Principles of Knowledge Auditing

Foundations for Knowledge Management Implementation

By: Patrick Lambe

Citation:

Principles of Knowledge Auditing: Foundations for Knowledge Management Implementation
By: Patrick Lambe
DOI: 10.7551/mitpress/14908.001.0001
ISBN (electronic): 9780262373166
Publisher: The MIT Press
Published: 2023

OA Funding Provided By:

OA Funding from MIT Press Direct to Open

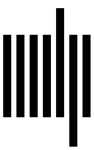

The MIT Press

15 The Inventory Audit: Auditing Knowledge Stocks

> An unbridled lucidity can destroy our understanding of complex matters.
> —Polanyi (1966, p. 18)

Having sought to disambiguate some of the more confusing metaphors we use to describe knowledge, my task in the remainder of this section of the book is to stir up the waters again by looking at how we identify knowledge in organizations.

I believe we work to overly simplistic models of how knowledge is deployed and used in organizations, and this compromises the effectiveness of how we can audit and inventory knowledge. So in the spirit of the quote from Polanyi above, I want to break apart these simplistic views, and in doing so explore a range of competing ways in which we might understand and describe knowledge. By the end of the following chapters, when the sediment has settled, we should have a clearer but less simplistic way of seeing, a way that is more suited to our purpose in a knowledge audit.

The inventory audit of knowledge stocks was the most frequently recognized type of knowledge audit in our 2017 survey covering 150 knowledge management (KM) professionals. It was recognized by 80 percent of respondents. It is almost always combined with another form of audit, the most frequent being a participative goal-setting audit, a discovery review audit, or a value audit. So the inventory audit functions as the preparatory, evidence-gathering stage for these other audit types (Lambe, 2017).

It follows that if we are going to compile useful inventories of knowledge stocks, we need a reliable and robust way of describing a set of knowledge types that we can easily differentiate and identify. This set of knowledge types is what we call a knowledge typology.

What We Need from a Typology of Knowledge

There are two unhelpful dualisms in how knowledge stocks are categorized in the general KM literature and in common practice. The first is the *individual knowledge*

versus collective knowledge distinction, and the second is the *tacit knowledge versus explicit knowledge* distinction (Hislop, 2013, pp. 21–23). They are unhelpful because they are too crude to adequately describe the variety that exists in the organizational knowledge landscape. Chapters 16 and 17 will examine each of these dualisms in turn, but first I want to clarify what a good typology of knowledge should provide.

Any good typology divides the landscape into a sufficiently diverse set of types to support sensemaking and action in relation to a given purpose. Binary typologies imply a very simple landscape, and this is manifestly not true in relation to knowledge management (Nissen & Jennex, 2005).

Moreover, every typology should have a clear purpose, and we measure its adequacy against its ability to inform the desired outcomes of that purpose. A typology is a representation of a landscape to serve an interpretive or sensemaking purpose. Without a purpose, a typology is just an intellectual fiction, bearing little utility in practice. Let us look at an example of an unhelpful binary typology.

Case Study: When Typologies Fail

Endometrial cancer is a cancer of the lining of the uterus. It is among the top four cancers in women in the US and in 2013 was estimated to account for over eight thousand deaths. Until recently, the cancer has been placed into one of two categories by specialists on the basis of a physical examination of thin slices of the tumors under a microscope. Type I tumors are considered likely to have a favorable outcome after surgery and radiation, while Type II tumors are more aggressive, have poor outcomes, and require chemotherapy.

However, the two types are difficult to distinguish, and there is disagreement among pathologists about how tumors should be classified. Disagreement about classification means inconsistency of classification and disagreement about the course of treatment. In short, the typology is not very helpful. It does not appear to match the complexity of the phenomena it is attempting to describe, and it does not provide robust recommendations for useful interventions.

In 2013 the Cancer Genome Atlas Research Network published a study of 373 endometrial tumors based on genome analysis, and they distinguished four distinct types of endometrial cancer, including variants that had structural and mechanical similarities with some types of colorectal and breast cancer, for which good therapeutic options had already been developed.

The findings, based on a new way of analyzing the tumors at the molecular level, provided much more fine-grained, reliable, and consistent methods for distinguishing endometrial cancer types, and this had immediate therapeutic implications, including the possibility of adapting prior therapies for the related cancer types.

This is what we should expect from an effective typology: greater differentiation of types based on real-world observable features, enabling practical, useful actions (Cancer Genome Atlas Research Network, 2013; Kolata, 2013).

So what is our purpose here, the purpose against which we need to evaluate the quality of our typologies of knowledge? We want to know what knowledge the organization has and what knowledge it needs to do its work. We want to be able to conduct an inventory audit of knowledge stocks in order to inform a discovery review, participative goal-setting, or assessment audit. Or perhaps even a value audit.

In simpler terms, we want to gain an understanding of how the organization depends upon (and produces) knowledge in relation to its core activities, so that we can design improvements to the way it uses, exploits, and produces knowledge. Extending the medical metaphor, we want to be able to design "therapeutic" interventions.

Case Study: Auditing Knowledge with Binary Typologies

Let us look at a case study of a knowledge audit based on a binary typology (tacit and explicit knowledge types or individual and corporate knowledge types) to see what happens when the typology is too simple. In 2001 a group of researchers at the Robert Gordon University in Aberdeen conducted a knowledge audit for the tax department of a large oil and gas company comprising twenty employees. They used an assessment survey followed up by interviews on knowledge processes, and they worked with the employees to develop knowledge maps of the tacit and explicit knowledge needed and used in their work (Burnett et al., 2004).

These knowledge maps had two problems. The first was that there was no common standard for reporting collective dependencies on knowledge, and in consequence, the individual knowledge maps were so disparate the audit team was not able to compile a common department-level knowledge map (Burnett et al., 2004, p. 33).

The second problem was in the maps themselves. In the examples reported by the team, we can discern fifteen distinct explicit knowledge types, while tacit knowledge (or "people") was a single category without any further differentiation (Burnett et al., 2004, pp. 32–33). Because explicit knowledge is more observable than tacit knowledge, when we rely on unguided self-reporting, it is easier to assign detailed labels to subtypes of explicit knowledge. For tacit knowledge, without an agreed-upon way of differentiating subtypes, a single large bucket has to do the job.

When the work has a heavy dependency on tacit knowledge, the consequences for intervention planning based on an audit's findings are profound. There is a strong bias toward overarticulation of explicit knowledge and low definition of tacit knowledge, which may not be consistent with the nature of the work. High levels of detail for explicit knowledge sources and vague references to tacit knowledge mean that recommendations will focus on explicit knowledge and have higher specificity for those resources, while interventions focused on tacit knowledge will be frustratingly vague. And indeed, the recommendations from this audit follow that pattern. While tacit knowledge was covered in the recommendations, the details were vague, and the level of detail in any follow-up would necessarily be biased toward the detail that was available to describe the explicit knowledge resources.

As a case in point, one of the recommendations in that study was for a taxonomy to improve access to, and the availability of, current knowledge. Where tacit knowledge is an important

resource, then any taxonomy must describe areas of tacit knowledge. When knowledge maps privilege explicit knowledge, so will the taxonomy, leaving great uncertainty about the capacity of the taxonomy to support access to tacit knowledge resources.

What does this case study tell us? First, that gathering data at the wrong level of granularity can pose problems. As Tom Stewart (2001) pointed out two decades ago, the first task of the knowledge manager is to determine the right unit of analysis, and in organizations that means being able to answer the question "what does the group need to know?" (p. 119; cf. Kogut & Zander, 1997, p. 312).

A typology that examines either the individual or the organization as a whole impedes the ability to gather data at the right level of detail—the group that does the work. The knowledge of the individual ("knowledge in people") is a distraction from what the group does and so is the pressure to describe knowledge at a higher organizational level. If that seems a controversial assertion, we will defend it later.

Second, the case tells us that the crude differentiation between tacit and explicit knowledge privileges the more observable portion of the duality (in this case explicit knowledge), resulting in an imbalance in how the follow-up attention and action is directed. The dualism biases knowledge management toward the explicit.

How then should we gather data about the knowledge in use within organizations? As with cancers, physical observation of the full spectrum of knowledge use is neither easy nor necessarily reliable. A great deal of knowledge use in organizations happens within people's heads and in transient interactions between people. There is a lot of noise in all this activity. The salient and most important knowledge use is often not the most easily or directly observable. We can only observe the full spread of knowledge use by looking at proxy signals of knowledge use. It follows that a good typology should help us find and describe good proxies for knowledge use.

By far the most authoritative witnesses to knowledge uses and needs are the people who use knowledge in their daily activities. While external consultants are frequently used to conduct knowledge audits, when it comes to inventorying knowledge resources, external facilitators are merely that: facilitators. They do not carry intimate knowledge of how the organization does its work. The staff who do the work do carry that knowledge, so we need to find a way of enabling self-reporting of knowledge use. This is also more sustainable in the long run for the ongoing maintenance of these knowledge inventories.

However, individual reports may be prone to bias, or to poor self-insight, and they may only represent partial views of the work. We need to find a way of using collective reports, oriented around the knowledge that work groups need and use. Our typology

should therefore be accessible to the people who work with knowledge and who understand how it feeds their performance. It should not be too abstract or artificial. If the typology is not accessible to them, if it does not enable a naturalistic description of the way they work, and if it is not capable of being discussed and agreed upon in a group format, then our respondents will not be able to report their knowledge uses and dependencies reliably, consistently, and sustainably. This implies an approach to inventory audits based on collective knowledge-mapping exercises and not individual interviews where responses need to be somehow integrated by analysts with secondhand knowledge.

In summary, to be useful in an inventory audit a typology of knowledge must demonstrate the following five characteristics:

- *Observable* The knowledge types need to be capable of being described and documented in consistent ways by different respondents. This condition goes to our desire for a broadly reliable and reproducible way of inventorying knowledge.
- *Naturalistic* The knowledge types need to represent distinctions that make functional sense to people in the enterprise. Respondents can readily identify knowledge resources as contributing factors in work and describe them in ways that will be consistently understood by their peers. This also goes to reliability and reproducibility.
- *Actionable* Following from our belief that a knowledge audit presupposes a theory of change, the knowledge types need to be relatively easily associated with actions to manage them, so that identification of knowledge resources can lead to decisions about how to conserve, grow, and manage them.
- *Comprehensive* The knowledge types need to cover the full range of knowledge resource types in common use within the enterprise.
- *Granular* A knowledge typology also needs to support inventorying knowledge resources at the right level of granularity—that is, knowledge as it is used in the context of work. Many of the typologies of knowledge we will discuss describe knowledge as it is used by individuals abstracted from specific tasks within their work group—that is, *personal knowledge*. Some of them describe higher-level typologies dealing with the way that organizations work with knowledge, often at a strategic level—that is, *organizational* or *strategic knowledge*. For an organizational knowledge inventory, we need reports that are contextualized to work.

As Joseph Horvath (2000) puts it, "Philosophers may define knowledge in *structural* terms (i.e., in terms of its relation to other concepts) but, in business settings, it makes more sense to define knowledge in *functional* terms (i.e., in terms of its use)" (p. 35).

Business functions are organized in the service of organizational objectives and are natural aggregators for the information and knowledge resources required to meet those

objectives (Orna, 2004, p. 71; Henczel, 2001, p. 64). Business-function-oriented audits capture knowledge in use at probably the most stable level of detail for naturalistic and reproducible discrimination of knowledge types, and they support action planning.

By business functions I do not mean organizational structure; I mean the more stable, underlying business functions that organization structures attempt to organize and connect. Organization structures may change relatively often, but business functions remain fairly stable over time. And business-function-oriented audits have the most authoritative witnesses we can hope to find as to knowledge use: the people who perform those business functions on a daily basis.

As we will see later in this section of the book, organizational or strategic knowledge is too broad based for our purposes in an inventory audit, and knowledge associated with individuals, while extremely fine-grained, can be difficult to associate with follow-up actions. Business functions persist as individuals come and go, and they reflect the detailed components of activity that depend upon, and produce, consistent supplies of knowledge resources. Function-oriented audits also have the advantage of being more easily connected to business performance measures and the bottom line when it comes to intervention planning (Hasanali et al., 2003, p. 15).

Many typologies of knowledge exist in the KM literature, all developed for different purposes, but few meet our five conditions for auditability, and few can be used to identify and describe knowledge that is used at the functional level. That is not to deny their usefulness for other purposes.

In the remaining chapters, we will look in detail at the way typologies of knowledge have been constructed in the past, and why many of them fail our criteria for auditability. We will propose a typology that does meet these criteria, and we will close with a case study using this typology, that shows how a more differentiated typology, collected at the work-group level, redresses the bias toward explicit knowledge that we saw in the oil company example above.

* * *

Summary

In this chapter we returned to the importance of the inventory audit as a foundational audit activity for other forms of knowledge audit. I laid out the importance of having a clear typology of knowledge types and the risks of using an inappropriate typology. Here is a summary of the main points:

1. A useful typology needs to describe observable forms of knowledge or observable proxies for knowledge use.
2. In order to get reliable witness reports, the typology needs to describe knowledge in naturalistic ways.
3. The typology of knowledge types should be actionable—that is, lead to inferences about how the knowledge types can be managed.
4. The typology should be comprehensive and cover the full range of knowledge use in organizations.
5. The typology should be sufficiently granular to reflect the types of knowledge work at a functional or operational level in the organization.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Lambe, Patrick, 1960– author.
Title: Principles of knowledge auditing : foundations for knowledge management implementation / Patrick Lambe.
Description: Cambridge, Massachusetts : The MIT Press, [2023] | Includes bibliographical references and index.
Identifiers: LCCN 2022015397 (print) | LCCN 2022015398 (ebook) | ISBN 9780262545037 (paperback) | ISBN 9780262373159 (epub) | ISBN 9780262373166 (pdf)
Subjects: LCSH: Knowledge management. | Organizational learning. | Organizational change.
Classification: LCC HD30.2 .L362 2022 (print) | LCC HD30.2 (ebook) | DDC 658.4/038—dc23/eng/20220718
LC record available at https://lccn.loc.gov/2022015397
LC ebook record available at https://lccn.loc.gov/2022015398