# Kashi research synthesis

Buyer economics, governance fit, and claim discipline for the current product direction

Prepared for the Kashi project • Tailored to the 2026-04-21 progress deck and concept note • 21 April 2026

**Working thesis.** Bottom line: the strongest economic case for Kashi is not "one leave case costs X."

It is that persistent manager-linked communication asymmetry creates a long period of hidden operating loss before anything becomes formal, and Kashi shortens that period by making repeated structural patterns visible earlier.

That claim fits Kashi's architecture. More inflated claims — legal harassment detection, exact productivity loss per manager, or direct prevention of leave — do not.

Purpose of this brief: convert scattered research into something decision-ready for the project — usable in the deck, concept note, buyer narrative, pilot plan, and internal claim discipline.

## Executive conclusion

This brief summarizes the research most relevant to Kashi's current product and go-to-market shape. It is not a generic wellness memo. It is a project-usable synthesis of what the external evidence actually supports, what it does not support, and how that maps onto Kashi's live architecture and pitch.

The central conclusion is simple. Kashi should be sold as a people-risk visibility layer that reduces the duration, severity, and escalation probability of hidden operating loss. The strongest buyer logic is not "AI detects harassment" and not "one stress-leave case costs ¥7.9M." Those lines are too late-stage, too legally loaded, and too easy for a skeptical CFO, GC, or procurement lead to attack. The better thesis is: repeated harmful interaction patterns are economically costly long before they become formal, and Kashi makes those patterns reviewable sooner using structural signals only.

For Kashi specifically, the research supports five high-confidence points:

1. The largest cost bucket is usually hidden productivity loss, especially presenteeism, not visible leave administration.
2. Regrettable attrition and late formal escalation are expensive enough that even partial risk reduction can justify spend.
3. Manager-linked team drag is economically real, but should be modeled conservatively and probabilistically rather than as precise causal savings.
4. Kashi's no-affect, no-content-classification, longitudinal structural design is well aligned with both the evidence and emerging regulation.
5. Pilot success should be proven first through leading indicators and persistence reduction, not by immediately promising reduced harassment claims or reduced psychiatric leave.

## 1. What the external research actually says

The evidence base points to a cost chain, not a single crisis event. Poor workplace conditions and unresolved conflict do not become expensive only when an employee takes leave, resigns, or files a complaint. They are expensive earlier, during the long period in which performance, participation, confidence, coordination, and managerial attention degrade while the issue is still deniable or diffuse. This matters for Kashi because the product is structurally strongest in exactly that earlier period: repeated interaction asymmetry across meetings, not legal adjudication after the fact.

### 1.1 Hidden productivity loss is the biggest economic bucket

WHO reports that depression and anxiety alone cost the global economy an estimated 12 billion working days and US$1 trillion per year, predominantly through lost productivity rather than only through visible absence.[3] Deloitte's 2024 UK employer research reaches the same shape of conclusion at employer level: poor mental health costs UK employers about £51 billion annually, with presenteeism the largest contributor at around £24 billion. Deloitte also reports an average return of about £4.70 for every £1 spent on workforce mental-health and wellbeing support.[4]

The Japan-side picture is directionally similar and more relevant for Kashi's domestic pitch. Yokohama City University reported in 2025 that the annual productivity loss associated with mental-health-related presenteeism and absenteeism in Japan was about ¥7.6 trillion, of which roughly ¥7.3 trillion came from

presenteeism and only about ¥0.3 trillion from absenteeism; the total equaled roughly 1.1% of GDP.[5] That finding matters because it supports the exact buyer logic Kashi needs: the direct visible absence event is not the main bill. The main bill is the hidden operating leakage while the employee is still nominally present and the organization still thinks nothing has formally happened.

Implication for Kashi: the product should be framed as reducing the half-life of invisible people-risk leakage, not as a leave-case reimbursement story. Kashi's strongest economic promise is earlier visibility into persistent interaction risk before productivity loss compounds. This directly matches the product's live emphasis on structural meeting telemetry, longitudinal aggregation, and review-worthy event construction rather than legal or emotional classification.[1][2]

### 1.2 Leave and long absence matter, but mostly as tail risk and visible receipt

Long absence is still economically important, but it should not be the centerpiece of the value case. The reason is simple: salary itself is often already committed cost. A CFO will immediately ask whether the company is truly incurring incremental expense or simply failing to receive expected output from a budgeted employee. The more defensible leave-cost model is therefore not "salary lost," but: lost productive capacity, cover or overtime cost, delivery slippage, manager and HR handling time, ramp-back inefficiency after return, and elevated risk of later resignation or formal procedure.

The Health and Safety Executive in Great Britain reports that work-related stress, depression, or anxiety resulted in an average of 22.9 days lost per case in 2024/25.[12] That provides a useful illustration of the duration that a severe case can consume, but even this should be treated as a tail indicator, not the whole story. The RIETI paper on Japanese firm data makes this explicit: sick leave and turnover due to bad mental health are only the "tip of the iceberg," and firms with higher mental-health-related sick-leave or turnover rates tend to show lower annual profit rates, especially when fixed employment costs are high.[6]

Implication for Kashi: use leave as a visible late-stage cost anchor, not as the main logic. The real message should be that Kashi helps institutions see the pattern before it matures into long absence, not that Kashi directly predicts or prevents psychiatric leave.

### 1.3 Regrettable attrition is one of the cleanest financial justifications

Regrettable attrition is a stronger buyer bucket than vague "culture improvement" language because replacement cost is legible to finance. Gallup estimates that replacing leaders and managers costs around 200% of their salary, replacing technical professionals costs around 80%, and replacing frontline workers costs around 40%, excluding additional losses in morale and institutional knowledge.[10] The exact ratio will vary by role and market, but the order of magnitude is large enough that one avoidable exit can justify meaningful spend.

For Kashi, this matters because the product is aimed at repeated managerial or team-level interaction asymmetry. Those patterns are plausible upstream drivers of regrettable exit, especially where one contributor is repeatedly interrupted, dismissed, chilled, or deprived of normal contribution uptake across many meetings. Kashi does not need to prove that a specific exit was legally caused by one meeting pattern. It only needs to support a more modest and defensible claim: unresolved repeated interaction asymmetry increases the probability of regrettable attrition, and earlier visibility makes earlier correction more likely.

Critical caution: Kashi should never talk about “reducing attrition” in the abstract. Some turnover is neutral or even healthy. The disciplined phrase is “reducing regrettable attrition associated with unresolved manager- or team-risk patterns.” That is financially meaningful without pretending that every departure is a failure or that Kashi can prove direct causality in every case.

### 1.4 Formal escalation is late, slow, and materially more expensive than earlier correction

Acas estimated the cost of workplace conflict to UK employers at about £28.5 billion per year and found that cases which moved into formal procedures cost more than three times those resolved informally.[8] That is exactly the escalation logic Kashi should lean into. The economic problem is not merely the grievance file itself. It is the extra manager time, HR time, investigation time, witness time, workflow disruption, relationship damage, and avoidable delay that accumulate once a pattern is allowed to harden into a formal matter.

This is one of the cleanest bridges between the external literature and Kashi’s product structure. Kashi does not need to promise that it will resolve every case informally or that formal procedures are bad. It only needs to show that persistent review-worthy patterns can be surfaced earlier, which should lower the share of cases that remain unmanaged until only formal handling is left. Earlier structured visibility is economically valuable even when later formal process remains necessary.

Implication for Kashi: do not pitch the product as an “investigation tool” first. Pitch it as an earlier warning and triage tool that reduces preventable formalization. Investigation support is secondary. Prevention of avoidable escalation is the stronger buyer story.

### 1.5 Manager-linked team drag is real, but attribution must stay conservative

Gallup’s engagement research is frequently cited because it shows that manager quality has system-level effects: managers account for at least 70% of the variance in engagement across teams, and top-quartile engaged business units outperform bottom-quartile units on profitability, productivity, absenteeism, turnover, and safety.[9] Specifically, Gallup’s meta-analysis reports 23% higher profitability, 14% higher productivity (production records and evaluations), 78% lower absenteeism, and lower turnover in highly engaged units.[9]

For Kashi, this evidence is directionally helpful but must be handled with discipline. Engagement is not the same as harassment, and no one should claim that Kashi will create a 14% productivity increase. That would be unserious. The valid use of this research is narrower: manager behavior scales economically beyond the individual target; repeated harmful meeting dynamics can damage the functioning of a whole team; therefore even a small recoverable fraction of team drag can be financially meaningful.

The correct modeling stance is scenario-based. Kashi should not say “this manager costs the team 14% productivity.” It should say that persistent concentration of interruption, chilling, ignored turns, or agreement asymmetry can be treated as a governance-relevant risk signal for team drag, and that a conservative recoverable share — for example 1% to 3% of team output over the unmanaged period — is already enough to matter. That is much harder to dismiss.

## 2. What the Japan-specific evidence says

The Japan evidence is not identical to the UK or global sources, but it points in the same strategic direction and is particularly useful for Kashi's domestic enterprise narrative.

First, the prevalence problem is real. The Ministry of Health, Labour and Welfare's 2024 summary of the fiscal-2023 workplace harassment survey reported that 19.3% of workers had experienced power harassment in the previous three years.[7] That is not a fringe event rate. Second, organizational response is often weak even when the issue is recognized. In the same materials, when employers had acknowledged power harassment, "did nothing" was still the most common response at 28.7%.[7] Third, employees' assessment of whether their employer actively takes harassment prevention seriously is correlated with markedly different experience rates: among respondents who rated their employer as "actively" addressing prevention and resolution, the power-harassment experience rate was much lower than among those who rated the employer as "not very active"; the public summary highlights 15.2% versus 35.1%.[7]

These points support a practical Kashi thesis. The problem is not only that harmful behavior exists. It is that organizations often react weakly, inconsistently, or too late, and that workers are left bearing the burden of recognition and proof. Kashi's value proposition — structural pattern visibility over time, private awareness, manager mirror, role-based escalation support, and no casual raw-content browsing — is tightly aligned with that institutional failure mode.[1][2]

A second Japan-specific point comes from the RIETI firm-level evidence and Yokohama City University's 2025 productivity-loss estimate. Together they strengthen a buyer narrative that is more CFO-usable than generic culture language: visible leave and resignation are the measurable tail of a much larger hidden productivity problem, and firms with worse mental-health-linked leave or turnover indicators perform worse financially.[5][6] That is exactly why Kashi should sell "earlier visibility into hidden loss" rather than "wellness" or "empathy."

## 3. Why this research is relevant to Kashi's actual architecture

Kashi's architecture matters because the product should be sold in ways that are congruent with what it really does. The project documents are already strong on this point. Kashi is positioned as governance infrastructure rather than a harassment classifier or employee-monitoring tool; it uses structural signals from meetings, longitudinal windows, review-worthy events rather than moral labels, per-speaker baseline calibration, and strong role-based visibility controls.[1][2]

This matters commercially and legally. The external research supports Kashi most strongly where Kashi is already strongest technically:

- repeated pattern recognition matters more than one-off moments;
- hidden cost accumulates before formal crisis;
- manager-linked behavior can scale to team-level outcomes;
- earlier visibility and earlier correction are economically meaningful;
- causal certainty is not required for governance value.

The same research is much weaker where Kashi is currently weakest or deliberately refuses to go:
• semantic harassment classification;
• emotion or affect inference;
• exact causal attribution from one meeting pattern to one legal outcome;
• automated HR decisions;
• precise monetary attribution at individual manager level.

That means Kashi's existing refusal posture is not just ethics. It is evidentiary discipline. The official European Commission AI Act page now explicitly lists emotion recognition in workplaces as a prohibited practice, effective from February 2025, and classifies AI tools for employment and worker management among the high-risk areas with phased rules coming into force from August 2026 and August 2027.[11] Kashi's refusal to infer affect and its insistence on structural metadata therefore strengthen, rather than weaken, its defensibility.

One internal coherence point still deserves tightening. Kashi's public materials sometimes say transcripts are the source data, while elsewhere stressing that the system never transcribes for analysis and operates only on structural metadata.[1][2] That can confuse legal and procurement readers. The cleaner formulation is: Kashi may ingest transcript and diarization artefacts made available by meeting platforms as source input, but employer-side scoring is performed on structural interaction metadata only; Kashi does not perform semantic harassment classification, emotion inference, or affect analysis. That wording is closer to the actual thesis and easier to defend.

## 4. Claim discipline: what Kashi can say, cannot say, and should say instead

A lot of enterprise products fail here. They find a real problem, then oversell the last 20% and poison trust. Kashi should be stricter than that.

| Claim | Status | Why | Recommended use |
|---|---|---|---|
| "Kashi reduces hidden labor-cost leakage by surfacing persistent interaction risk earlier." | Strong | Matches external economics on presenteeism, escalation cost, and the importance of earlier visibility; matches Kashi's longitudinal structural design. | Use this as the main buyer thesis. |
| "Kashi helps reduce the probability, duration, and severity of regrettable attrition and formal escalation." | Moderately strong | Reasonable expected-loss framing; avoids pretending to prove causality for every case. | Good in deck and concept note. |
| "Kashi identifies repeated manager-linked patterns that may be suppressing | Strong | Direct fit with Kashi's detectors: interruption concentration, chilling delta, ignored turns, agreement | Use in product explanation and manager mirror framing. |

| Claim | Status | Why | Recommended use |
|---|---|---|---|
| contribution and team challenge quality.” | | asymmetry, floor-time inequality. | |
| “Kashi detects harassment / power harassment.” | Weak / avoid | External research and regulation do not support automatic legal or moral classification from these signals alone. | Replace with “surfaces review-worthy interaction patterns.” |
| “Kashi reduces leave cases by X%.” | Weak / avoid | Too lagging, too causal, and too easy for finance/legal to reject. | Talk about earlier visibility before leave or burnout becomes formal. |
| “This manager costs the team 14% productivity.” | Avoid | Gallup’s engagement figures are not a direct per-manager loss meter and not a Kashi-specific result. | Model conservative recoverable drag scenarios instead. |
| “Kashi improves culture.” | Too vague | Not wrong, but too soft for budget approval and too fuzzy for measurement. | Replace with “reduces persistence of unresolved communication-risk patterns.” |

## 5. Economic model Kashi should actually use

The buyer model should be portfolio-based, not single-incident theatre. A CFO is more likely to trust an expected-loss framework than a hero-number about one catastrophic case.

Recommended model:

Expected annual preventable loss
= team-drag loss
+ regrettable-attrition loss
+ formal-escalation premium
+ long-leave / absence tail risk

Where:
• team-drag loss = team payroll or value base × conservative recoverable drag % × unmanaged duration
• regrettable-attrition loss = number of regrettable exits linked to unresolved people-risk patterns × replacement cost
• formal-escalation premium = formal-route handling cost − earlier-correction handling cost
• long-leave / absence tail risk = number of severe cases × incremental lost-output and coverage cost

This structure is superior to a single leave-case story for four reasons. First, it recognizes that most of the cost is hidden before the crisis. Second, it uses a portfolio view that finance understands. Third, it allows conservative scenario modeling instead of fake precision. Fourth, it fits Kashi's longitudinal, pattern-based architecture much better than a case-management or benefits-administration story.

Recommended verbal formulation:

Kashi does not primarily save money by "catching harassment." It reduces the time that repeated manager- or team-linked interaction asymmetries remain invisible and unmanaged. That lowers hidden productivity leakage, reduces the likelihood of regrettable exits and avoidable formal escalation, and helps the organization act before weak signals mature into expensive cases.

Illustrative modeling discipline:
• Use low / base / high scenarios, not single-point ROI.
• Treat leave as visible tail risk, not the central bucket.
• Treat team drag as small recoverable percentages, not dramatic engagement deltas.
• Distinguish economic sponsor from operational buyer: the pain sits with CEO/CFO/business head, but HR/Legal/Compliance usually govern deployment.
• Do not present external averages as direct realized savings. Present them as category anchors that justify internal baseline modeling and pilot measurement.

## 6. Pilot implications: what Kashi should prove first

A short pilot is unlikely to prove reductions in leave or litigation. Kashi should therefore stage proof in three layers.

Layer A — leading indicators (prove these first)
• time from repeated pattern emergence to human review
• persistence of directional asymmetry across 30/90/180-day windows
• concentration of interruptions, ignored turns, unanswered questions, or agreement shifts
• manager acknowledgement and follow-through rate after mirror feedback
• false-positive discipline and cold-start suppression performance

Layer B — intermediate operating outcomes
• fewer repeated patterns reaching manual complaint stage
• lower manager/HR hours spent on late-stage triage per flagged unit
• fewer repeat concerns in previously flagged teams
• better contribution distribution after intervention windows

Layer C — long-lag business outcomes
• lower regrettable attrition in flagged teams versus matched baselines
• lower stress-related absence in flagged teams over longer windows
• reduced share of cases requiring full formal procedure
• higher participation stability and lower baseline drift among previously suppressed contributors

This sequence matters. If Kashi tries to sell lagging hard-outcome ROI too early, it will sound like every

other HR-tech overclaim. If it proves that it can surface patterns earlier, with good governance, and that those patterns become less persistent after human action, the larger business case becomes much more credible later.

| Proof layer | What to measure | Why it matters |
|---|---|---|
| A | Leading indicators | Prove that Kashi surfaces repeated patterns earlier, with acceptable false-positive discipline and useful persistence logic. |
| B | Intermediate operating outcomes | Show that human action after visibility changes team pattern persistence and reduces repeat concern load. |
| C | Long-lag business outcomes | Reserve for later windows; use matched baselines and do not overpromise short-pilot proof of leave or litigation reduction. |

## 7. Concrete recommendations for the project

1. Re-center the deck around hidden operating loss, not around one leave case. Keep dramatic case numbers as optional examples, not as the main logic.

2. Strengthen the sponsor/gatekeeper framing. Say that the economic pain sits with CEO/CFO/business leadership while operational deployment usually routes through HR, Legal, Compliance, and labor relations. Do not overstate “not HR.”

3. Use “review-worthy interaction pattern” consistently instead of legal or moral labels. That is already the strongest internal construct.

4. Tighten the transcript/metadata wording so procurement does not read the deck as self-contradictory.

5. Add one slide or one document section called “What Kashi can prove vs what Kashi does not claim.” This increases trust.

6. For Japan-facing sales, explicitly connect Kashi to two institutional failures shown in public evidence: non-trivial prevalence and weak employer response even after recognition.[7]

7. For CFO-facing materials, add the portfolio formula and conservative scenario logic. Remove any phrasing that implies exact per-manager productivity accounting.

8. Keep the no-affect, no-content-classification posture. It is not a feature sacrifice; it is a defensibility advantage under both evidence constraints and the EU AI Act posture.[11]

9. Use Kashi’s scenario 2 and scenario 3 outputs as a rhetorical teaching device: structural signals are

powerful for persistent asymmetry, but not omniscient for all hostile content. That strengthens trust because it admits boundary conditions.[1]

10. Make pilot success criteria explicit before sales conversations. Otherwise every buyer will silently substitute their own and later declare the pilot inconclusive.

## 8. Deck-ready language you can reuse

**Main buyer line:**

Kashi is not a harassment classifier and not a wellness dashboard. It is a governance layer that shortens the time that persistent people-risk patterns remain invisible and unmanaged.

**Finance line:**

The largest cost is usually not the visible leave case. It is the long period before that, when contribution is suppressed, delivery slows, managers spend time firefighting, attrition risk rises, and avoidable formal escalation becomes more likely.

**Defensibility line:**

Kashi does not infer emotion, intent, or illegality. It surfaces repeated structural interaction asymmetries for human review, using longitudinal patterns rather than one-off moments.

**Pilot line:**

A pilot should not be judged first by rare late-stage outcomes. It should be judged by whether Kashi surfaces repeated patterns earlier, whether those patterns become less persistent after human action, and whether fewer issues drift into late formal handling.

**Japan-facing line:**

In the Japanese workplace context, the core problem is often not the absence of rules but the lateness and weakness of institutional visibility. Kashi addresses that visibility gap.

## 9. Final verdict

The research does not justify saying that Kashi detects harassment, proves legal wrongdoing, quantifies exact productivity loss per manager, or directly prevents leave. It does justify something narrower and, in practice, more valuable: persistent interaction asymmetry is economically relevant long before formal crisis; hidden productivity leakage is larger than visible absence; regrettable attrition and formal

escalation are expensive enough to matter; and a system that surfaces repeated structural patterns earlier can create real governance and financial value.

That narrower claim is not a downgrade. It is the serious version. It matches the evidence, matches Kashi's architecture, and gives the project a clearer buyer story than generic culture-tech language ever could.

## 10. Source notes

The sources below are the basis for the claims in this brief. External studies are used as category anchors and evidence boundaries, not as direct realized ROI for any specific customer. Internal Kashi documents are used to test architectural fit and wording discipline.

| Source | Use in this brief |
|---|---|
| [1] Internal: Kashi — Progress & Project Overview (2026-04-21). | Internal project PDF. Used for current product architecture, detector list, governance stance, and existing buyer language. |
| [2] Internal: Transparency That Drives Institutional Accountability concept note. | Internal project DOCX. Used for current positioning, governance model, permissions, and product thesis. |
| [3] World Health Organization, "Mental health at work" (2024). | WHO fact sheet stating depression and anxiety cost 12 billion working days and US$1 trillion in lost productivity globally. https://www.who.int/news-room/fact-sheets/detail/mental-health-at-work |
| [4] Deloitte UK, "Poor mental health costs UK employers £51 billion a year for employees" (2024). | States £51bn annual cost, presenteeism around £24bn, and average £4.70 return per £1 spent. https://www.deloitte.com/uk/en/about/press-room/poor-mental-health-costs-uk-employers-51-billion-a-year-for-employees.html |
| [5] Yokohama City University, "メンタル不調の影響、年間7.6 兆円の生産性損失に" (2025). | States approx. ¥7.3tn presenteeism loss, ¥0.3tn absenteeism loss, ¥7.6tn total, around 1.1% of GDP. https://www.yokohama-cu.ac.jp/news/2025/20250611hara.html |
| [6] RIETI Discussion Paper 16-E-016, "Does Mental Health Matter for Firm Performance? Evidence from longitudinal Japanese firm data" (2016). | Finds mental-health-related sick leave and turnover are the tip of the iceberg and are associated with lower firm profit rates. https://www.rieti.go.jp/jp/publications/dp/16e016.pdf |
| [7] MHLW, "令和 5 年度 職場のハラスメントに関する実態調査 結果概要" (2024). | Used for prevalence of power harassment, weaker response patterns, and lower experience rates where anti-harassment efforts are rated active. https://www.mhlw.go.jp/content/11909000/001259093.pdf |
| [8] Acas, "Workplace conflict: estimating the cost to employers" (2021). | States formal procedures cost more than three times informal resolution. https://www.acas.org.uk/workplace-conflict-estimating-the-cost-to-employers |
| [9] Gallup, "Q12 Meta-Analysis \| 11th Edition" and related | Used for 23% profitability, 14% productivity, 78% lower absenteeism, and lower turnover in highly engaged units. |

| Source | Use in this brief |
| --- | --- |
| engagement research. | https://www.gallup.com/workplace/321725/gallup-q12-meta-analysis-report.aspx |
| [10] Gallup, “Employee Retention Depends on Getting Recognition Right” (2024). | Used for replacement-cost estimates: ~200% managers, ~80% technical roles, ~40% frontline. https://www.gallup.com/workplace/650174/employee-retention-depends-getting-recognition-right.aspx |
| [11] European Commission, AI Act page (accessed April 2026). | Used for prohibition of emotion recognition in workplaces and timeline of prohibited/high-risk rules. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai |
| [12] UK Health and Safety Executive, “Working days lost in Great Britain” (2025). | Used for average 22.9 days lost per stress, depression, or anxiety case in 2024/25. https://www.hse.gov.uk/statistics/dayslost.htm |