

Why “no signal” must never mean “no problem”

Research synthesis tailored to Kashi’s current design: meeting-only, structural-signal detection, no content classification, no affect inference, and no HR decisions from the tool.

Prepared for: product / governance / pitch refinement

Date: 2026-04-21

Bottom line

False negatives are not peripheral noise for Kashi. They are a central governance risk.

Kashi can reduce one important class of institutional blindness — repeated structural asymmetry in recorded meetings — without proving that the absence of signal means the absence of harm. If the product fails to state that clearly, its interface can become a new denial mechanism.

Executive summary

- The strongest external evidence concerns underreporting, retaliation risk, organizational tolerance, and the broad scope of workplace harm beyond formal meetings.
- The literature supports several Kashi-relevant blind spots: harm outside recorded meetings, omission-based exclusion, coalition or mobbing behavior, and retaliation after a person uses a reporting or support channel.
- The evidence does not support any inference that “no meeting-pattern signal” equals “no problem.” In fact, the reporting literature suggests the opposite: institutions often remain blind precisely because many targets do not report and many harms do not present as a clean event.
- Kashi’s current architecture is defensible, but its bounded scope must be made explicit in product copy, dashboard states, escalation flows, and governance language.
- The safest posture is: Kashi is a visibility layer, not an exoneration engine.

1. Why this matters specifically for Kashi

Kashi’s current posture is intentionally narrow. It focuses on recorded meetings, structural interaction metadata, deterministic detectors, and role-limited presentation. That narrowness is a strength: it avoids content surveillance, affect inference, and overclaiming. But it also creates structured blind spots.

Those blind spots matter because Kashi is not a private diary or a general workplace-intelligence system. It observes only a bounded slice of work. If a manager, investigator, or executive sees “no qualifying signal detected” and silently upgrades that to “the person is probably fine,” the product has failed its own governance purpose.

The core risk is not merely technical under-detection. It is institutional misuse of under-detection. A false negative can become more dangerous when it is wrapped in a tidy dashboard state, because the absence of alarm may be read as an institutional permission to discount the lived experience.

2. False-negative taxonomy for Kashi

The table below separates high-confidence literature-backed blind spots from one important inferential risk. This is the usable project view.

Risk bucket	What Kashi may miss	Evidence level	Why the miss occurs	Main implication
Scope false negatives	Exclusion, harassment, or retaliation outside recorded meetings	High	Convention 190 / ILO treat the world of work as broader than the formal meeting room	No-meeting signal cannot rule out harm
Omission false negatives	“Allowed to speak” but ignored, out of the loop, not meaningfully heard, not staffed, not acted on	High–Medium	Exclusion often operates through omission rather than overt interruption	Kashi must not reduce harm to audible confrontation
Coalition / mobbing false negatives	Distributed or quiet group behavior rather than one obvious aggressor	Medium	Coordinated silence, alignment, or social freezing may leave weak dyadic traces	Pattern absence in one dyad may conceal group-based targeting
Reporting / retaliation false negatives	People who do not use the tool, under-report, or are harmed after using it	High	Fear of retaliation and disbelief suppresses reporting and distorts observable behavior	No report + no signal is not a clean state
Metric-adaptation false negatives	Behavior becomes more polished after people learn what the system notices	Inference	People may shift from overt interruption to agenda control, access denial, or civil non-response	Kashi should describe this as a plausible limitation, not a proven finding

3. Research-backed failure mode 1 — harm outside recorded meetings

The International Labour Organization’s framing is directly relevant here. Convention No. 190 treats violence and harassment as occurring not only in the workplace in the narrow physical sense, but also during work-related trips, travel, training, events or social activities, through work-related communications including ICT-enabled communications, in employer-provided accommodation, and when commuting to and from work.[1]

That matters because Kashi’s current design is intentionally centered on recorded meetings. This means the system can surface meeting-based power asymmetries without claiming visibility into the full ‘world of work.’ A person can be safe inside the recorded meeting and unsafe outside it. A team can perform politely on-platform and exclude someone in side channels, staffing conversations, or social coordination that never enters Kashi’s observation window.

Therefore, Kashi should never permit a move from “no structural meeting signal detected” to “no workplace harm likely occurred.” That is not a conservative conclusion; it is an unsupported one.

4. Research-backed failure mode 2 — omission-based exclusion and “never meaningfully heard”

A second blind spot is omission. Some harmful patterns do not look like sharp interruption or overt verbal aggression. They look like being left out of the loop, not being given consequential work, not being answered, not being credited, or being allowed to speak without being allowed to matter.

The ostracism literature supports this. Research on workplace ostracism in academia describes information exclusion — being ‘left out of the loop’ — as an important form of exclusion with negative effects on trust and liking, even when the target is not visibly expelled from the group.[2] Related work distinguishes classic ostracism from being out of the loop and suggests these are distinct exclusion experiences rather than trivial variants.[3]

This matters for Kashi because a user may technically have floor time and still be organizationally neutralized. The system’s current detectors already move in this direction through unanswered-question rate, topic-credit ignored turns, and chilling delta. That is good. But the product should still state openly that many omission-based harms remain only partially observable from meeting structure alone, especially when exclusion occurs through staffing, information routing, or follow-up inaction rather than in-meeting suppression.

5. Research-backed failure mode 3 — quiet coalition or mobbing behavior

The literature on bullying and mobbing is useful here because it expands the picture beyond one aggressor versus one target. ILO materials distinguish mobbing as a collective or group phenomenon in which several people ‘gang up’ on a target worker.[4] Other literature similarly describes mobbing as a sophisticated group wear-down process rather than a single dramatic act.[5]

For Kashi, the relevance is straightforward: a quiet coalition may not produce a clean dominant-interruptor signature. Harm can be distributed across several people through non-response, side alignment, repeated social non-support, selective endorsement of one another, or coordinated devaluation of one person’s contributions. In such cases, each individual actor may look only mildly problematic while the combined pattern remains severe.

So the UI and governance language should not encourage users to think only in dyads. A weak signal for one manager-target pair does not falsify the possibility of broader group-based targeting.

6. Research-backed failure mode 4 — underreporting and retaliation

This is the strongest part of the evidence base. The EEOC Select Task Force report states that roughly three out of four individuals who experienced harassment never even talked to a supervisor, manager, or union representative about it.[6] The same EEOC materials say the fears that suppress reporting are ‘well-founded’: one cited study found that 75 percent of employees who spoke out against workplace mistreatment faced some form of retaliation, and reporting is often followed by indifference, trivialization, hostility, or reprisals.[7]

This has two direct implications for Kashi. First, non-use of the tool, non-escalation, or silence after a harmful pattern is not a clean negative. It may be a fear-shaped negative. Second, if the product provides a private explainer, evidence vault, or report-generation path, it must assume that the period after use may itself become a retaliation-risk window.

The product should therefore treat retaliation as a second-order blind spot rather than as a solved downstream process. Kashi may help a person make sense of a pattern, but that does not mean the institution becomes safe immediately after the pattern is surfaced.

7. Research-backed failure mode 5 — organizational tolerance makes absence of signal dangerous

The National Academies’ synthesis on harassment is useful here not because Kashi is a sexual-harassment classifier, but because it identifies a deep organizational pattern: organizational climate is, by far, the greatest predictor of harassment occurrence, and perceived tolerance matters through at least three channels — the risk to people who complain, lack of sanctions against offenders, and the perception that complaints will not be taken seriously.[8]

This matters because if Kashi’s outputs are misread as proof of absence, the product can accidentally reinforce organizational tolerance. A manager or HR function may say, implicitly or explicitly: ‘the system saw nothing.’ In a climate already tilted toward disbelief or procedural defensiveness, a polished ‘no signal’ state can harden denial rather than weaken it.

So the governance issue is not just detector performance. It is interpretive control. Kashi must govern what institutional actors are allowed to infer from missing or weak signals.

8. Important inferential risk — behavior may adapt around the metric

This point is logically strong but empirically weaker than the categories above. The reviewed sources do not provide Kashi-specific evidence that meeting-governance metrics produce adaptation effects in exactly this form. So this should be presented as a reasoned limitation, not as a proven empirical claim.

Still, the concern is credible. Once an institution understands what the system tends to capture, overt interruption can decline while exclusion shifts into agenda control, side-channel alignment, selective access denial, or a more civil style of systematic non-response. In other words, the same power asymmetry may persist in a less countable form.

The correct Kashi posture is therefore: this is a plausible blind spot requiring monitoring and humility. It should not be sold as established fact, but it should be acknowledged in product and governance documentation.

Core product rule

No detected signal is not evidence that no harm is occurring.

At most, it means that no qualifying signal was detected within this system’s observable scope and threshold logic.

9. What this means for Kashi product, governance, and UI

- **Do not let “no signal” equal “healthy.”** Those are different states. A safer state model is: signal detected; no qualifying signal detected in observed data; insufficient observation window; or concern may be out of scope for this layer.
- **Never use missing signal as a closure condition.** A complaint, concern, or lived experience should not be closed merely because Kashi did not surface a qualifying meeting-pattern event.

- **Put scope warnings where users infer meaning.** The warning belongs on dashboards, case summaries, and report-generation screens, not only in a governance page nobody reads.
- **Name out-of-scope harms explicitly.** Examples: staffing exclusion, access denial, private-channel retaliation, social freezing, off-platform harassment, commuting and event-related harm.
- **Treat post-use periods as special risk windows.** If a person opens the explainer, generates a report, or escalates, the organization should assume retaliation risk may increase rather than disappear.
- **Keep Kashi framed as an assistive visibility layer.** The product should help people structure and surface patterns without pretending to certify either wrongdoing or safety.

10. Drop-in language Kashi can reuse

Governance-page sentence

Kashi observes only a bounded subset of workplace interaction — primarily repeated structural patterns in recorded meetings. No detected signal should be interpreted as proof that no harm occurred, because exclusion, retaliation, coalition behavior, access denial, and other harms may occur outside the system’s observable scope or below its detection threshold.

Dashboard disclaimer

No qualifying signal detected in observed meeting data. This is not evidence of safety, and it does not rule out harms occurring outside recorded meetings or through omission-based exclusion.

Case-review rule

Absence of Kashi signal may inform, but must not determine, the handling of a concern. Missing signal is not exculpatory evidence.

Short internal slogan

Kashi is a visibility layer, not an exoneration engine.

11. Claims Kashi should avoid

- “No alert means no problem.”
- “Kashi can tell whether a person’s concern is real.”
- “Missing structural evidence weakens the person’s account.”
- “Meeting safety implies workplace safety.”
- “The product sees retaliation once reporting happens.”

Safer alternative: Kashi can reduce one important kind of blindness, but it cannot close the world of work into a single clean signal state.

12. Sources used

- [1] International Labour Organization. Convention No. 190 and related explanatory materials on the scope of the ‘world of work’. Key point used here: violence and harassment may occur during work-related trips, training, events or social activities, through work-related communications, in employer-provided accommodation, and while commuting.
- [2] Zimmerman, C.A. et al. (2016). Examining Workplace Ostracism Experiences in Academia. Key point used here: information exclusion / being ‘left out of the loop’ has negative effects and should not be treated as trivial.
- [3] Jones, E.E. et al. (2020). Being Ostracized Versus Out of the Loop: Redundant or Unique Forms of Exclusion? Key point used here: ostracism and being out of the loop are distinct exclusion experiences.
- [4] International Labour Organization materials on violence, bullying, and mobbing. Key point used here: mobbing may involve a group of workers ganging up on a target worker.
- [5] Khoo, S.B. (2010). Academic Mobbing: Hidden Health Hazard at Workplace. Key point used here: mobbing can take the form of sophisticated group wear-down behavior.
- [6] U.S. Equal Employment Opportunity Commission (EEOC), Select Task Force on the Study of Harassment in the Workplace (2016). Key point used here: roughly three out of four individuals who experienced harassment never even talked to a supervisor, manager, or union representative about it.
- [7] EEOC Select Task Force materials and associated prevention deck. Key point used here: fears of retaliation are well-founded; one cited study found that 75% of employees who spoke out against workplace mistreatment faced some form of retaliation.
- [8] National Academies of Sciences, Engineering, and Medicine (2018), Sexual Harassment of Women: Climate, Culture, and Consequences. Key point used here: organizational climate / perceived tolerance is the strongest predictor of harassment occurrence; risk to complainants, lack of sanctions, and not being taken seriously are central dimensions.
- [9] Kashi — Progress & Project Overview (2026-04-21). Internal project document used to tailor this memo to Kashi’s current scope, detector posture, and governance stance.

Usable conclusion for the project:

Kashi should explicitly refuse the inference that clean UI equals clean reality. The product is strongest when it says: we surface one important class of hidden pattern early; we do not certify that everything else is fine.

13. Reference URLs

- **ILO / Convention No. 190 scope of the world of work:** <https://www.ilo.org/media/7476/download>
- **EEOC Select Task Force report (PDF):**
https://www.eeoc.gov/sites/default/files/migrated_files/eeoc/task_force/harassment/report.pdf
- **EEOC task force webpage:** <https://www.eeoc.gov/select-task-force-study-harassment-workplace>
- **National Academies — Sexual Harassment of Women (chapter pages):**
<https://www.nationalacademies.org/read/24994/chapter/4>
- **National Academies summary finding on organizational climate:**
<https://www.nationalacademies.org/read/24994/chapter/9>
- **Workplace ostracism in academia (PMC):** <https://pmc.ncbi.nlm.nih.gov/articles/PMC4885318/>
- **Being Ostracized Versus Out of the Loop (PMC):** <https://pmc.ncbi.nlm.nih.gov/articles/PMC8297684/>
- **Academic Mobbing: Hidden Health Hazard at Workplace (PMC):** <https://pmc.ncbi.nlm.nih.gov/articles/PMC4170397/>

- **ILO work-related violence / mobbing overview:** <https://www.ilo.org/media/167791/download>