INSIGHTS · THE INSIGHT FUNCTION

Insight Management Platforms and Knowledge Repositories

How organisations store, find and reuse what they have learned about customers, what research on organisational memory and knowledge management says about why repositories succeed or fail, and how to design one that people actually use.

ModQu research article · 2,531 words · 23 references (APA 7th edition)

Most organisations that have commissioned research for some years hold a large body of knowledge about their customers. Much of it is effectively lost. Reports sit in shared drives, email attachments and agency portals; the people who commissioned them move on; and new projects begin by asking questions that earlier studies have already answered. The cost is not only the money spent on repeated research. It is also the slower and poorer decisions that result when existing knowledge is not used.

Insight management platforms and knowledge repositories are intended to solve this problem. They store research findings and related material in one place, make them searchable, and help people find and reuse what the organisation already knows. This article explains the problem these tools address, draws on research on organisational memory and knowledge management to explain why repositories often fail, and describes the practices that make them work.

The problem: organisational forgetting

Organisations learn, but they also forget. Levitt and March (1988) described organisational learning as the encoding of inferences from history into routines that guide behaviour, and noted that what is learned must be stored and retrieved if it is to be used. Walsh and Ungson (1991) defined organisational memory as stored information from an organisation's history that can be brought to bear on present decisions. They identified several places where it is held: in individuals, in the culture, in transformations such as procedures, in structures and roles, in the physical workplace, and in external archives. Much of what an organisation knows about its customers is held in the first of these, the memories of individual employees, and is lost when those employees leave.

Huber (1991) distinguished four processes in organisational learning: knowledge acquisition, information distribution, information interpretation and organisational memory. Market research departments have traditionally been strong at the first. The other three are where knowledge is most often lost. Sinkula (1994) applied this view to market information and argued that organisational memory is a central part of how firms process what they learn about markets.

The consequences are measurable. Moorman and Miner (1997) studied new product development and found that organisational memory affected both the performance and the creativity of new products, and that its effects depended on how dispersed the memory was across the organisation. Memory can also be a constraint: highly consistent memory can lead teams to repeat what has worked before rather than consider new possibilities. The task is therefore not simply to store more, but to store knowledge in a form that informs new decisions without closing off new thinking.

Why knowledge does not flow by itself

It might seem that knowledge, once recorded, would spread naturally to those who need it. Research suggests otherwise. Szulanski (1996) studied the transfer of best practices within firms and found that the main barriers were not lack of motivation but knowledge-related factors: the recipient's lack of capacity to absorb new knowledge, ambiguity about why a practice worked, and difficult relationships between source and recipient. Argote and Ingram (2000) argued that the ability to transfer knowledge within the firm is a source of competitive advantage precisely because it is difficult, and that knowledge embedded in people, tools and routines transfers differently from knowledge written down.

Hansen (1999) showed that how hard knowledge is to transfer depends partly on its nature. Codified knowledge, which can be written down and understood without much context, can be shared through weak ties between units. Complex, tacit knowledge requires strong relationships for effective transfer. Customer insight contains both kinds. A survey result can be codified easily; an understanding of why a customer segment behaves as it does may depend on context that is hard to capture in a report.

Nonaka (1994) described organisational knowledge creation as a continuing conversion between tacit knowledge, which is personal and hard to express, and explicit knowledge, which can be recorded and shared. Nonaka and Takeuchi (1995) argued that new knowledge is created as organisations move repeatedly between these forms: sharing experience, articulating it, combining it with other explicit knowledge, and internalising it through use. A repository mainly supports one part of this cycle, the combination of explicit knowledge. It does little for the others unless it is connected to people and conversations.

Two strategies for managing knowledge

Hansen et al. (1999) identified two broad strategies for managing knowledge in professional organisations. In a **codification** strategy, knowledge is captured in documents and databases where it can be found and reused by anyone. In a **personalisation** strategy, knowledge stays mainly with the people who developed it, and the organisation's role is to help people find and talk to each other. They argued that successful firms emphasise one strategy, supported by the other, and that the choice should follow the nature of the work. Firms that solve similar problems repeatedly benefit from codification; firms whose work is highly customised benefit more from personalisation.

Insight functions need both. Many customer questions recur in slightly different forms, such as how satisfied customers are, how well a brand is known, or why customers leave, and codified findings serve them well. Other questions are new, and the most valuable resource is a colleague who understands the relevant customers deeply. An effective insight platform therefore stores findings and also shows who produced them, so that a user can follow up with a person when a document is not enough. Cross and Sproull (2004) found that the information relationships people find most useful for action provide not only answers but also help in framing problems, validating plans and gaining legitimacy, which documents cannot easily supply.

What an insight repository holds

A typical insight repository contains several kinds of material. The most basic is the research report itself, with its data, methods and findings. More useful repositories also hold short summaries of key findings, written so that someone who did not commission the research can understand them; tagged insights, in which individual findings are extracted and labelled by topic, customer segment, product and date; raw or processed data where permitted; and materials such as interview recordings, verbatim comments and images. Some also hold a record of the decisions the research informed.

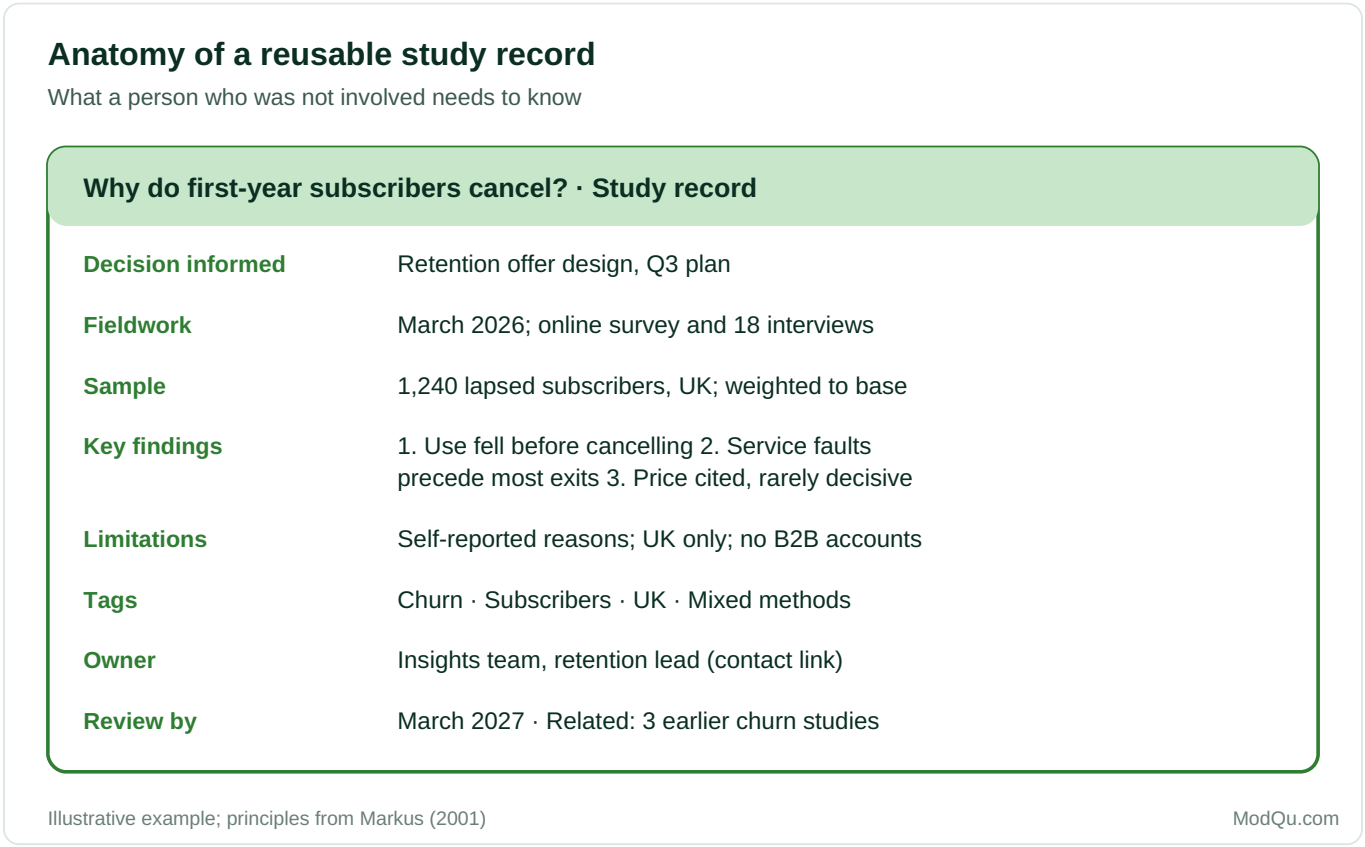


Figure 1. The fields a well-designed study record contains. Short summaries, clear limitations and an owner who can be contacted make findings reusable. Illustrative example, diagram by ModQu.

The structure of each record matters more than the technology. Markus (2001) studied knowledge reuse and found that what makes reuse succeed depends on who the reuser is. When the person reusing knowledge is the person who produced it, or a close colleague, little documentation is needed. When the reuser is a distant expert or a novice, knowledge must be documented with much more context: why the work was done, what it assumed, and where its results apply. Most insight repositories serve the second group: product managers, marketers and executives who were not involved in the research. Records must therefore explain context and limitations, not only results.

Why repositories fail

Many organisations have built repositories that few people use. Research on knowledge management systems helps explain why.

People do not contribute. Adding material to a repository takes time and offers little immediate reward. Kankanhalli et al. (2005) studied contributions to electronic knowledge repositories and found that the costs of codification, such as time and effort, discouraged contribution, while enjoyment in helping others and perceptions that one's knowledge was useful encouraged it. Organisational rewards had a positive effect only under some conditions. Wasko and Faraj (2005) found that people contributed knowledge in electronic networks when doing so enhanced their professional reputation and when they were structurally embedded in the network. For insight teams, the implication is that contributing should be built into the normal process of completing a project rather than treated as an additional task.

People cannot find what they need. Search is often poor, tagging inconsistent, and titles uninformative. Pirolli and Card (1999) described information seeking as a form of foraging, in which people follow cues that signal where useful information is likely to be found and abandon sources that seem unproductive. If the first few searches in a repository return irrelevant or outdated results, users stop searching and commission new research instead.

What they find is not trusted or not usable. A user who finds a five-year-old report with no summary and no indication of its sample cannot tell whether its findings still hold. Findings without context are hard to apply to a new decision.

The repository is disconnected from work. Alavi and Leidner (2001) reviewed knowledge management systems research and argued that such systems are effective only when they support the knowledge processes of the organisation, creation, storage and retrieval, transfer and application, rather than storage alone. A repository that exists separately from the processes by which research is commissioned and decisions are made will be bypassed.

Haas and Hansen (2005) added a further caution. In a study of sales teams in a management consulting firm, they found that using more electronic documents did not always improve performance; for experienced teams, reliance on documents could even hurt by consuming time without adding value. Knowledge repositories help most when the knowledge they hold is relevant and the user lacks the experience to proceed without it.

Designing a repository that works

Research and practice suggest several design principles.

Build capture into the research process

The best time to create a record is when a project ends, while the people involved still remember its context. Make the completion of a standard record, with summary, key findings, limitations and tags, a required step before a project is closed. Agencies can be asked to supply summaries in the required format as part of their deliverables.

Write for the distant reuser

Records should be understandable by someone who was not involved. This means a plain statement of the question and the decision it was meant to inform, a description of the sample and method, three to five key findings in plain language, the main limitations, and the date of fieldwork. Markus (2001) showed that this context is what distant reusers most need.

Use a consistent taxonomy

Tag each record with a controlled set of terms covering topics, customer segments, products, markets and methods. A consistent taxonomy makes search and browsing more reliable and allows the insight team to see where knowledge is thick and where it is thin.

Connect findings to people

Each record should name an owner who can be contacted. This combines codification with personalisation, following Hansen et al. (1999), and allows users to obtain the tacit context that a record cannot hold.

Curate and retire

Customer knowledge ages. A finding about media habits from several years ago may be misleading today. Records should carry a review or expiry date, and the insight team should periodically review older material, marking findings that remain valid, those that have been superseded, and those that should no longer be relied on.

Link insights across studies

The most valuable function of a repository is often to show what many studies together say about a topic. Synthesis pages that summarise the organisation's knowledge on key themes, with links to the underlying studies, help users see the big picture and identify gaps. Grant (1996) argued that a firm's main role is integrating specialised knowledge held by individuals into its products and decisions; synthesis is one way an insight function performs this integration.

Measure use

Track how often records are searched, viewed and cited in decisions, and how often new research requests are resolved by existing knowledge. These measures show whether the repository is working and where it needs improvement.

A simple test of a repository's value is the number of new research requests that can be answered, wholly or in part, from what the organisation already knows. If the insight team routinely checks the repository before commissioning new work, and records when existing studies were sufficient, it builds both a measure of savings and a habit of reuse. Over time, the record of which studies are reused most also shows which kinds of research have lasting value and which answer only the question of the moment.

The insight lifecycle

Taken together, these principles describe a lifecycle through which each piece of insight passes.

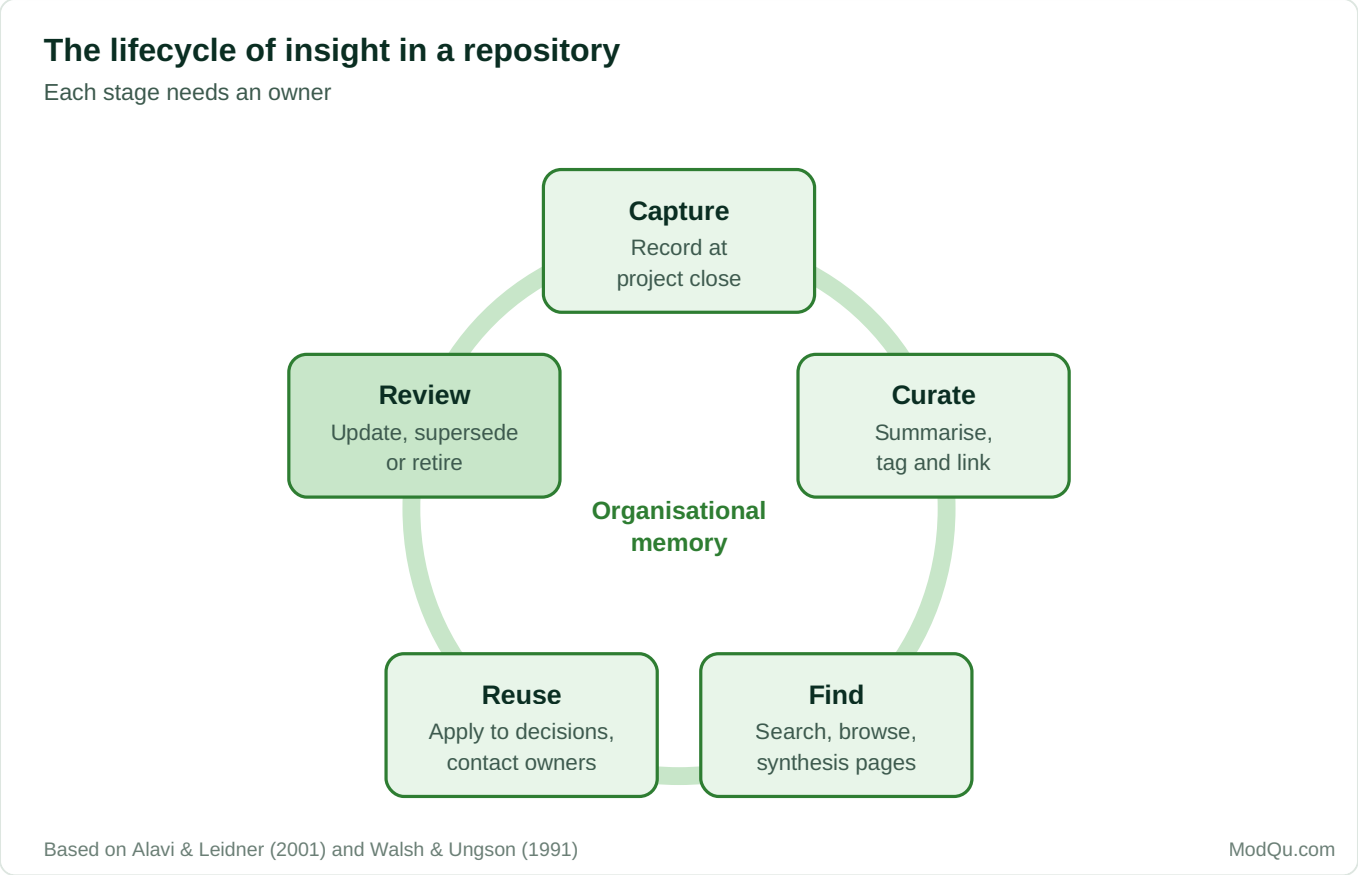


Figure 2. The lifecycle of insight in a repository. Each stage needs an owner; the review stage is the one most often neglected. Diagram by ModQu.

Insight platforms and artificial intelligence

Commercial insight management platforms have added features that address some of the problems described above, including automatic transcription and tagging of qualitative material, natural-language search, and automatic summaries. Recent systems use large language models combined with retrieval from a document store, an approach known as retrieval-augmented generation, to answer questions in plain language by drawing on the documents in a repository (Lewis et al., 2020). Such tools can make existing knowledge easier to find and use, particularly for people who would not otherwise search a repository.

They also raise new questions. A generated answer may combine findings from studies with different samples, dates and definitions without making the differences clear. It may present a finding with more confidence than the original study supports, or omit the limitations a careful reader would note. The principles that make a repository trustworthy, clear

records, consistent definitions, visible dates and limitations, and named owners, become more important rather than less when an automated system is summarising the content. Users should be able to see which studies an answer draws on and check them.

Repositories and the use of research

A repository's value lies in whether its contents are used in decisions. Research on the use of market research shows that use depends on more than access. Deshpandé and Zaltman (1982) found that managers were more likely to use research when they trusted the researchers and when the findings were presented in a way that fitted their needs. Menon and Varadarajan (1992) developed a model of how marketing knowledge is used within firms and distinguished between direct, instrumental use and more conceptual use, in which research shapes how managers think about an issue. A good repository supports both: it lets a manager find a specific number for a specific decision, and it lets anyone build a broader understanding of customers by reading across studies. Said et al. (2015) found that managers valued insight because it was new, shared and acted upon. A repository addresses the second of these directly and makes the third more likely.

Key takeaways

Organisations forget much of what they learn about customers, because knowledge is held by individuals, scattered across systems, and recorded without the context needed to reuse it. Research on organisational memory and knowledge transfer shows that knowledge does not flow by itself and that repositories succeed only when they are connected to work, written for people who were not involved, and linked to the people who produced the knowledge. Effective repositories build capture into the research process, use consistent records and tags, name owners, retire outdated findings, synthesise across studies, and measure use. New tools based on artificial intelligence can make repositories easier to search, but they increase rather than reduce the need for well-structured, dated and qualified records.

References

Alavi, M., & Leidner, D. E. (2001). Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS Quarterly, 25*(1), 107–136. https://doi.org/10.2307/3250961

Argote, L., & Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes, 82*(1), 150–169. https://doi.org/10.1006/obhd.2000.2893

Cross, R., & Sproull, L. (2004). More than an answer: Information relationships for actionable knowledge. *Organization Science, 15*(4), 446–462. https://doi.org/10.1287/orsc.1040.0075

Deshpandé, R., & Zaltman, G. (1982). Factors affecting the use of market research information: A path analysis. *Journal of Marketing Research, 19*(1), 14–31. https://doi.org/10.1177/002224378201900102

Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal, 17*(S2), 109–122. https://doi.org/10.1002/smj.4250171110

Haas, M. R., & Hansen, M. T. (2005). When using knowledge can hurt performance: The value of organizational capabilities in a management consulting company. *Strategic Management Journal, 26*(1), 1–24. https://doi.org/10.1002/smj.429

Hansen, M. T. (1999). The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly, 44*(1), 82–111. https://doi.org/10.2307/2667032

Hansen, M. T., Nohria, N., & Tierney, T. (1999). What's your strategy for managing knowledge? *Harvard Business Review, 77*(2), 106–116.

Huber, G. P. (1991). Organizational learning: The contributing processes and the literatures. *Organization Science, 2*(1), 88–115. https://doi.org/10.1287/orsc.2.1.88

Kankanhalli, A., Tan, B. C. Y., & Wei, K.-K. (2005). Contributing knowledge to electronic knowledge repositories: An empirical investigation. *MIS Quarterly, 29*(1), 113–143. https://doi.org/10.2307/25148670

Levitt, B., & March, J. G. (1988). Organizational learning. *Annual Review of Sociology, 14*, 319–340. https://doi.org/10.1146/annurev.so.14.080188.001535

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems 33* (pp. 9459–9474). Curran Associates.

Markus, M. L. (2001). Toward a theory of knowledge reuse: Types of knowledge reuse situations and factors in reuse success. *Journal of Management Information Systems, 18*(1), 57–93. https://doi.org/10.1080/07421222.2001.11045671

Menon, A., & Varadarajan, P. R. (1992). A model of marketing knowledge use within firms. *Journal of Marketing, 56*(4), 53–71. https://doi.org/10.1177/002224299205600405

Moorman, C., & Miner, A. S. (1997). The impact of organizational memory on new product performance and creativity. *Journal of Marketing Research, 34*(1), 91–106. https://doi.org/10.1177/002224379703400108

Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science, 5*(1), 14–37. https://doi.org/10.1287/orsc.5.1.14

Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. Oxford University Press.

Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review, 106*(4), 643–675. https://doi.org/10.1037/0033-295X.106.4.643

Said, E., Macdonald, E. K., Wilson, H. N., & Marcos, J. (2015). How organisations generate and use customer insight. *Journal of Marketing Management, 31*(9–10), 1158–1179. https://doi.org/10.1080/0267257X.2015.1037785

Sinkula, J. M. (1994). Market information processing and organizational learning. *Journal of Marketing, 58*(1), 35–45. https://doi.org/10.1177/002224299405800103

Szulanski, G. (1996). Exploring internal stickiness: Impediments to the transfer of best practice within the firm. *Strategic Management Journal, 17*(S2), 27–43. https://doi.org/10.1002/smj.4250171105

Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review, 16*(1), 57–91. https://doi.org/10.5465/amr.1991.4278992

Wasko, M. M., & Faraj, S. (2005). Why should I share? Examining social capital and knowledge contribution in electronic networks of practice. *MIS Quarterly, 29*(1), 35–57. https://doi.org/10.2307/25148667