before the deadline or recover when the first route fails.

The method provides a shared process for examining how these responsibilities connect.

It asks:

- What must travel from the authoritative source?
- What is each carrier authorised and able to support?
- Who may need the expertise, and in what real-world conditions?
- What must a person understand, decide, do or avoid?
- What action or service pathway is being proposed?
- Who receives responsibility at the next handover?
- What happens when the standard pathway does not work?
- Who owns the correction, review, escalation or system change required?

The method therefore treats the person's experience as an end-to-end test of institutional coordination. The person encounters one pathway, not the organisation's internal divisions.

Purpose of the method

The proposed purpose of Safe to Say, Safe to Reach, Safe to Use is:

To help institutions identify and address what is missing between issuing consequential expertise and enabling a person to understand, judge and use it responsibly in real life.

This includes making system failure visible when it cannot be repaired through communication alone.

A finding may concern unclear wording or hidden conditions. It may instead concern inaccessible channels, missing interpreters, unrealistic deadlines, unavailable services, referral rejection, lack of review, inadequate staffing, broken digital systems, unclear authority, unsafe contact requirements or an underlying policy or legal barrier.

The method does not imply that every barrier can be solved locally. It requires that the barrier be identified accurately, assigned to an appropriate owner and not misrepresented as a failure of the person to understand, comply or try hard enough.

The problem and proposal are distinct

The literature synthesised in this paper provides substantial grounds for examining comprehension, legal and health capability, information seeking, administrative burden, implementation barriers, accessibility, professional boundaries, procedural fairness, appropriate reliance and recovery after pathway failure.

Those bodies of work support the importance of the problems the framework addresses.

They do not establish that this particular integrated method is effective.

The Paper-to-Person Gap, the three conditions, the ten-step method and the 4:47pm Test are proposed organising tools. Whether their application improves responsible reach, practical understanding, appropriate reliance, informed judgement, feasible action, successful handover, access to help, contestability or recovery requires evaluation.

The framework's practical value depends on what it helps teams recognise, what they are able to change and whether the resulting resource and pathway work better for the people expected to use them.

How the method works

The method is the practical process for making regulated expertise work beyond the person, institution or system that originally holds it.

It is used when a law, policy, research finding, professional judgement, formal decision, service pathway or AI-supported resource has to travel to a person who must understand, judge or act on it.

The method does not begin by rewriting a document.

It begins by asking:

What has to travel with this expertise for a real person to use it safely?

The answer is rarely only the conclusion.

A person may also need the source, reasoning, conditions, exceptions, uncertainty, authority, relevance, hidden questions, practical pathway, handover and route to challenge or recover.

The method is:

Orient → Capture → Anchor → Bound → Reach → Translate → Connect → Verify → Test → Maintain

It has five phases.

Phase	Steps	What the phase does
1. Start in the real situation	Orient	Establish what the expertise is for, who needs it, what is at stake, who can carry it and what conditions shape real use
2. Protect what must remain true	Capture · Anchor · Bound	Surface the source, reasoning, evidence, authority, conditions, uncertainty and professional boundaries that must survive transfer
3. Make it reach the person	Reach	Design the practical and human route through which expertise reaches the people who need it and truthfully shows why it may matter
4. Make it usable in life	Translate · Connect	Make practical meaning, hidden questions, reliance, action, help, contestability and recovery available to the person
5. Check that it survived	Verify · Test · Maintain	Confirm that truth, authority, reach and use survive release and continue to work when conditions change

The sequence gives the method a disciplined order:

First know what may travel. Then make it reach the person. Then make it usable. Then check that nothing important was lost.

The process is not mechanically one-way. Testing may reveal that a claim needs further qualification, a channel does not reach the intended person, the carrier is being asked to go beyond authority or the pathway does not actually work.

But one rule remains fixed:

0. Orient

What is this expertise for, who is it for, what is at stake, who can carry it and what might the person not know to ask?

Before changing the resource, define the transfer problem.

A court notice, research summary, hospital discharge instruction, policy, professional guideline, regulator FAQ, service referral and AI-supported response should not all be carried in the same way.

The acceptable degree of simplification, action, personalisation, verification and escalation changes according to the function, carrier and stakes.

Establish the function

Decide what the expertise is meant to do.

It may be intended to:

- inform;
- warn;
- explain;
- instruct;
- enable action;
- obtain consent;
- record a formal decision;
- support a person to prepare for a qualified conversation;
- advise; or
- determine an outcome.

Some functions are conditional.

For many carriers, including most frontline workers, public resources and general-purpose AI tools, **advise** and **decide** are not available functions. The resource may explain the rule, name the pathway and prepare a person for the next step. It must not drift into applying the rule to that person's facts, assessing their prospects or directing their personal choice where the carrier lacks authority.

Establish the stakes

The greater the consequence of misunderstanding, inappropriate reliance or failed action, the stronger the requirements for source authority, verification, testing, traceability, escalation and maintenance.

Stakes	Likely consequence if the transfer fails	Typical response
Low	Inconvenience or minor error	Proportionate source check and usability review
Moderate	Money, employment, access, eligibility or service delay	Stronger source control, user testing and pathway check
High	Legal rights, safety, health, liberty, housing, family or major financial consequences	Formal verification, real-user testing, clear handover and recovery route
Critical	Irreversible or potentially catastrophic harm	Specialist review, strong authority controls, escalation, monitoring and maintenance

Name the person and the real conditions of use

Design for the person who is likely to need the expertise, not an ideal reader who has time, confidence, complete documents, stable internet, privacy and ready access to advice.

Ask:

- Who is expected to use this?
- What may be happening in their life when they receive it?
- What language, disability, technology, safety, money, privacy, time, document, trust or access barriers may affect use?
- What might make it unsafe or impossible for them to receive, read, act on or follow up?
- What do they already know?
- What do they not know to ask?
- What might they reasonably assume that would be unsafe or wrong?

Name the authority and carrier

Separate two questions:

- 1. Who has authority over the outcome?** Who can advise, assess, determine, approve, refuse, review or formally decide?
- 2. Who has authority to carry the expertise?** Who may explain, translate, personalise, illustrate, direct, refer or hand over?

The same underlying expertise may lawfully travel further through one carrier than another.

A practising lawyer may apply law to a client's facts. A public resource may explain the usual rule and name a pathway. A frontline worker may explain process and help someone reach the right service. A general-purpose AI tool may retrieve or explain approved general material only within carefully designed limits.

The carrier's authority sets the ceiling for what Translate and Connect may do.

Identify barriers and enablers

Ask what will stop the expertise from being used as intended.

Examples include:

- The person does not know the resource exists.
- The pathway requires documents people commonly do not have.
- The delivery route is inaccessible or unsafe.
- The service has no capacity.
- The phone line closes before the person can call.
- The online form cannot be completed on a phone.
- Staff do not know the resource exists or cannot perform what it promises.
- Another service uses different categories or has conflicting requirements.
- The institution rewards speed, risk avoidance or closed files rather than safe use.
- The person has good reason not to trust the source.

Diagnose the existing failure

Where a resource or pathway already exists, begin with evidence of where it breaks down. A record such as “information provided; referral made; did not proceed” describes events without establishing whether the expertise became usable.

Ask:

- Could the person encounter the expertise in time and recognise its relevance?
- Could they understand the practical situation and judge what they could rely on?
- Could they identify and take a feasible next step?
- Could they reach help, question an error or recover when the first route failed?
- What did they have to infer, guess, navigate or risk alone?
- Did the breakdown occur in the underlying rule, template, resource, conversation, handover or follow-up?
- What should have caught it, and why did it not?
- Who has the authority and capacity to repair it?

These questions locate failures in the passage of expertise and its surrounding pathway. They do not establish that a person lacks intelligence, motivation or capacity. Where the person understands but cannot act, examine safety, resources, access, service availability and meaningful choice before proposing a communication fix.

Produces: an **orientation record** setting out function, intended person, stakes, authority, carrier, conditions of use, likely barriers, enablers and the questions the person may not know to ask.

1. Capture

What does the actual source say, mean and qualify?

Start with the source, not with a summary, marketing copy, inherited wording or an assumed interpretation.

For documented expertise, identify the authoritative version and the parts carrying legal, factual, evidentiary, ethical or professional weight.

For expertise held mainly in people's heads, capture not only the headline conclusion but also the reasoning that makes it safe and useful.

Surface:

- the central message;
- the source material;
- assumptions;
- evidence and its limits;
- qualifications;
- exceptions;
- examples;
- recurring questions;
- practical judgement;
- tacit distinctions;
- common misunderstandings;
- decision points;
- escalation triggers; and
- the outcome the expertise is intended to enable.

Ask:

- What does this actually mean?
- What feels obvious to the expert because they know the field?
- What would the expert immediately correct if someone repeated this back wrongly?
- What conditions or distinctions would change the answer?
- What questions do people usually ask only after something has gone wrong?
- What does a person need to do, know or prepare in order to use this safely?
- What would the person be at risk of assuming if only the headline travelled?

Produces: a **source and expert-reasoning record**.

2. Anchor

What supports each important claim, and who has authority?

Connect each material claim to the right source.

The source may be:

- legislation;
- regulation;
- formal policy;

- professional standard;
- research evidence;
- organisational record;
- a formal decision;
- a recognised practice guide;
- clearly identified professional judgement; or
- a clearly labelled lived-experience account.

Distinguish what is:

Claim type	Meaning
Established	Directly supported by law, evidence, policy, standard or formal record
Determined	Formally decided by the authorised decision-maker
Interpreted	A reasoned reading of the source, not the source itself
Inferred	A conclusion drawn from available facts or evidence
Recommended	A proposed course, professional view or practice suggestion
Experienced	A person's or community's account of how something works in life
Unknown or contested	Not established, incomplete, disputed or dependent on further facts

Then name authority.

- Who can explain the process?
- Who can give individual advice?
- Who can assess the evidence?
- Who can make the decision?
- Who can review or change the decision?
- Who can provide independent support?

Authority should not be hidden in passive language.

“The application may be refused” hides who decides. “The Department may refuse the application” identifies the actor and makes authority visible.

Produces: a **claim, source and authority map**.

3. Bound

What must the expertise not accidentally promise, imply or decide?

Bound protects the person from inappropriate reliance and protects the carrier from crossing an authority boundary.

Identify before redesigning:

- What is not established.
- What facts could change the answer.
- What remains uncertain, disputed or incomplete.
- What must remain qualified.
- What the resource cannot determine.
- What may be rewritten freely.
- What must remain intact or be reviewed by the appropriate owner.
- Where general explanation must stop and individual advice, assessment or a formal decision must begin.

Small changes can create large drift:

Safe source meaning	Unsafe drift
May	Will
Associated with	Causes
Usually	Always
One professional view	The organisation's conclusion
General information	Individual advice
Possible pathway	Guaranteed outcome

The last row often changes not through the wording but through the interaction.

A resource may begin as general information. But if a person gives their own facts and the carrier applies the general rule to them, assesses their position, predicts an outcome, compares options or recommends a course of action, the transfer may have become individual advice.

That matters whether the carrier is a person, chatbot, website, form or AI system.

A disclaimer does not alter the legal or professional character of what was actually done.

A disclaimer is not a shield.

The boundary must be designed into the pathway.

If the carrier cannot safely go further, the resource must name what happens next, who can take over and how the person can reach them.

Produces: a **boundary map** marking preserved conditions, reliance limits, personalisation limits, carrier limits and handover points.

4. Reach

Can this expertise reach the people who need it in time, and will they see why it may matter?

Reach comes after Bound for a reason.

First establish what may responsibly travel.

Then design how that bounded expertise will reach attention, relevance and human significance without distorting the truth.

Reach has two jobs.

Practical reach

Can the person find, access and receive the expertise through a route that works in their life?

Ask:

- Where will the person look first?
- Does the expertise appear there?
- Will it reach them before the deadline, decision, risk or missed opportunity?
- Can they use the delivery route with their device, language, disability, time, privacy and safety constraints?
- Is it available in an accessible format?
- Can it be read on a phone?
- Can it be received safely where coercion, surveillance or family violence may be relevant?
- Is an interpreter, advocate, support person or trusted intermediary needed?
- Is the information hidden behind a portal, search term, paywall, form or institutional label the person is unlikely to know?

Human reach

Can the person see why it may matter and stay with the expertise long enough to use it?

Ask:

- What is happening to the person or likely to affect them?
- What right, risk, choice, cost, consequence, opportunity or deadline is involved?
- What do they need to know first?
- Will they recognise themselves in the opening?
- What facts or situations would help them see whether it may apply?

- What may make the expertise hard to receive: fear, shame, grief, anger, distrust, urgency or overload?
- How can the resource acknowledge those realities without using them to force an outcome?
- Is the urgency real and proportionate?
- Does the person have a truthful reason to continue?
- Is the next move visible?

Safe to Reach may use:

- clear openings;
- recognisable situations;
- relevant examples;
- human stakes;
- proportionate urgency;
- trustworthy and accountable ownership;
- structure that helps a person enter quickly; and
- a visible next move.

It may not use anything on the ethical governor list under Safe to Reach in Section 5.

Produces: a **reach plan** setting out intended people, delivery routes, timing, access conditions, entry point, genuine human stake, likely barriers, truthful urgency, visible next move and ethical guardrails.

5. Translate

What does the person need in order to understand what this means, see the questions that matter and work out what they can rely on?

Translation changes the form of expertise without changing its meaning, function or authority.

It is not simply replacing difficult words.

A person may understand every word in:

“You may be eligible. Your circumstances may affect the outcome.”

and still not know:

- what circumstances matter;
- whether the information applies to them;
- what “may” means;
- who decides;
- what deadline applies;
- what they should do first; or

- what happens if they cannot reach the service.

Translation makes the invisible expert reasoning visible without converting it into unauthorised individual advice.

Design for the person and the conditions of use

Consider:

- time;
- stress;
- safety;
- privacy;
- language;
- disability;
- access to technology;
- money;
- documents;
- service access;
- trust;
- workload;
- emotional conditions; and
- what the person does not yet know to ask.

People do not receive consequential expertise as detached reasoning machines.

Fear, urgency, shame, uncertainty, distrust, grief, cognitive overload and anticipated consequences can affect whether they can engage with it at all.

This does not mean creating an emotional response. It means:

- acknowledging likely concerns and human stakes;
- explaining why uncertainty exists;
- avoiding unnecessary coldness, blame or shame;
- making urgency understandable without manufacturing panic; and
- providing choices or support where the expertise may be difficult to receive or act on.

Translation may involve

- Changing the order so the most relevant point appears first.
- Grouping related material together.
- Explaining tacit reasoning and practical implications.
- Bringing forward common questions, hidden conditions and likely misunderstandings.
- Defining technical or legal terms where explanation helps.
- Using ordinary words where they carry the same meaning.
- Keeping technical terms where they carry meaning an ordinary word would lose.

- Making conditions, uncertainty, evidence limits, scope, authority and decision-makers visible.
- Adding examples, diagrams, headings, summaries, glossaries, accessible formats, translated versions or spoken alternatives.
- Designing layers so a person can begin with what matters and go deeper where needed.

Start with sentence shape

- 1. Un-nest.** Move a qualification out of the middle of a sentence and give it its own sentence.
- 2. Prefer the common word** where the meaning survives.
- 3. Say who does what.** Name the actor and decision-maker where known.
- 4. Use ordinary capitalisation.**

For example:

Before A person who, having been given notice under section 12 and not having responded within the period specified in that notice, fails to attend may be dealt with under section 14.

After This applies to a person who was given a section 12 notice and did not respond in time. If that person then fails to attend, they may be dealt with under section 14.

Nothing necessary has been removed. The qualification survives. The reader can now hold it.

When full complexity cannot safely fit into one layer, do not erase the problem. Use layers.

For example:

- 1.** What is happening and why it may matter.
- 2.** Conditions, limits, uncertainty and exceptions.
- 3.** Source and authority.
- 4.** Individual advice, assessment or formal decision boundary.
- 5.** Help, action, fallback and recovery.

Produces: a **translated expertise resource**.

6. Connect

What does the person need to do, ask, access, question or recover from?

Understanding alone is not enough.

A person needs a route from expertise to their own judgement and action.

Connect builds that route.

It should make visible:

- Why the expertise may matter and how urgent it is.
- Rights, choices, risks, consequences and decision-makers.
- What is known, conditional, uncertain or outside scope.
- Questions the person may need to ask.
- What documents, records or evidence may be relevant.
- A realistic next step.
- What to do, by when, through which route, with what documents and using what words.
- Reachable help.
- A backup if the first route fails.
- A route to question, correct, review, appeal, complain or escalate.
- Recovery after a failed contact, rejected referral, missed step or incorrect outcome.

The action boundary

Choices and next steps can be provided at two levels.

Level	What can be provided
General pathway level	What exists for people in this kind of situation: process, deadline, public service, usual documents, questions to ask and how to access help
Personal advice level	Which option this person should choose, assessment of their evidence, prediction of their outcome or recommendation about their personal position

A carrier authorised only to provide general information should fully build the first level and stop before the second.

The stop is not a gap in the resource.

It is where the handover goes.

Make the handover real

A handover has four parts:

- 1. A named service:** not merely “a lawyer” or “a professional.”
- 2. A route:** phone number, website, counter, form, webchat or appointment pathway.
- 3. Access conditions:** hours, cost, documents, eligibility, language and accessibility.
- 4. A backup:** what to do if the first route cannot help or cannot be reached.

A referral is not a handover. A disclaimer is not a plan.

Support follow-through without steering the decision

A person may understand and intend to act but still not follow through because life intervenes.

Where appropriate, support the person to carry out a decision they have already made:

- When will they take the step?
- Where will they do it?
- What do they need to have ready?
- What is the first likely obstacle?
- What is the backup if that obstacle occurs?

This is not a mechanism for directing a person's choice. It is support for carrying out their own choice where the resource is entitled to do so.

Identify when the issue is not expertise

Do not treat a person's inability to act as a failure of understanding if the real barrier is:

- safety;
- money;
- time;
- disability;
- trauma;
- privacy;
- lack of documents;
- lack of authority;
- digital exclusion;
- language;
- geographic isolation;
- service capacity;
- an unjust requirement; or
- absence of any meaningful option.

Those are system findings.

Produces: a **usable pathway** with action, follow-through, help, handover, contestability, fallback, recovery and a record of system changes needed for the pathway to work.

7. Verify

Has the expertise remained accurate and within proper boundaries?

Return the translated resource and pathway to the people accountable for the underlying expertise.

Verification is not cosmetic proofreading.

It checks whether the final transfer still accurately represents:

- law, policy, evidence or formal decision;
- conditions, exceptions and deadlines;

- uncertainty, scope and missing facts;
- authority and decision-making roles;
- privacy, ethics and safety requirements;
- professional boundaries;
- the distinction between general information, individual advice and formal determination; and
- the limits of what the carrier may personalise.

The relevant source owner, legal owner, policy owner, professional owner or decision owner must check the substance.

The operational owner must confirm that the action, help, handover and fallback routes actually exist.

Produces: a **substantive verification record** identifying what was reviewed, who reviewed it, what changed and what remains outside scope.

8. Test

Can people find it, see why it matters and use it in real life?

Do not ask:

“Was that clear?”

Give intended users and frontline professionals the resource or pathway and ask them to show what they think is happening and what they would do.

Test whether they can:

- Find or receive the expertise in time.
- See why it may matter to them.
- Recognise whether it may apply to them.
- Find the part that matters.
- Explain the main point in their own words.
- Identify what is established, conditional, uncertain or outside scope.
- Identify what they can appropriately rely on.
- See a question, document, condition or choice they might otherwise not have known mattered.
- Identify choices, rights and decision-makers.
- Take a safe next step.
- Describe when and how they would realistically take that step.
- Reach help.
- Find a backup if the first route fails.
- Find a review, correction, complaint or escalation route.

- Tell where the resource stops and who takes over.

Test reliance in practice

Being able to repeat a resource's limits does not establish that a person will use it within those limits.

Test both what the person understands and what they would conclude or do. For example:

- Would they treat general guidance as a decision about their own eligibility?
- Would they act on a headline while overlooking a condition that changes it?
- Would they assume that a confident answer has assessed their circumstances?
- Would they recognise when further facts, qualified advice or a formal decision are needed?
- Would they dismiss supported information because its uncertainty is more visible than that of a competing source?

Use scenarios that include supported claims, conditional claims and questions the carrier cannot determine. Observe the basis for the person's judgement, rather than treating confidence, satisfaction or recall of a disclaimer as evidence of appropriate reliance.

The aim is warranted reliance: using the resource for what it can responsibly support and recognising where another source of judgement must take over.

Test under ordinary conditions, not only calm workshop conditions.

A resource fails if people can explain the wording but cannot complete the task.

If the resource names a phone number, call it. If it promises an online form, try it on a phone. If it directs people to a service, confirm the service can accept them. If it offers review, test whether the route is visible and usable.

Produces: evidence of reach, use and failure, including a list of system problems that better transfer design cannot solve alone.

9. Maintain

Who keeps this accurate, reachable and useful when the law, evidence, service or pathway changes?

Stale expertise is both a Safe-to-Say and Safe-to-Use failure.

A resource that was accurate when issued can become wrong or unsafe when:

- the law changes;
- the evidence moves;
- a policy is replaced;
- a service closes;
- a pathway changes;
- a phone line becomes inaccessible;
- an AI system changes;

- the intended audience changes; or
- repeated real-world use reveals an unanticipated misunderstanding or failure.

Assign maintenance ownership at the beginning.

For high-stakes resources, set review dates as well as change triggers.

Monitor:

- Are people finding it?
- Are they recognising why it matters?
- What questions are they still not seeing until too late?
- Where are they getting stuck, dropping out or asking the same question?
- Are they missing deadlines or relying too far?
- Are people intending to act but unable to follow through? Why?
- Are help routes actually reachable?
- Are review and escalation pathways working?
- Are staff creating workarounds?
- Has the source, decision, pathway or delivery route changed?

Produces: a named **maintenance owner**, review trigger, review date, monitoring plan and change record.

The method in one view

Phase	Step	Focus
Start in the real situation	0. Orient	Purpose · person · stakes · authority · carrier · barriers
Protect what must remain true	1. Capture	Source · reasoning · conditions · tacit knowledge
	2. Anchor	Claims · evidence · authority · decision-maker
	3. Bound	Limits · uncertainty · advice boundary · handover point
Make it reach the person	4. Reach	Practical access · relevance · human stake · truthful reason to attend
Make it usable in life	5. Translate	Meaning · structure · hidden questions · reliance cues
	6. Connect	Action · follow-through · help · challenge · fallback · recovery
Check that it survived	7. Verify	Truth · authority · limits · routes
	8. Test	Real people · real tasks · real conditions
	9. Maintain	Ownership · monitoring · review · repair

The method is complete only when the expertise remains safe to say, can reach the people who need it and gives them enough of the source, reasoning and pathway to use it safely in real life.

Essential requirements and adaptable implementation

The framework distinguishes between **essential requirements** and **adaptable implementation**.

Essential requirements are the conditions that must be addressed if the framework is being applied. They express the minimum integrity of Safe to Say, Safe to Reach and Safe to Use.

Adaptable implementation concerns how an institution meets those requirements in a particular setting. It allows the framework to be used across sectors without treating every resource as if it requires the same format, workforce, budget, channel or level of personalisation.

Element	Essential requirement	Adaptable implementation
Source	Expertise must be traceable to an appropriate source or identified judgement	Citation, link, policy reference, version note, named decision-maker, evidence summary or internal source register
Authority	The carrier's role and boundary must be clear	Disclaimer, role label, script, interface cue, escalation rule, training, approval workflow or user-facing explanation
Conditions and uncertainty	Material limits, conditions and uncertainty must not be lost	Layered content, expandable detail, examples, decision trees, "what this does not tell you" sections, verbal explanation or interpreter support
Reach	Intended users must have a realistic route to encounter the expertise in time	SMS, letter, phone, website, community partner, in-person contact, translated material, outreach, reminder or professional intermediary
Accessibility	People must have a reasonable route to access and operate the resource or service	Accessible digital design, plain language, translation, captioning, alternative formats, telephone support, assisted completion, flexible channels or reasonable adjustments
Practical meaning	A person must be able to understand what matters and what the expertise does not establish	Examples, scenarios, visual explanations, scripts, FAQs, short summaries, layered detail, conversations or guided forms
Action	The stated next step must be feasible enough to attempt	Online form, phone number, booking link, walk-in option, document checklist, reminder, callback, assisted service, advocate or interpreter
Handover	A person must have a defined route to qualified help when the boundary is reached	Warm referral, direct booking, case transfer, priority route, shared information with consent, service directory or escalation contact
Recovery	A person must be able to identify an available correction, review, complaint, escalation or alternative pathway	Review notice, complaints process, supervisor contact, ombuds route, alternative service, re-entry point, emergency option or correction workflow
Testing and accountability	Failures must be discoverable and owned	User testing, frontline review, mystery shopping, accessibility testing, call data, complaint analysis, audit, monitoring dashboard, incident review or governance report

The distinction is important.

A statutory decision letter may need formal wording, identified review rights, a phone number, translated versions, an online portal and a staffed support route. A public research summary may instead need clear sourcing, transparent uncertainty, a brief statement of relevance and a link to further information. A chatbot may need visible source links, explicit limits, escalation triggers, conversation logging and a human-support route.

The form changes. The underlying responsibilities do not.

The minimum safe transfer

The framework does not require every resource to carry the entire body of underlying expertise. It requires a **minimum safe transfer**: enough of the expertise must travel for the person to avoid foreseeable misunderstanding, inappropriate reliance or avoidable abandonment.

The minimum safe transfer will usually include:

- What this is and where it comes from.
- Why it may matter, and for whom.
- What it establishes, and what it does not establish.
- Any critical conditions, deadlines, uncertainty or limits.
- What a person can reasonably do next.
- What requires individual advice, assessment, decision-making or another authorised form of help.
- How to reach that help.
- What to do if the first route fails or the information appears wrong.

The appropriate level of detail depends on the consequences of error. A higher-stakes context requires more careful work on source integrity, reach, accessibility, action, handover, contestability and recovery.

Applying the method

The method can be used at different levels.

Level of application	Example	Typical use of the method
Single resource	A decision letter, web page, form, FAQ, script or public warning	Review source, boundaries, wording, findability, accessibility, next action and review route
Single interaction	A call-centre conversation, consultation, intake interview or chatbot exchange	Clarify what the carrier can say, what must be asked, what must be explained and how handover occurs
Service pathway	A referral route, benefit application, discharge pathway or complaints process	Map stages, authority, handovers, barriers, capacity, recovery points and ownership
Digital or AI-supported system	A portal, eligibility tool, automated notice or AI assistant	Review sources, output boundaries, user interpretation, escalation, accessibility, monitoring and correction
System change or policy implementation	A new law, program, regulatory requirement or service model	Identify affected people, map how expertise will travel, test reach and use, assign owners and monitor failures

A team may begin with a high-risk point in an existing pathway: a recurring complaint, missed deadline, failed referral, disproportionate call volume, repeated user confusion, incorrect AI reliance, low completion rate or pattern of escalation.

The method then asks the team to trace backward and forward:

- What expertise was the person expected to use?
- Where did it come from?
- What did the person receive?
- What did they reasonably understand?
- What action was expected?
- What barrier intervened?
- Which team owned that barrier?
- What route existed when the expected path failed?
- What change would make the transfer safer without exceeding authority or overstating certainty?

What successful application looks like

The framework should not be treated as a certification that a resource, service or institution is “safe.” It is a disciplined way to identify and reduce foreseeable failures in the transfer of consequential expertise.

A credible application produces evidence that the institution has:

- Identified the source, authority, conditions and limits of the expertise.
- Distinguished general information, advice, assessment, recommendation and formal decision-making.
- Considered who needs the expertise and the conditions in which they will encounter it.
- Designed truthful relevance without manipulative urgency or pressure.
- Made practical meaning, next steps and boundaries visible.
- Tested accessibility, feasibility, handovers and recovery under realistic conditions.
- Identified where a pathway depends on another service, system or decision-maker.
- Assigned owners for content, access, delivery, escalation, correction and review.
- Monitored recurring failures and revised the pathway where evidence shows that people are being left behind.

The intended result is not perfect communication. It is a more accountable transfer of expertise: accurate enough to issue, reachable enough to matter and usable enough to support safe real-world judgement and action.

The 4:47pm Test

Does the expertise work when life is not ideal?

The real question is not whether regulated expertise works for a calm person with time, confidence, stable internet, complete documents and a highly available professional.

The real question is whether it works at **4:47pm on a Tuesday**.

“4:47pm” names a **condition**, not a literal testing time: limited capacity, incomplete information, competing demands and a narrowing window for help. Tuesday keeps the situation ordinary and recognisable. This is not a test of whether a resource works during business hours. It is a test of whether expertise still works when life is already happening around it.

At 4:47pm, a person may be tired, frightened, angry, grieving or overloaded. Some facts may be missing. A deadline may be close. Another service may close at five. Their child may need collecting. Their phone may be nearly flat. English may not be their first language. They may not have a private place to make a call. They may not know whether they can trust the source.

The professional helping them may also be under pressure: incomplete information, limited authority, another person waiting and no ability to guarantee the outcome or availability of the next service.

This is not an exceptional test case.

It is often the ordinary condition in which people need law, policy, research and professional expertise to work.

The 4:47pm Test asks whether expertise still works when the person cannot supply the time, knowledge, privacy, confidence, documents or emotional distance the system quietly assumes.

If it works only for an ideal reader, the Paper-to-Person Gap remains open.

What the test checks

The 4:47pm Test checks all three conditions of the proposed standard.

Condition	The 4:47pm question
Safe to Say	Does the expertise still preserve truth, limits, conditions, uncertainty, authority and professional boundaries under pressure?
Safe to Reach	Can the person find it, access it safely and see why it matters before time or attention runs out?
Safe to Use	Can they understand it, identify what they do not know to ask, judge what they can rely on, take a feasible next step, reach help, question an error and recover if the first route fails?

The test is not asking whether every person experiences every constraint at once.

It asks teams to test **plausible combinations of constraints** relevant to the people who will use the expertise.

For example:

Scenario	What it tests
A low phone battery, a deadline tomorrow and a service closing in 13 minutes	Findability, visible urgency, phone usability, next step and fallback
An interpreter need, missing documents and a person unfamiliar with the system	Meaning, hidden questions, accessible support and reliance limits
Monitored communications, unsafe mail and a need for specialist family-violence support	Safe contact, privacy, delivery route, handover and emergency distinction
A person in regional Victoria without reliable internet or transport	Alternative routes, service access, realistic action and recovery
A frontline worker with incomplete facts and three people waiting	Carrier authority, clear scripts, safe handover and whether the system expects staff to repair missing design

The scenarios should be chosen from evidence, lived expertise, frontline experience and the actual conditions in which the resource will be used.

They should not turn disadvantage into a generic fictional persona.

Adapting the test to the setting

The test asks:

If a person encountered this at 4:47pm (tired, rushed, worried, on a phone, missing information and unable to get through to the first contact), could they still understand what matters, identify a safe next step and recover if the route fails?

The test can be adapted to the actual circumstances of the setting. For example:

- A court or tribunal may test a notice received shortly before a filing deadline.
- A hospital may test discharge instructions used at night, when the usual clinic is closed.
- A government service may test an online application completed on a low-data mobile connection.
- A financial-services provider may test hardship information used by a person who cannot call during business hours.
- An AI-supported tool may test a user who presents an unusual or high-risk situation and needs human escalation.

- A regulator may test a safety warning received by a person who does not speak the dominant language or cannot safely receive messages at home.

The test is not a claim that every pathway must resolve every difficulty immediately. It is a way to make failure conditions visible before they are experienced by people with the least capacity to absorb them.

A pathway does not pass merely because it has a contact number. It passes more credibly when the person can identify:

- Whether the issue is urgent.
- What they can safely rely on.
- What information or documents they need.
- Which channel they can use now.
- What to do if the usual channel is closed or inaccessible.
- Whether there is an alternative, escalation or emergency route.
- What happens if they receive no response, are refused help or discover an error.

Test questions

The following questions can be used as a checklist, scenario-based test or facilitated workshop tool.

Can the person find it?

- Can they find the right resource or service without already knowing the system's language?
- Does it appear in the place they are likely to look first?
- Can they access it on a phone?
- Is it available when they need it, not only through a narrow window they cannot reach?
- Can they use the delivery route safely, privately and in their preferred or required language?
- Is an accessible format, interpreter, support person or alternative route available where needed?
- If the person cannot use the obvious route, can they find another verified route?

Can they see why it matters?

- Does the first screen, heading, letter or sentence tell them what is happening?
- Can they see whether it may apply to them?
- Is the real right, risk, choice, deadline, consequence or opportunity visible?
- Is the urgency clear and proportionate?
- Does the resource give them a truthful reason to continue?
- Does it explain the real stakes without manufacturing fear, shame, urgency or pressure?
- Is there a visible next move?

Can they understand what is happening?

- Can they explain the main point in their own words?
- Can they identify who made the decision, who can explain the process and who can decide the outcome?
- Can they see what conditions, facts or documents may matter?
- Can they distinguish what is known from what remains uncertain?
- Can they identify what the resource can explain and what it cannot decide?

Can they tell what they can rely on?

- Can they distinguish a formal decision from general guidance?
- Can they tell whether the resource has assessed their own situation or merely explained the usual process?
- Can they identify what is established, conditional, uncertain or outside scope?
- Can they see when a qualified person or formal decision-maker must take over?
- Does the resource avoid inviting reliance that the carrier is not authorised to support?

Can they do something realistic?

- Can they find the main point, deadline and next step, with any conditions that materially change their meaning or safety visible alongside them?
- Do they know what to do, by when, through which route, with what documents and using what words?
- Can they realistically take that step with the time, money, privacy, access and safety they have?
- Can they decide when and how they will carry it out?
- Does the pathway account for the likely first obstacle?
- If the person cannot take the step now, does the resource distinguish what can wait from what needs escalation, where it is authorised to do so?

Can they reach help?

- Is there a named service rather than a generic instruction to “seek advice”?
- Is there a phone number, website, counter, appointment route or other real way to contact it?
- Are the hours, eligibility, cost, language and accessibility conditions visible?
- Can the person reach that service before it closes?
- Is the service actually able to do what the resource says it can do?
- If help is unavailable, does the resource make that limit clear and offer a verified fallback where one exists?

Can they question and recover?

- Can they find a route to correct an error, request review, challenge a decision, complain or escalate?
- Do they know what to do if they cannot get through?

- Is there a backup if the first service cannot help?
- Can they continue after a failed referral, a missed step or an unavailable service?
- Does the resource explain where it stops and who takes over?

The professional test

The 4:47pm Test is not only for the person receiving expertise.

It also asks:

Can an ordinary professional safely help the person under these conditions?

A frontline worker may have:

- incomplete facts;
- limited authority;
- a crowded waiting room;
- no access to the decision-maker;
- no ability to guarantee the next service; and
- little time to reconstruct an unclear pathway.

The resource or pathway should not require that professional to:

- guess what the expert meant;
- quietly give advice beyond their authority;
- invent a referral route;
- promise an outcome;
- explain a deadline they cannot see;
- repair a missing handover; or
- become the person's only way of making the system work.

If the resource works only because an exceptional frontline worker repeatedly fills its gaps, the resource itself is incomplete.

The system test

The 4:47pm Test also reveals failures that cannot be fixed by rewriting.

For example:

- The published phone number has a two-hour wait.
- The service closes before a person can call after work.
- The online form does not work on a phone.

- The person cannot safely receive email, text or post.
- The named service has no capacity.
- The form requires documents people commonly do not have.
- There is no interpreter pathway.
- The court, agency or service does not record safe contact preferences.
- The person is sent between services and must repeat their story.
- The review route exists on paper but cannot be found or used.
- The person has no meaningful option because the rule, service or pathway is structurally unavailable.

Those are not failures of understanding, motivation or engagement.

They are **system findings**.

The correct response may be a change to workflow, staffing, service agreements, technology, referral protocols, escalation routes, role clarity or the underlying rule, not merely clearer wording.

The test in practice

Use the 4:47pm Test at three points:

- 1. Before release:** while designing the resource, interaction or pathway.
- 2. During testing:** with intended users and frontline professionals.
- 3. After release:** through monitoring, service data and lived experience.

Before release

Ask the design team:

- What scenario are we testing?
- Which constraints are plausible for this intended audience?
- Can the person find the main point, deadline and next step, with any conditions that materially change their meaning or safety visible alongside them?
- Does the resource assume knowledge, privacy, money, technology or confidence that many people will not have?
- Can the person reach a real service before it closes?
- Is there a verified backup?
- Does the carrier have the authority and practical tools to do what the pathway asks of them?

During testing

Do not ask only:

“Was that clear?”

Give people the resource and ask:

- What happened?
- Does this apply to you or someone like you?
- What would you do next?
- By when?
- Who decides?
- What can you rely on?
- What would you do if the first contact route failed?
- What question would you ask that the resource has not answered?
- What would stop you from taking the next step?
- What information, support or route would make the next step more realistic?

After release

Monitor:

- Missed deadlines.
- Repeat questions.
- Abandoned forms.
- Failed referrals.
- Unsafe assumptions.
- Service wait times.
- Unsuccessful contact attempts.
- Escalation patterns.
- Frontline workarounds.
- Use of backup routes.
- The questions people ask only after the pathway has already failed.
- Whether particular groups are disproportionately unable to reach, use or recover through the pathway.

What the test produces

The 4:47pm Test does not produce a general judgement that something “works” or “does not work.”

It produces evidence, identified failures, ownership and a required response.

Finding	Required response
Meaning, limits or authority are misunderstood	Revise the resource, return it to the relevant source or decision owner and verify it again
The person cannot find or access the expertise	Change the delivery route, entry point, format, timing or access support
The next step, handover or fallback fails	Repair the pathway, confirm the responsible service and retest it
A structural barrier prevents action	Record the system finding, name the responsible owner and escalate it for service, workflow, policy or technology change
A material risk remains unresolved	Record the limitation and decide whether release requires a restriction, stronger warning, alternative carrier, further work or non-release
The carrier is being asked to exceed authority	Redesign the boundary, script and handover; do not solve the gap through unauthorised personalisation
Help is unavailable and no verified fallback exists	Make the absence visible, record it as a system finding and do not present the pathway as complete

A resource does not pass because it contains the right words.

It passes when the team can show what was tested, what failed, what changed, who owns the remaining barrier and when the pathway will be retested.

The final question

The 4:47pm Test reduces the framework to one practical question:

Could the person use this when it mattered, with its meaning, limits and authority intact, and a feasible route forward?

If not, the test identifies what must change in the resource, delivery route or system.

Worked example

From a police-issued safety notice to a pathway that exposes what still fails

Illustrative notice created for this example. This fictional composite is not based on an actual court, police or service form. It is not legal advice, safety advice or a substitute for individual risk assessment, safety planning, legal advice or emergency help.

This example concerns a police-issued family violence safety notice. Police have applied to the Magistrates' Court for a family violence intervention order. The notice records conditions that currently apply to the respondent and gives Maya a court date.

The example shows how a court, legal service, specialist family-violence service or public system might transfer general process expertise without:

- telling a person whether they should apply;
- assessing their risk;
- predicting whether a court will make an order;
- telling them the safest personal course of action; or
- turning general information into individual legal or safety advice.

The example demonstrates a central principle of the method:

A pathway does not pass because it contains a referral. It passes only when the person can actually use the pathway, or when the pathway's absence has been identified, owned and escalated as a system finding.

A family violence intervention order can protect a person, children or property from a family member, partner or former partner. Police can issue safety notices and apply for intervention orders; the court decides whether to make an interim or final order. If a person is in immediate danger, they should call 000.

The situation

Maya has experienced controlling and threatening behaviour from her former partner. They have two children. They share a phone plan, and Maya worries that he can see messages, call records or browser history.

Police attended an incident and issued a family violence safety notice. Police have applied to the court for an intervention order. The notice says:

- a court date has been listed;

- conditions set out in the notice apply to the respondent from the time specified in the notice;
- those conditions continue until the court deals with the matter or the notice otherwise ends.

Maya reads the notice late at night after the children are asleep.

She does not know:

- What document she has received.
- Whether the conditions are legally in force now.
- Who the conditions apply to.
- Where the conditions begin and end.
- What the court date is for.
- Whether she is required to attend.
- What happens if she cannot attend.
- Whether she can bring someone.
- Whether her former partner will see what she tells a legal or support service.
- What happens if she cannot safely receive calls, texts, email or mail.
- How to ask for help without explaining everything from the beginning.
- What to do if the first support route cannot be used.

The issue is not simply that Maya needs a clearer notice.

It is that she needs to establish:

- What the source already knows.
- What needs process clarification.
- What requires legal or specialist family-violence support.
- Whether a usable route to that support actually exists.

The illustrative notice

Family Violence Safety Notice

Police have applied to the Magistrates' Court for a family violence intervention order.

Your court date is [date] at [location].

Conditions set out in this notice apply to the respondent from [time/date] until the court deals with the matter or the notice otherwise ends.

You may be required to attend court.

If you need help, contact the court or seek legal advice.

The notice may be legally accurate.

But it still leaves Maya carrying work that the source, pathway and surrounding expertise could have carried:

- locating the conditions;
- identifying who is bound by them;
- understanding their current legal effect;
- knowing what the court date means;
- working out which service can answer which question;
- establishing confidentiality and contact practices before disclosure;
- and recovering when the available contact route is unsafe or inaccessible.

That is the Paper-to-Person Gap.

Step 0: Orient

What is this expertise for, who is it for, what is at stake and how will it reach Maya?

Question	Answer for this example
Document	Police-issued family violence safety notice
Current legal position recorded by source	Conditions listed in the notice apply to the respondent from the stated commencement time until the court deals with the matter or the notice otherwise ends
Court process	Police have applied to the court for an intervention order; a court date is listed
Function	Record the current legal position, notify Maya of the court process, explain general process and connect her to legal and specialist support
Person	A person who may be experiencing coercion, trauma, monitored communications, limited privacy, caring responsibilities, financial hardship, language barriers, disability or distrust of systems
Stakes	Critical: safety, children's safety, privacy, court attendance, legal protection, housing, finances and future risk
Decision owner	Court or magistrate decides whether to make, vary or continue an intervention order
Carrier	Safety notice, court registry, police, legal service, court-support worker and specialist family-violence service
Carrier limit	Registry staff may explain general court process. They cannot give legal advice or assess Maya's personal safety. A general resource can explain the notice and name support pathways; it cannot tell Maya what is safest or predict the court's outcome
Likely barriers	Monitored phone, unsafe device access, inability to speak privately, children, fear of escalation, limited transport, language, disability, financial hardship and trauma response
What Maya may not know to ask	Which conditions apply now, safe-contact arrangements, attendance, interpreter/accessibility support, court support, confidentiality limits, what to do if a route fails

Orientation finding

The source can answer the legal-status questions.

The court or registry can answer general process and access questions.

Legal and specialist family-violence services may answer different personal questions within their authority.

But the pathway must first establish whether Maya can safely reach any of them.

Produces: an orientation record.

Steps 1–3: Capture, Anchor and Bound

What the source can establish

Maya's question	What the source can answer	What should travel with the notice
What document have I received?	It is a family violence safety notice	"You have received a family violence safety notice."
Who applied to the court?	Police applied	"Police have applied to the court for an intervention order."
Are conditions currently in force?	Yes, if the notice states they apply from the stated date/time	"The conditions in this notice currently apply to the respondent from [time/date]."
Who do the conditions apply to?	The respondent	"The conditions tell the respondent what they must not do."
Where are they recorded?	The conditions section of the notice	"You can find them under [heading] on page [x]."
When do they end?	As stated in the notice: when the court deals with the matter or the notice otherwise ends	"The notice states when the current conditions end or change."
What is the court date for?	The court will deal with the police application	"The court date is when the court will consider the application."
Who decides what happens next?	The court	"Only the court can decide what order, if any, to make."

What cannot be established by the general resource

Question	Why it needs qualified or specialist help
What is safest for Maya personally?	Requires individual safety assessment and Maya's circumstances
What legal option should Maya choose?	Requires individual legal advice
What will the court decide?	Only the court decides
Will the respondent comply with conditions?	Cannot be predicted
Is a particular phone, email or web route safe for Maya?	Depends on individual circumstances and risk
Will a service keep every detail confidential?	Services have different privacy, confidentiality and information-sharing responsibilities

Boundary rule

The resource can explain what the notice records, what process follows and where to seek help. It cannot assess Maya's safety, advise her on her legal position or predict the court's outcome.

Produces: a source map, claim-and-authority map and boundary map.

Step 4: Reach

Can the expertise reach Maya safely, and will she see why it matters?

The rebuilt notice must not assume that phone, text, email, post, shared devices or online forms are safe for every person.

It must also avoid treating Maya's inability to use a contact route as a personal failure.

Rebuilt opening

You have received a family violence safety notice.

Police have applied to the court for a family violence intervention order.

The conditions listed under **[heading]** on page **[x]** currently apply to the respondent from **[time/date]**.

Your court date is **[date]** at **[time]**.

This page explains what the notice means, what the court will decide and where you can get support. It also explains what to do if the contact details in the notice are not safe or possible for you to use.

This opening tells Maya:

- what document she has;
- what currently applies;
- who the conditions bind;
- where to find the conditions;
- what happens next;
- and that unsafe contact is a recognised pathway issue.

It does not promise a court result or tell Maya what is safest.

Draft emergency wording

This wording requires specialist review before publication:

If you are in immediate danger

Call 000.

This page cannot assess what is safest for you. If you cannot use the options listed here safely, tell a specialist family-violence service or another trusted support person that you need help finding a safer way to contact support.

This is deliberately labelled as draft wording requiring specialist review. It is not presented as a finished safety instruction.

Produces: a reach plan, including a specialist-review requirement for emergency and safe-contact wording.

Step 5: Translate

What does Maya need to understand what the notice means and what she can rely on?

The original illustrative notice said:

“A hearing has been listed at the Magistrates’ Court. You may be required to attend. The court may make orders under the Family Violence Protection Act.”

That wording does not answer what document Maya has received, whether conditions are currently in force or who those conditions bind.

Rebuilt explanation

What this notice is

You have received a family violence safety notice.

Police have applied to the court for a family violence intervention order.

What applies now

The conditions listed under **[heading]** on page **[x]** currently apply to the respondent from **[time/date]**.

The conditions tell the respondent what they must not do.

What the court date means

Your court date is **[date]** at **[time]** at **[location]**.

The court will consider the police application. Only the court can decide what order, if any, to make.

Attendance

Check the section headed **[Attendance / Court date]** in your notice for what it says about attendance.

If you are unsure what the notice requires, the court registry can explain general process and attendance arrangements. A legal service can advise you about your own situation.

What this page can help with

This page explains what the notice records, the general court process and where to find legal and specialist family-violence support.

What this page cannot decide

It cannot tell you what the court will decide, what is safest for you personally or what legal option you should take.

Produces: a translated notice and process explainer.

Step 6: Connect

What does Maya need to do, ask, access, question or recover from?

A handover must include:

1. A named service.
2. A route.
3. Access conditions.
4. A backup if the route fails.

The example demonstrates two different outcomes:

- one completed handover; and
- one unresolved pathway failure that must be escalated.

Completed handover: specialist family-violence support

Specialist family-violence support

A specialist family-violence service can help you think about safety, practical support and safer ways to contact court, legal and other services.

1800RESPECT Phone: 1800 737 732 Text: 0458 737 732 Online chat and video call: 1800respect.org.au Available: 24 hours a day, 7 days a week Cost: Free

Before sharing personal details, you can ask:

“Can you explain your privacy and information-sharing arrangements, and how you will contact me?”

You can also ask:

“I have received a family violence safety notice and a court date. I cannot safely use the contact details in the notice. Can you help me understand what support or safer contact options may be available?”

This is a completed handover only after the service owner verifies:

- the contact details;
- hours;
- cost;
- service scope;
- accessibility;
- privacy/information-sharing explanation; and
- the support route’s current capacity.

1800RESPECT provides 24/7 phone, text, online-chat and video-call support. It should not be represented as universally safe for every person or as a substitute for individual safety planning.

Legal-help route: proposed, subject to verification

Legal help

A legal service may be able to give advice about your own situation and explain legal options.

Victoria Legal Aid Phone: 1300 792 387 Website: legalaid.vic.gov.au

Before sharing personal details, you can ask:

“Can you tell me your service hours, whether you can help with this type of matter, what language or accessibility support is available, how you keep information private and what happens if you cannot assist me?”

This is not yet a complete handover in the example because it still requires verification of:

- current operating hours;
- service eligibility;
- capacity;
- interpreter and accessibility arrangements;
- safe-contact options; and
- an alternative if the service cannot assist.

That is intentional. The framework does not allow a service name to stand in for a functioning route.

Victoria Legal Aid provides family-violence intervention-order information and legal assistance subject to its guidelines and available capacity.

Failed route: Maya cannot safely use the court phone or form

Maya cannot safely call from home.

She tries the court’s online form using a borrowed device. The form requires an email address that she cannot safely use. The notice provides no other verified route for someone who cannot call, email or use the form.

At this point, the pathway fails.

The framework does not solve that by writing:

“Use another safe route.”

No verified alternative has been established in this scenario.

The correct result is:

Finding	Status	Required response
Maya cannot safely use the listed phone route	Pathway failure	Court/service owner must provide or verify an alternative safe contact process
Maya cannot use the online form without a safe email address	Pathway failure	Digital and court owners must assess an alternative submission/contact route
No verified alternative route is named	Unresolved material risk	Do not represent the pathway as complete; record owner, escalation and review date
Specialist service may assist with planning	Partial support route	Verify whether it can help connect Maya to a court route without requiring unsafe contact

The resource can honestly say:

If the contact route in this notice is not safe or possible for you to use

This notice does not provide a verified alternative route in every circumstance.

A specialist family-violence service may be able to help you identify safer options for contacting support. The court is responsible for ensuring there is a workable process for people who cannot use the listed contact route.

If no alternative is available, this is a pathway problem that must be raised with the court or responsible service. It is not a failure by you to engage.

This is not a satisfying outcome for Maya. That is the point.

The framework has exposed a system failure rather than disguising it as a personal problem or a completed pathway.

Questions Maya can take to a service

- What do the current conditions in my notice mean?
- Do I need to attend court, and what happens if I cannot?
- Can I bring a support person?
- Can I ask for an interpreter or accessibility support?
- What are your privacy and information-sharing arrangements?
- How will you contact me?
- What can I do if I cannot safely use the phone, email or online form in the notice?
- What support may be available for my children?
- What happens after the court date?

Produces: a usable pathway, a completed specialist-support handover, a proposed legal-help route requiring verification and a documented court-contact system failure.

Step 7: Verify

What must be checked before this pathway can be released?

Owner	Verification task
Court / decision owner	Notice type, current conditions, court date, attendance wording and process explanation
Legal owner	Legal accuracy; no unauthorised advice; correct distinction between notice, process and court decision
Family-violence specialist owner	Emergency wording, safe-contact wording, privacy prompts, referral design and safety assumptions
Operational court owner	Whether an actual alternative contact process exists for people unable to call, email or use the online form
Digital owner	Whether the form requires unsafe assumptions such as a private email address and whether another route exists
Service owners	Current contact details, hours, eligibility, capacity, language/access support and privacy explanations
Translation owner	Whether people can establish legal status, locate conditions, understand court purpose and identify what requires help

The court-contact pathway may not be described as complete until the court or responsible service verifies and publishes an alternative route.

Produces: a verification record identifying both verified elements and unresolved material risks.

Step 8: Test

What must testing establish?

Testing should be trauma-informed, voluntary, supported and appropriately compensated. It should not require people with lived experience to disclose personal details or relive harm.

The resource should be tested with specialist practitioners, court-support workers, legal practitioners and people with relevant lived expertise.

Task-based tests

Task	What it tests
"What document has Maya received?"	Practical meaning
"Do any conditions apply now? Who do they apply to? Where would you find them?"	Current legal status and source use
"What is the court date for, and who decides what happens?"	Process and authority
"What would Maya do if she did not understand whether she had to attend?"	Attendance clarification route
"What would Maya ask before disclosing details to a service?"	Confidentiality and information-sharing understanding
"Who can help with specialist safety support? Who can give legal advice?"	Authority and handover
"What would Maya do if she could not safely call or use the online form?"	Failed-route recovery
"Does the page tell Maya what is safest for her personally?"	Boundary and appropriate reliance
"What would you do if Maya felt in immediate danger tonight?"	Emergency distinction

The test should establish whether people can locate the current legal status, understand attendance, distinguish court process from personal advice, find support and recognise when the court-contact pathway remains unresolved.

Produces: evidence of use, boundary understanding, failed-route navigation and system failure.

Step 9: Maintain

Who owns the unresolved route and what triggers review?

The court or responsible agency owns the unresolved contact-route problem.

The resource should include a maintenance record:

Item	Owner	Trigger
Notice wording and legal status	Court / legal owner	Law, form or process change
Specialist support details	Service owner	Hours, capacity or contact change
Legal-help details	Legal service owner	Eligibility, capacity, hours or referral change
Safe-contact and alternative court route	Court operations and digital owners	User testing, complaint, incident, accessibility finding or service failure
Emergency wording	Specialist family-violence owner	Updated practice guidance or safety review
Whole pathway	Translation owner	Missed dates, failed contacts, repeat questions, complaints or lived-experience feedback

The pathway remains incomplete until the court verifies a workable alternative contact route for people who cannot use the standard phone, email or online form.

That limitation must remain visible.

Produces: a maintenance record with an unresolved-risk owner, escalation pathway and review date.

What this example demonstrates

The notice does not become personalised legal advice or individual safety planning.

The court's authority does not change.

The resource does not promise a particular outcome, absolute confidentiality or a universally safe contact method.

What changes is that the pathway distinguishes:

What Maya can now establish	What still needs qualified help	What remains a system failure
What notice she has received	What is safest for her personally	No verified alternative court contact route
Whether conditions are currently recorded as applying to the respondent	What legal option she should take	A form requiring an unsafe email address
Where to find the conditions	Whether she should attend or what the court is likely to decide	No confirmed fallback when standard channels cannot be used
What the court date is for	How a particular service will handle her details and circumstances	The responsible court/service must design and publish a workable alternative
Who decides the court outcome	Individual legal advice and specialist safety planning	
What questions to ask before disclosure		

This is the point of the method.

It does not make the missing route disappear.

It makes the missing route visible, gives it an owner and stops the system from recording Maya's inability to use it as a failure to engage.

Safe to Say preserves the legal position and boundaries.

Safe to Reach helps Maya establish what the notice means and why it matters.

Safe to Use gives her a route into support, and identifies, rather than hides, the point where the pathway still fails.

Discussion, limitations and evaluation agenda

Proposed contribution

The proposed contribution is not that existing disciplines fail to consider accuracy, authority, accessibility, action, ethics or implementation. Those matters are already addressed, in different ways and to different degrees, across the fields on which the framework draws.

Its proposed contribution is a shared method and accountability structure for making selected responsibilities explicit and connected: preserving accuracy and authority, enabling truthful reach, and supporting practical use across the whole pathway, from source to person, through action, help, contestability and recovery.

How far this integrated framework offers a distinct or effective contribution remains an empirical question. It requires literature review, application and evaluation across regulated settings, with intended users, frontline practitioners, source owners, decision-makers and institutions responsible for the pathways in which expertise is carried.

Limitations

Safe to Say, Safe to Reach, Safe to Use is a proposed framework, not a validated theory, professional standard or assurance scheme. It offers a structured way to identify and discuss failures in the transfer of consequential expertise, but it does not yet establish that applying the framework improves outcomes, reduces harm or produces consistent results across settings.

Its claims should therefore remain proportionate to the current evidence: the framework draws together established concerns from several fields and offers a testable integration. Its distinctiveness, usefulness, reliability, cost and effects require further conceptual work, co-design, application and empirical evaluation.

Evidence limitations

It is not yet empirically validated

The framework has not yet been tested through controlled studies, multi-site implementation, longitudinal evaluation or independent replication. It should not be presented as proven to improve comprehension, service access, safety, fairness, trust, outcomes or organisational coordination.

The framework's three conditions may be intuitively persuasive while still proving difficult to distinguish, measure or apply reliably. For example:

- Teams may disagree about what counts as a material qualification or sufficient uncertainty.

- “Truthful relevance” may be interpreted differently by subject-matter experts, communications teams and intended users.
- A feasible action for one person may be inaccessible or unsafe for another.
- A successful handover may depend on service capacity, funding, geography or eligibility criteria outside the control of the team using the framework.
- The presence of a review or complaints process does not demonstrate that it is accessible, fair or effective.

Evaluation of complex interventions is inherently difficult because outcomes depend on context, implementation, interactions among components and system conditions. The UK Medical Research Council and NIHR framework distinguishes development, feasibility, evaluation and implementation as separate but connected phases, rather than assuming that a plausible intervention is already effective. Context-dependent interventions may also be difficult to standardise or generalise across settings.⁴

It does not establish causal effects

If a team applies the framework and outcomes improve, the improvement may result from other changes: additional staffing, clearer policy, increased service capacity, legal reform, new funding, better technology, public attention, staff training or changed user circumstances.

Conversely, a well-applied framework may not improve outcomes where the decisive problem is outside the expertise-transfer pathway, for example, an unavailable service, an unaffordable requirement, a statutory exclusion, an unjust decision rule or a shortage of authorised professionals.

The framework should therefore be evaluated using mixed methods and, where feasible, comparison groups, staged rollout, before-and-after measures or other designs that can distinguish implementation effects from broader change.

Its concepts require operational definitions

Several core concepts need further definition and testing before they can support strong comparison across sectors:

- **Consequential expertise:** What threshold of likely impact brings a resource, interaction or pathway within scope?
- **Regulated expertise:** Which forms of formal authority, professional duty or evidence standard are sufficient?
- **Safe:** Is this a threshold, a relative judgement or a risk-management standard?
- **Appropriate reliance:** What reliance is reasonable for a particular source, carrier and context?
- **Truthful relevance:** How should teams distinguish warranted urgency from manipulative framing?
- **Feasible action:** What level of effort, cost, delay or burden makes a stated next step unrealistic?
- **Qualified handover:** What evidence is needed to show that a referral has become a workable transfer?
- **Recovery:** Which forms of correction, review, complaint, escalation or alternative help are necessary in different settings?

Until these concepts are operationalised, different teams may use the same words while applying different thresholds.

Scope limitations

It cannot solve structural failure by communication alone

The framework can reveal that a pathway is inaccessible, under-resourced, unjust or impossible to complete. It cannot itself create service capacity, change legislation, remove poverty, resolve housing shortages, provide specialist advice, shorten statutory delays or correct discriminatory institutional rules.

A clear explanation of an inaccessible service remains an explanation of an inaccessible service. A well-designed referral does not guarantee that the next service can accept the person. A properly qualified notice cannot make an unfair decision fair.

The framework must therefore not be used to recast structural failures as communication problems or to imply that a person would succeed if only they understood the system better.

It is not a substitute for specialised obligations

The framework does not replace:

- Legal, clinical, financial, migration, tax, therapeutic or other specialist judgement.
- Professional regulation, scope-of-practice rules or ethical codes.
- Legal compliance, procedural fairness or statutory decision-making duties.
- Accessibility, privacy, safeguarding, cybersecurity or equality assessments.
- Evidence review, research governance, clinical governance or quality assurance.
- Service design, workforce planning, commissioning or operational management.
- AI safety, model validation, data governance, assurance or human-oversight requirements.
- Individual advocacy, representation, informed consent or safety planning.

A resource may meet the framework's process questions while still be legally defective, clinically unsafe, inaccessible, unfair or technically unreliable. Conversely, compliance with a specialist regime does not necessarily establish that expertise has reached and supported the person safely.

It may not suit every context equally

The framework is most suited to expertise that is consequential and regulated, especially where people must understand, judge, act, seek help or challenge an outcome.

It may be less useful, or require a lighter application, for low-consequence information, routine internal communication, purely technical exchanges between specialists or situations in which no practical action, reliance or service pathway is involved.

It also needs careful adaptation in settings involving:

- Acute emergencies, where speed may limit extensive testing or layered explanation.
- Crisis, violence, child-protection or safeguarding matters, where disclosure, contact and handover may create safety risks.

- Complex collective decisions, where the “person” is a community, organisation or public rather than an individual.
- Cultural, linguistic or legal contexts that differ substantially from the framework’s initial assumptions.
- Matters in which a person’s preferred action conflicts with legal duties, safety requirements or the rights of others.

Adaptation should retain the framework’s underlying functions (integrity, truthful reach, practical use and recovery) while allowing implementation to fit local conditions. Implementation guidance similarly emphasises intentional adaptation, stakeholder involvement, evaluation and transparent reporting of changes, rather than treating adaptation as failure.⁵

Implementation limitations

It may add work and delay

Applying the framework well may require cross-functional review, user research, accessibility testing, pathway mapping, service-capacity checks, governance discussion, monitoring and ongoing updates. These activities require time, expertise, funding and leadership support.

In under-resourced organisations, the framework could become:

- Another checklist without power to change the underlying pathway.
- An additional approval stage that delays urgent communication.
- A burden placed on frontline staff without operational authority to fix service failures.
- A documentation exercise that rewards evidence of process rather than real-world improvement.
- A risk-management tool that prioritises institutional defensibility over a person’s practical needs.

Proportionate application is therefore essential. A high-stakes clinical instruction, benefit decision, safety warning or AI-supported triage pathway warrants more intensive review than a low-consequence informational update. The framework should not imply that every item needs the same process.

Cross-functional accountability may remain difficult

The framework assumes that institutions can identify and coordinate source owners, communications teams, accessibility specialists, frontline workers, digital teams, service managers, complaints functions and governance owners.

In practice, authority may be fragmented across agencies, contractors, jurisdictions, funding arrangements and technology providers. The team that writes a notice may not control the service capacity behind its referral. The organisation hosting a chatbot may not control the policy it explains. A regulator may issue guidance without controlling how local services implement it.

Naming an accountable owner may expose a real accountability gap without resolving it. This is useful diagnostically, but it can also create conflict or inertia where no actor has both the responsibility and resources to act.

Fidelity and adaptation create a tension

The framework is intended to be adaptable across sectors. But extensive adaptation may make it impossible to know whether teams are applying the same framework or whether observed effects are comparable.

Too much standardisation would be equally problematic: it could ignore local law, culture, service structures, language, user conditions and risk. Implementation research recognises this tension between fidelity to core functions and context-appropriate adaptation. Transparent documentation of what was adapted, why, by whom and with what effect is needed if the framework is to develop an evidence base.⁶

Risks of misuse

“Safe” could be mistaken for a guarantee

The word “safe” may imply certainty or an assurance level that no communication, service or institution can provide. A resource can be carefully designed and still be misunderstood, missed, misapplied or overtaken by changing circumstances.

“Safe to Say, Safe to Reach, Safe to Use” should therefore be understood as a disciplined inquiry into foreseeable risks and responsibilities, not a certification that no harm will occur, that all users will succeed or that all outcomes are fair.

The framework should not be used as a shield against accountability after harm. A completed assessment does not prove that a system acted reasonably, nor should it displace legal rights, complaint mechanisms, professional duties or independent review.

It could become paternalistic

A framework concerned with safe use can be misused to justify withholding information, narrowing choices, increasing gatekeeping or treating people as incapable of making their own decisions.

The framework must support informed judgement and agency, not institutional control. It should distinguish between helping a person understand risk and steering them toward the institution’s preferred outcome. It should also avoid assuming that professionals always know what is best for people, particularly where lived experience, culture, community knowledge and personal priorities matter.

The principle of “appropriate reliance” should not become a rationale for excessive disclaimers, refusal to explain or unnecessary barriers to information.

It could shift responsibility back to the person

The framework explicitly seeks to identify system failure. Yet it could be misapplied if institutions use it to ask whether a person “used the information correctly” rather than asking whether the system created realistic conditions for understanding, action and recovery.

Measures such as form completion, referral uptake, website engagement or reduced contact volume can be misleading if they are treated as proof of success without examining comprehension, informed choice, access barriers, service capacity, equity and outcomes.

A person's inability to complete a pathway may be evidence of an inaccessible or unjust system, not a lack of motivation, literacy, capability or compliance.

It could privilege measurable failures

Institutions may be more able to measure web clicks, call volumes, form completion and response times than hidden harms such as fear, loss of trust, stigma, coercion, exclusion, cultural unsafety or the decision not to seek help.

Qualitative evidence, lived experience, complaint narratives, community feedback and frontline observation are therefore essential. A pathway should not be judged successful only because people complete it quickly or because complaints decline.

AI-specific limitations

Where used with AI-supported systems, the framework cannot substitute for technical validation, model governance or human oversight.

An AI system may present information fluently, confidently and in a personalised style that encourages users to treat it as authoritative or individually applicable. Research on public-sector decision-making warns of automation bias: undue deference to automated systems even where warning signs or contrary information exist. It also notes risks that algorithmic systems can reproduce systemic bias and be difficult to diagnose because of opacity or complexity.⁷

The framework can help make source, limits, authority boundaries, escalation and contestability more visible to users. It cannot establish that an AI model is accurate, robust, unbiased, secure, privacy-preserving or appropriate for a particular high-stakes use.

AI-specific application should therefore include, at minimum:

- Independent source and output validation.
- Clear limits on what the system may explain, recommend, assess or decide.
- Monitoring for error, drift, inequitable performance and inappropriate reliance.
- Accessible routes to qualified human support and meaningful review.
- Procedures to correct harmful outputs and preserve records where accountability requires it.
- Governance that considers automation bias, transparency, privacy, security, bias and differential impact.

Assumptions and external conditions

Applying the framework depends on more than the quality of its questions. It requires access to sufficiently trustworthy sources; participation by relevant source, professional and operational owners; involvement of intended users and frontline practitioners; and organisational willingness to make material conditions, uncertainty and authority visible.

It also requires someone to own the response to identified failures. A source map, user test or failed referral record has limited practical value if no one can change the resource, repair the pathway or escalate the underlying barrier. Maintenance requires continuing capacity, not only an owner named at release.

These are implementation conditions to examine, not assumptions that every organisation already meets.

The framework operates within circumstances it cannot control alone: the fairness of the underlying rule, service capacity and funding, geographic availability, safety, money, time, privacy, language, documents, technology and historical reasons for distrust. Some barriers may also persist because institutional arrangements preserve discretion or transfer work and risk onto people. Those mechanisms require investigation rather than an assumption of deliberate obstruction.

Where these conditions prevent use, clearer communication may reveal the problem without resolving it. The framework should identify the barrier, its consequences and the responsible owner, while distinguishing what can be repaired through design from what requires service, governance or policy change.

Implications for use

These limitations support a cautious use position.

The framework should be used as:

- A shared language for diagnosing transfer failures.
- A structured prompt for design, review, governance and improvement.
- A way to surface dependencies among accuracy, reach, usability, handovers and recovery.
- A basis for proportionate testing and empirical evaluation.
- A complement to, not replacement for, established professional, legal, accessibility, service-design and assurance processes.

It should not be used as:

- Proof that a resource, service, pathway or AI system is safe.
- A substitute for specialist judgement or legal and professional obligations.
- A way to avoid investment in service capacity, accessibility, staffing or structural reform.
- A one-time checklist detached from monitoring, lived experience and correction.
- A basis for blaming people when systems remain difficult or impossible to use.

The limitation is also the framework's next task: to specify its concepts, test its application, document adaptations, assess outcomes and determine where it genuinely improves the transfer of consequential expertise, and where existing methods already do the work better.

Evaluation agenda

The framework proposes a relationship between institutional practice and people's ability to use consequential expertise. That relationship requires testing.

Its proposed theory of change is that, when institutions preserve source meaning and authority, design truthful reach, make necessary reasoning visible and connect expertise to feasible pathways, people will be better supported to recognise relevance, understand their situation,

judge what they can rely on and find an appropriate next step. Verification, user testing and monitoring should also make resource and system failures more visible and assign responsibility for addressing them.

Each link is a proposition. Producing a verified resource does not establish improved understanding; improved understanding does not establish feasible action; and identifying a pathway does not establish that a service can receive the person.

What should be evaluated

Level	Proposed indicators
Outputs	Verified sources and claims; visible authority and boundaries; reach plan; translated resource; documented handover and recovery routes; accountable owners and maintenance arrangements
Immediate outcomes	People encounter the expertise, recognise its relevance, understand practical meaning and distinguish what is established, conditional, uncertain or outside scope
Practical outcomes	Feasible task completion, appropriate help-seeking, connection to suitable services and recovery after a failed route
Institutional outcomes	Clearer ownership, repaired pathways, maintained resources and better recognition of failures previously recorded as individual non-engagement
Longer-term contribution	Informed judgement, warranted trust and more equitable use of expertise, assessed alongside other influences

Evaluation should examine accuracy and professional boundaries alongside reach and use. An easier task is not an improvement if it depends on losing a material condition or inviting unsupported reliance.

How to test the proposed effects

Begin with a defined resource or pathway, intended population and observable problem. Record the existing conditions and, where feasible, compare the redesigned pathway with the existing one.

Use complementary forms of evidence:

- Substantive review against agreed sources and carrier authority.
- Observation of how people find and interpret the resource.
- Explanations in participants' own words.
- Scenario-based tests of judgement and appropriate reliance.
- Realistic tasks, including failed-contact and recovery scenarios.
- Operational evidence about service access, referral completion and pathway failures.
- Accounts from intended users and frontline professionals.

Select established measures where they fit the outcome and population. Confidence measures can complement evaluation, but confidence is not equivalent to demonstrated capability or completed action. Likewise, fewer calls, complaints or abandoned forms should not automatically count as success: they may reflect either improved use or reduced access.

Document who was included, what conditions were tested, what adaptations were made and whose needs remain unexamined. Examine differences between groups rather than relying only on average results.

What remains to be established

Evaluation should address three questions:

- 1. Conceptual value:** Does the framework help teams identify and connect responsibilities or failures they previously considered separately?
- 2. Practical feasibility:** Can teams apply it with available time, skills, authority and resources?
- 3. Effectiveness:** Does its application improve responsible reach, understanding, warranted reliance, feasible action, help or recovery?

It should also examine unintended effects, including information overload, additional review work, maintenance burden, delayed release and checklist-driven approval.

Neither capacity outcomes nor behavioural outcomes should be treated as established before evaluation. The framework states what it intends to support; pilots must investigate whether, for whom and under what conditions it does so.

A useful first application is one consequential pathway with a visible failure, accountable owners and a manageable evaluation plan. Record what changed, what remained unresolved and what requires institutional action.

A small, well-observed pilot is more valuable than a large rewrite programme that is never tested in practice.

Conclusion and key points

Conclusion

Safe to Say, Safe to Reach, Safe to Use addresses a recurring failure in consequential systems: expertise can be accurate, approved and professionally responsible at its source, yet still fail the person whose rights, safety, health, finances, access to services or future choices depend on being able to use it.

This paper names that failure the Paper-to-Person Gap. It proposes a framework, ten-step method and 4:47pm Test for examining whether regulated expertise can travel safely from source to person: through a carrier, across a service pathway and into real-world judgement and action.

The framework's contribution is not to replace plain language, legal design, knowledge translation, service design, accessibility, professional responsibility, risk communication, health literacy, behavioural science or AI governance. It brings selected responsibilities from these fields into one shared accountability structure for consequential expertise.

Its central proposition is straightforward:

Expertise is not safely transferred merely because it has been approved, published, sent, explained, made accessible or referred onward. It must also remain safe to say, safe to reach and safe to use.

The framework is proposed rather than validated. Its value and limits should be tested through literature review, application and empirical evaluation across different regulated settings, with intended users, frontline professionals, source owners, decision-makers and institutions that own the pathway.

Key points

1. Correct at the source is not enough

A legally accurate notice, evidence-based guideline, approved policy, professional explanation or AI-supported response may still fail if a person cannot find it in time, see why it matters, understand its limits, judge what they can rely on, take a feasible next step or obtain help when the first route fails.

The relevant unit of analysis is not only the document, message or tool. It is the full pathway:

Source → carrier → person → judgement → action → help, review or recovery

2. The three conditions must hold together

Safe to Say, Safe to Reach and Safe to Use are interdependent conditions, not sequential stages or alternative options.

Safe to Say requires accuracy, traceable source, appropriate qualification, visible uncertainty, proper authority and clear boundaries on what a person may rely on.

Safe to Reach requires timely, accessible and discoverable delivery, together with truthful relevance and attention proportionate to the evidence and practical choices available.

Safe to Use requires practical meaning, visible conditions and limits, feasible next steps, authorised handovers and routes to correction, review, complaint, escalation or recovery.

A resource that satisfies only one or two conditions remains vulnerable to failure. Clear but unauthorised content is unsafe. Accurate but unreachable content is insufficient. Engaging content without a feasible pathway can create attention without support.

3. A referral is not a complete handover

A system cannot assume that a person can carry the work of coordination after being given a phone number, web link or generic instruction to “seek independent advice.”

Where an individual assessment, decision, representation or specialist intervention is required, the responsible institution should make the boundary visible and provide a practical route to the appropriate source of help. That includes identifying what the next service can do, how it can be reached, what information may be needed and what options remain if the person cannot use that route.

A referral becomes a more credible handover when the person can understand where they are going, why, what will happen next and what to do if the service cannot help.

4. Failure often belongs to the system

When people do not understand, act, complete a process or reach help, the cause may be more than individual knowledge, motivation or effort.

It may be an authority problem, a timing problem, inaccessible delivery, missing context, unavailable service capacity, a failed handover, an opaque digital system, inadequate escalation, an unworkable deadline or an underlying policy failure.

The framework asks institutions to make those points visible, assign them to an accountable owner and distinguish preventable transfer failures from limits that require broader reform.

Next step

The next step is practical evaluation.

The framework should be applied to real examples across law, public administration, health, financial services, social services, research translation, regulation and AI-supported public-facing systems. Evaluation should examine whether the method:

- Identifies failures that existing quality, legal, communications or service-design processes miss.

- Improves people's ability to recognise relevance, understand limits, identify a next step and reach appropriate help.
- Reduces avoidable delay, abandonment, failed referrals, inappropriate reliance, repeat contacts and preventable escalation.
- Helps institutions assign responsibility across legal, policy, content, accessibility, operational, digital and governance functions.
- Operates proportionately in low- and high-consequence settings without creating unnecessary process or diluting professional boundaries.

The 4:47pm Test offers a practical starting point. Take one consequential letter, form, website, service pathway or AI-supported interaction and ask:

If a person encountered this late in the day, under stress, on a phone, without all the information they need and unable to reach the first service they try, could they still tell what matters, what they can rely on, what to do next and how to recover if the route fails?

References

Legal design and people-centred justice

Chung, C., & Kim, S. H. (2022). A systematic literature review of legal design. *The Design Journal*, 25(1), 3–27.

Hagan, M. (2018). A human-centered design approach to access to justice: Generating new prototypes and hypotheses for intervention to make courts user-friendly. *Indiana Journal of Law and Social Equality*, 6(2), 199–239. <https://www.repository.law.indiana.edu/ijlse/vol6/iss2/2/>

Hagan, M. (2019). Participatory design for innovation in access to justice. *Daedalus*, 148(1), 120–127. https://doi.org/10.1162/daed_a_00531

Malik, N. (2026). Human-centered design and people-centered justice. American Bar Foundation.

Organisation for Economic Co-operation and Development. (2021). OECD framework and good practice principles for people-centred justice. OECD Publishing. <https://doi.org/10.1787/cdc3bde7-en>

Organisation for Economic Co-operation and Development. (2025). Toolkit for access to justice and people-centred justice systems. OECD Publishing. <https://doi.org/10.1787/aecf7f78-en>

Access to justice, legal capability and expertise

Balmer, N. J., Pleasence, P., McDonald, H. M., & Sandefur, R. L. (2024). Public understanding of law survey: Volume 2. Victoria Law Foundation.

Organisation for Economic Co-operation and Development, & Open Society Foundations. (2019). Legal needs surveys and access to justice. OECD Publishing. <https://www.oecd.org/gov/legal-needs-surveys-and-access-to-justice-g2g9a36c-en.htm>

Pleasence, P., & Balmer, N. J. (2019). Development of a general legal confidence scale: A first implementation of the Rasch measurement model in empirical legal studies. *Journal of Empirical Legal Studies*.

Sandefur, R. L. (2015). Elements of professional expertise: Understanding relational and substantive expertise through lawyers' impact. *American Sociological Review*, 80(5), 909–933. <https://doi.org/10.1177/0003122415601157>

Plain language, information design and cognition

Assy, R. (2011). Can the law speak directly to its subjects? The limitation of plain language. *Journal of Law and Society*, 38(3), 376–404. <https://doi.org/10.1111/j.1467-6478.2011.00543.x>

International Organization for Standardization. (2023). Plain language, Part 1: Governing principles and guidelines (ISO Standard No. 24495-1:2023). <https://www.iso.org/standard/78907.html>

International Organization for Standardization. (2025). Plain language, Part 2: Legal communication (ISO Standard No. 24495-2:2025). <https://www.iso.org/standard/85774.html>

Stoll, M., Kerwer, M., Lieb, K., & Chasiotis, A. (2022). Plain language summaries of health research: A systematic review. PLOS ONE, 17(8), e0272629. <https://doi.org/10.1371/journal.pone.0272629>

Sweller, J., van Merriënboer, J. J. G., & Paas, F. (2019). Cognitive architecture and instructional design: 20 years later. Educational Psychology Review, 31, 261–292. <https://doi.org/10.1007/s10648-019-09465-5>

Information behaviour and everyday information seeking

Dervin, B. (1998). Sense-making theory and practice: An overview of user interests in knowledge seeking and use. Journal of Knowledge Management, 2(2), 36–46. <https://doi.org/10.1108/13673279810249369>

Kuhlthau, C. C. (1991). Inside the search process: Information seeking from the user's perspective. Journal of the American Society for Information Science, 42(5), 361–371.

Lloyd, A. (2010). Information literacy landscapes: Information literacy in education, workplace and everyday contexts. Chandos Publishing.

Savolainen, R. (1995). Everyday life information seeking: Approaching information seeking in the context of “way of life.” Library & Information Science Research, 17(3), 259–294. [https://doi.org/10.1016/0740-8188\(95\)90048-9](https://doi.org/10.1016/0740-8188(95)90048-9)

Wilson, T. D. (1999). Models in information behaviour research. Journal of Documentation, 55(3), 249–270. <https://doi.org/10.1108/EUM00000000007145>

Wilson, T. D. (2000). Human information behavior. Informing Science, 3(2), 49–56. <https://doi.org/10.28945/576>

Human-centred and service design

Design Council. (2019). Framework for innovation: Design Council's evolved Double Diamond. <https://www.designcouncil.org.uk/our-resources/framework-for-innovation/>

International Organization for Standardization. (2019). Ergonomics of human-system interaction, Part 210: Human-centred design for interactive systems (ISO Standard No. 9241-210:2019). <https://www.iso.org/standard/77520.html>

Knowledge translation and implementation science

Aarons, G. A., Hurlburt, M., & Horwitz, S. M. (2011). Advancing a conceptual model of evidence-based practice implementation in public service sectors. Administration and Policy in Mental Health and Mental Health Services Research, 38(1), 4–23. <https://doi.org/10.1007/s10488-010-0327-7>

Damschroder, L. J., Reardon, C. M., Widerquist, M. A. O., & Lowery, J. (2022). The updated Consolidated Framework for Implementation Research based on user feedback. Implementation Science, 17, Article 75. <https://doi.org/10.1186/s13012-022-01245-0>

Field, B., Booth, A., Ilott, I., & Gerrish, K. (2014). Using the Knowledge to Action Framework in practice: A citation analysis and systematic review. *Implementation Science*, 9, Article 172. <https://doi.org/10.1186/s13012-014-0172-2>

Graham, I. D., Logan, J., Harrison, M. B., Straus, S. E., Tetroe, J., Caswell, W., & Robinson, N. (2006). Lost in knowledge translation: Time for a map? *Journal of Continuing Education in the Health Professions*, 26(1), 13–24. <https://doi.org/10.1002/chp.47>

Moullin, J. C., Dickson, K. S., Stadnick, N. A., Rabin, B., & Aarons, G. A. (2019). Systematic review of the Exploration, Preparation, Implementation, Sustainment framework. *Implementation Science*, 14, Article 1. <https://doi.org/10.1186/s13012-018-0842-6>

Proctor, E. K., Silmere, H., Raghavan, R., Hovmand, P., Aarons, G. A., Bunger, A. C., Griffey, R., & Hensley, M. (2011). Outcomes for implementation research: Conceptual distinctions, measurement challenges, and research agenda. *Administration and Policy in Mental Health and Mental Health Services Research*, 38(2), 65–76. <https://doi.org/10.1007/s10488-010-0319-7>

Administrative burden

Halling, A., & Bækgaard, M. (2024). Administrative burden in citizen–state interactions: A systematic literature review. *Journal of Public Administration Research and Theory*, 34(2), 180–195.

Moynihan, D. P., Herd, P., & Harvey, H. (2015). Administrative burden: Learning, psychological, and compliance costs in citizen–state interactions. *Journal of Public Administration Research and Theory*, 25(1), 43–69. <https://doi.org/10.1093/jopart/muu009>

Health literacy and checking understanding

Agency for Healthcare Research and Quality. (2024). AHRQ health literacy universal precautions toolkit (3rd ed.). <https://www.ahrq.gov/health-literacy/improve/precautions/index.html>

Ha Dinh, T. T., Bonner, A., Clark, R., Ramsbotham, J., & Hines, S. (2016). The effectiveness of the teach-back method on adherence and self-management in health education for people with chronic disease: A systematic review. *JBIR Database of Systematic Reviews and Implementation Reports*, 14(1), 210–247. <https://doi.org/10.11124/jbisrir-2016-2296>

Nutbeam, D., & Lloyd, J. E. (2021). Understanding and responding to health literacy as a social determinant of health. *Annual Review of Public Health*, 42, 159–173. <https://doi.org/10.1146/annurev-publhealth-090419-102529>

Sørensen, K., Van den Broucke, S., Fullam, J., Doyle, G., Pelikan, J., Slonska, Z., Brand, H., & the HLS-EU Consortium Health Literacy Project European. (2012). Health literacy and public health: A systematic review and integration of definitions and models. *BMC Public Health*, 12, Article 80. <https://doi.org/10.1186/1471-2458-12-80>

Talevski, J., Wong Shee, A., Rasmussen, B., Kemp, G., & Beauchamp, A. (2020). Teach-back: A systematic review of implementation and impacts. *PLOS ONE*, 15(4), e0231350. <https://doi.org/10.1371/journal.pone.0231350>

Risk and uncertainty communication

Bonner, C., Trevena, L. J., Gaissmaier, W., et al. (2021). Current best practice for presenting probabilities in patient decision aids: Fundamental principles. *Medical Decision Making*, 41(7), 821–833. <https://doi.org/10.1177/0272989X211048412>

Kerr, J. R., Schneider, C. R., Recchia, G., Dryhurst, S., Sahlin, N. E., Dufouil, C., & van der Linden, S. (2022). Transparent communication of evidence does not undermine public trust in science. *PNAS Nexus*, 1(3), pgac135. <https://doi.org/10.1093/pnasnexus/pgac135>

Trevena, L. J., Zikmund-Fisher, B. J., Edwards, A., Gaissmaier, W., Galesic, M., Han, P. K. J., King, J., Lawson, M. L., Linder, S. K., Lipkus, I., Ozanne, E., Peters, E., Timmermans, D., & Woloshin, S. (2013). Presenting quantitative information about decision outcomes: A risk communication primer for patient decision aid developers. *BMC Medical Informatics and Decision Making*, 13(Suppl. 2), S7. <https://doi.org/10.1186/1472-6947-13-S2-S7>

van der Bles, A. M., van der Linden, S., Freeman, A. L. J., Mitchell, J., Galvao, A. B., Zaval, L., & Spiegelhalter, D. J. (2019). Communicating uncertainty about facts, numbers and science. *Royal Society Open Science*, 6(5), 181870. <https://doi.org/10.1098/rsos.181870>

van der Bles, A. M., van der Linden, S., Freeman, A. L. J., & Spiegelhalter, D. J. (2020). The effects of communicating uncertainty on public trust in facts and numbers. *Proceedings of the National Academy of Sciences*, 117(14), 7672–7683. <https://doi.org/10.1073/pnas.1913678117>

Community engagement and public-interest communication

Covello, V. T. (2003). Best practices in public health risk and crisis communication. *Journal of Health Communication*, 8(Suppl. 1), 5–8. <https://doi.org/10.1080/713852406>

Reynolds, B., & Seeger, M. W. (2005). Crisis and emergency risk communication as an integrative model. *Journal of Health Communication*, 10(1), 43–55. <https://doi.org/10.1080/10810730590904571>

World Health Organization. (2017). Communicating risk in public health emergencies: A WHO guideline for emergency risk communication policy and practice. <https://www.who.int/publications/i/item/9789241550208>

World Health Organization. (2021). WHO competency framework: Risk communication and community engagement.

Behaviour change and intervention evidence

Maier, M., Bartoš, F., & Stanley, T. D. (2022). No evidence for nudging after adjusting for publication bias. *Proceedings of the National Academy of Sciences*, 119(31), e2200300119. <https://doi.org/10.1073/pnas.2200300119>

Mertens, S., Herberz, M., Hahnel, U. J. J., & Brosch, T. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *Proceedings of the National Academy of Sciences*, 119(1), e2107346118. <https://doi.org/10.1073/pnas.2107346118>

Michie, S., van Stralen, M. M., & West, R. (2011). The behaviour change wheel: A new method for characterising and designing behaviour change interventions. *Implementation Science*, 6, Article 42. <https://doi.org/10.1186/1748-5908-6-42>

Shared decision-making and decision support

National Institute for Health and Care Excellence. (2021). Standards framework for shared decision-making support tools, including patient decision aids.

<https://www.nice.org.uk/corporate/ecd8/resources/standards-framework-for-shareddecisionmaking-support-tools-including-patient-decision-aids-pdf-1124019137221>

Stacey, D., Lewis, K. B., Smith, M., Carley, M., Volk, R. J., Douglas, E. E., & Pacheco-Bueckert, L. (2024). Decision aids for people facing health treatment or screening decisions. *Cochrane Database of Systematic Reviews*, 2024(1), CD001431. <https://doi.org/10.1002/14651858.CD001431.pub6>

Procedural and administrative justice

Adler, M. (Ed.). (2010). *Administrative justice in context*. Hart Publishing.

Blader, S. L., & Tyler, T. R. (2003). A four-component model of procedural justice: Defining the meaning of a “fair” process. *Personality and Social Psychology Bulletin*, 29(6), 747–758.

<https://doi.org/10.1177/0146167203029006007>

Lind, E. A., & Tyler, T. R. (1988). *The social psychology of procedural justice*. Plenum Press.

Mashaw, J. L. (1983). *Bureaucratic justice: Managing social security disability claims*. Yale University Press.

Tyler, T. R. (2006). *Why people obey the law*. Princeton University Press.

Human factors, safety and handovers

Agency for Healthcare Research and Quality. (n.d.). Tool: Handoff. TeamSTEPPS.

<https://www.ahrq.gov/teamstepps-program/curriculum/communication/tools/handoff.html>

Agency for Healthcare Research and Quality. (n.d.). Handoffs. Patient Safety Network.

<https://psnet.ahrq.gov/primer/handoffs>

Hollnagel, E. (2014). *Safety-I and Safety-II: The past and future of safety management*. Ashgate.

Manser, T., & Foster, S. (2011). Effective handover communication: An overview of research and improvement efforts. *Best Practice & Research Clinical Anaesthesiology*, 25(2), 181–191.

<https://doi.org/10.1016/j.bpa.2011.02.006>

Starmer, A. J., Spector, N. D., Srivastava, R., Allen, A. D., Landrigan, C. P., Sectish, T. C., et al. (2014). Changes in medical errors after implementation of a handoff program. *New England Journal of Medicine*, 371(19), 1803–1812. <https://doi.org/10.1056/NEJMsa1405556>

Wears, R. L., Hollnagel, E., & Braithwaite, J. (Eds.). (2015). *Resilient health care, volume 2: The resilience of everyday clinical work*. Ashgate.

Trauma-informed practice and cultural safety

Curtis, E., Jones, R., Tipene-Leach, D., Walker, C., Loring, B., Paine, S.-J., & Reid, P. (2019). Why cultural safety rather than cultural competency is required to achieve health equity: A literature review and recommended definition. *International Journal for Equity in Health*, 18, Article 174.

<https://doi.org/10.1186/s12939-019-1082-3>

Dudgeon, P., Milroy, H., & Walker, R. (Eds.). (2014). *Working together: Aboriginal and Torres Strait Islander mental health and wellbeing principles and practice*. Commonwealth of Australia.

Ramsden, I. (2002). *Cultural safety and nursing education in Aotearoa and Te Waipounamu* (Doctoral thesis, Victoria University of Wellington).

Substance Abuse and Mental Health Services Administration. (2014). SAMHSA's concept of trauma and guidance for a trauma-informed approach (HHS Publication No. SMA 14-4884). U.S.

Department of Health and Human Services. <https://store.samhsa.gov/sites/default/files/sma14-4884.pdf>

Epistemic justice and co-production

Beresford, P. (2002). User involvement in research and evaluation: Liberation or regulation? *Social Policy and Society*, 1(2), 95–105. <https://doi.org/10.1017/S1474746402000222>

Dotson, K. (2011). Tracking epistemic violence, tracking practices of silencing. *Hypatia*, 26(2), 236–257. <https://doi.org/10.1111/j.1527-2001.2011.01177.x>

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.

Ocloo, J., Garfield, S., Franklin, B. D., & Dawson, S. (2021). Exploring the theory, barriers and enablers for patient and public involvement across health, social care and patient safety: A systematic review of reviews. *BMJ Quality & Safety*, 30(2), 176–190. <https://doi.org/10.1136/bmjqs-2018-008263>

Oliver, K., Kothari, A., & Mays, N. (2019). The dark side of coproduction: Do the costs outweigh the benefits for health research? *Health Research Policy and Systems*, 17, Article 33.

<https://doi.org/10.1186/s12961-019-0432-3>

Accessibility, language access and digital inclusion

Flores, G. (2005). The impact of medical interpreter services on the quality of health care: A systematic review. *Medical Care Research and Review*, 62(3), 255–299.

<https://doi.org/10.1177/1077558705275416>

Karliner, L. S., Jacobs, E. A., Chen, A. H., & Mutha, S. (2007). Do professional interpreters improve clinical care for patients with limited English proficiency? A systematic review of the literature. *Health Services Research*, 42(2), 727–754. <https://doi.org/10.1111/j.1475-6773.2006.00629.x>

United Nations. (2006). *Convention on the Rights of Persons with Disabilities*.

<https://www.un.org/disabilities/documents/convention/convoptprot-e.pdf>

World Wide Web Consortium. (2023). *Web Content Accessibility Guidelines (WCAG) 2.2*.

<https://www.w3.org/TR/WCAG22/>

Professional responsibility and information governance

International Organization for Standardization. (2016). Information and documentation, Records management, Part 1: Concepts and principles (ISO Standard No. 15489-1:2016).

<https://www.iso.org/standard/62542.html>

For this category, the final bibliography should also include the specific professional codes, statutory duties, administrative-law requirements, privacy law, procedural-fairness obligations, information-access duties and service standards governing each application. A single generic source cannot establish the authority, confidentiality, competence or scope-of-practice requirements in every legal, justice or public-serving setting.

AI governance and appropriate reliance

Buçinca, Z., Malaya, M. B., & Gajos, K. Z. (2021). To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), Article 188, 1–21. <https://doi.org/10.1145/3449287>

Dahl, M., Magesh, V., Suzgun, M., & Ho, D. E. (2024). Large legal fictions: Profiling legal hallucinations in large language models. *Journal of Legal Analysis*.

Magesh, V., et al. (2025). Hallucination-free? Assessing AI legal research tools. *Journal of Empirical Legal Studies*.

National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce.

<https://doi.org/10.6028/NIST.AI.100-1>

Schemmer, M., Hemmer, P., Nitsche, M., Köhl, N., & Vössing, M. (2023). A meta-analysis of appropriate reliance on AI advice: The role of task difficulty and AI confidence. *Proceedings of the 28th International Conference on Intelligent User Interfaces*.

Vasconcelos, H., Jörke, M., Grunde-McLaughlin, M., Gerstenberg, T., Bernstein, M. S., & Krishnan, M. (2023). Explanations can reduce overreliance on AI systems. *Proceedings of the ACM on Human-Computer Interaction*.

Other sources cited in the text

- Victoria Legal Aid, Community Legal Services Program report 2024–25: https://www.legalaid.vic.gov.au/sites/default/files/2026-02/vla-clsp-2024-25-report_0.pdf
- Victoria Legal Aid, “Over 490,000 legal aid services for Victorians in 2024–25”: <https://www.legalaid.vic.gov.au/over-490000-legal-aid-services-victorians-2024-25>
- Complex-intervention development and evaluation framework: <https://pubmed.ncbi.nlm.nih.gov/34590577/>
- Intervention adaptation guidance, BMJ: <https://www.bmj.com/content/374/bmj.n1679>
- Frameworks for monitoring fidelity and adaptations: <https://rethinkingclinicaltrials.org/chapters/conduct/monitoring-intervention-fidelity-and-adaptations/frameworks-for-monitoring-fidelity-and-adaptations/>

- Automation bias in public-sector decision-making, Journal of Public Administration Research and Theory: <https://academic.oup.com/jpart/article/33/1/153/6524536>

-
1. Victoria Legal Aid, Community Legal Services Program report 2024–25, https://www.legalaid.vic.gov.au/sites/default/files/2026-02/Vla-clsp-2024-25-report_0.pdf ↵
 2. Victoria Legal Aid, “Over 490,000 legal aid services for Victorians in 2024–25”, <https://www.legalaid.vic.gov.au/over-490000-legal-aid-services-victorians-2024-25> ↵
 3. Agency for Healthcare Research and Quality, Tool: Handoff, TeamSTEPPS, <https://www.ahrq.gov/teamstepps-program/curriculum/communication/tools/handoff.html> ↵
 4. Complex-intervention development and evaluation framework, <https://pubmed.ncbi.nlm.nih.gov/34590577/> ↵
 5. Intervention adaptation guidance, BMJ, <https://www.bmj.com/content/374/bmj.n1679> ↵
 6. Frameworks for monitoring fidelity and adaptations, <https://rethinkingclinicaltrials.org/chapters/conduct/monitoring-intervention-fidelity-and-adaptations/frameworks-for-monitoring-fidelity-and-adaptations/> ↵
 7. Automation bias in public-sector decision-making, Journal of Public Administration Research and Theory, <https://academic.oup.com/jpart/article/33/1/153/6524536> ↵

lejacie

**Could the person use this when it
mattered?**

CHAMI RUPASINGHE

Admitted Australian lawyer
(non-practising). Founder,
Lejacie.
chami@lejacie.com · lejacie.com

WHAT THIS DOCUMENT IS

Safe to Say, Safe to Reach, Safe
to Use: an emerging framework
for closing the Paper-to-Person
Gap in regulated expertise.

DRAFT V0.1

30 September 2026. Proposed
practice framework for
discussion, testing and empirical
evaluation.