Knowledge Centers & Flexibility: How People Think About What They Know

Marc Demarest
marc@noumenal.com

July 1997

ABSTRACT

Commercial knowledge managers work routinely with subject matter experts (SMEs): people who are "in possession" of knowledge that the knowledge manager needs to embody and disseminate in the interest of producing more and better commercial performances. Each SME has, in addition to subject matter knowledge, a deeply-seated particular conceptual model of knowledge and how knowledge is developed and maintained (what this paper refers to as a knowledge center) and also more or less of an ability to understand other models of knowledge and how it is produced (what this paper refers to as a knowledge flexibility index, or KFI). Understanding SMEs' knowledge centers and knowledge flexibility indices is critical to successful knowledge engineering activities, since the tools, methods and approaches a knowledge manager uses must be tuned to the KC and KFI of the SMEs being interviewed.

This essay describes a typology for describing knowledge centers and knowledge flexibility, in the interest of allowing commercial knowledge managers to size and scope knowledge engineering activities in which they are engaged, and in the interest of associating certain kinds of knowledge engineering tools and techniques with particular kinds of knowledge engineering problems.

Some Preliminary Definitions

Commercial knowledge management is a **discipline**, practiced by **knowledge managers**, that takes as its primary objective **the (re)production of commercially-valuable performances** on the part of the firm's employees, suppliers, partners and customers, and therefore the performance of the firm itself in its chosen marketplaces.

A **discipline**, within this definition, is a series of techniques and tools for identifying or constructing, embodying, disseminating and facilitating the use of commercially-valuable knowledge. Within this four-phase process, each area – identification, embodiment, dissemination and facilitated use – has different techniques and tools, and problems, associated with it

Knowledge managers are charged with owning and executing this discipline. They may be formally charged by an organization – as is the case with knowledge officers, and, I would argue, with all managers – or they may be self-appointed, having recognized the key role knowledge pays in the execution of their responsibilities. The key identifier of a knowledge manager is therefore not that (a) she has an organizational charter to be a knowledge manager or (b) that she knows something (is a subject matter expert) but (c) that she executes a recognizable discipline that allows others to know and therefore to perform better.

Commercially-valuable performances are produced by people who are in possession, consciously or otherwise, of commercially-valuable knowledge. Commercially-valuable performances are in general those performances that tend toward the getting and keeping of more, better and more valuable customers for the firm. Specific measurements of commercial value are always ultimately financial measurements, but frequently have interim expressions having to do with quality, customer care, efficiency, effectiveness, or some other aspect of t he firm's performance. Knowledge managers are concerned with two broad classes of commercial performances: (a) innovation, which is the reproduction of commercially-valuable performances at larger and larger scale until innovative performances become the norm (standard operating procedure) and (b) exnovation, which is the systematic elimination of performances (or procedures or processes or behavior) which are no longer commercial valuable, or which are commercially debilitating, to the firm. Performances that are commercially valuable to a firm may not take place within the firm's boundaries. In fact, as disintermediation becomes the norm, rather than the exception, in commercial value systems within an increasingly globla economy, it is more likely that the performances a firm finds commercially-valuable will be executed within the boundaries of one of the firm's mates: its suppliers, partners, and customers. The knowledge manager's domain is therefore not circumscribed by the firm's boundaries (the value chain), but by the boundaries of the interfirm value system itself.

In such a definitional framework, defining what **knowledge** is in itself, or, worse, typing knowledge (know-how, know-what, know-who, etc.) as is currently the rage in KM circles, is (a) pointless (because these definitions become words about words, and nothing more), and (b) confusing, since it invites us to classify when we should be identifying, constructing, embodying, disseminating and facilitating use. As far as definitions of knowledge go, we should stop here: **we know commercially-valuable knowledge not by what it is, or what kind it is, but by what it does, namely produce commercially-valuable performances.**

Knowledge Engineering: Then, And Now

The term ***knowledge engineering*** used to refer to a set of techniques for "extracting" the knowledge "in the head" of a subject matter expert, and encoding that knowledge in a knowledge base within a software environment referred to as an expert system, which was, theoretically, capable of producing the behavior of the SME when given the SME's base of knowledge. Although this model – very popular in the early 1980s – has proven to be practically unworkable,

the term knowledge engineering still has great value for commercial knowledge managers in that it can stand as short-hand for a common scene in which knowledge managers have to practice their discipline.

That scene is, typically, as follows:

1. One, or a few, individuals routinely produce high-value performances of a particular kind in a particular area. These individuals are considered "top performers" and "subject matter experts." These performances may be of any sort at all: CSRs who consistently get rated well by customers, software designers who consistently produce programs with few problems per thousand lines of code (KLOCs), marketing managers whose outbound marketing programs consistently produce increases in awareness and market share, and so forth.

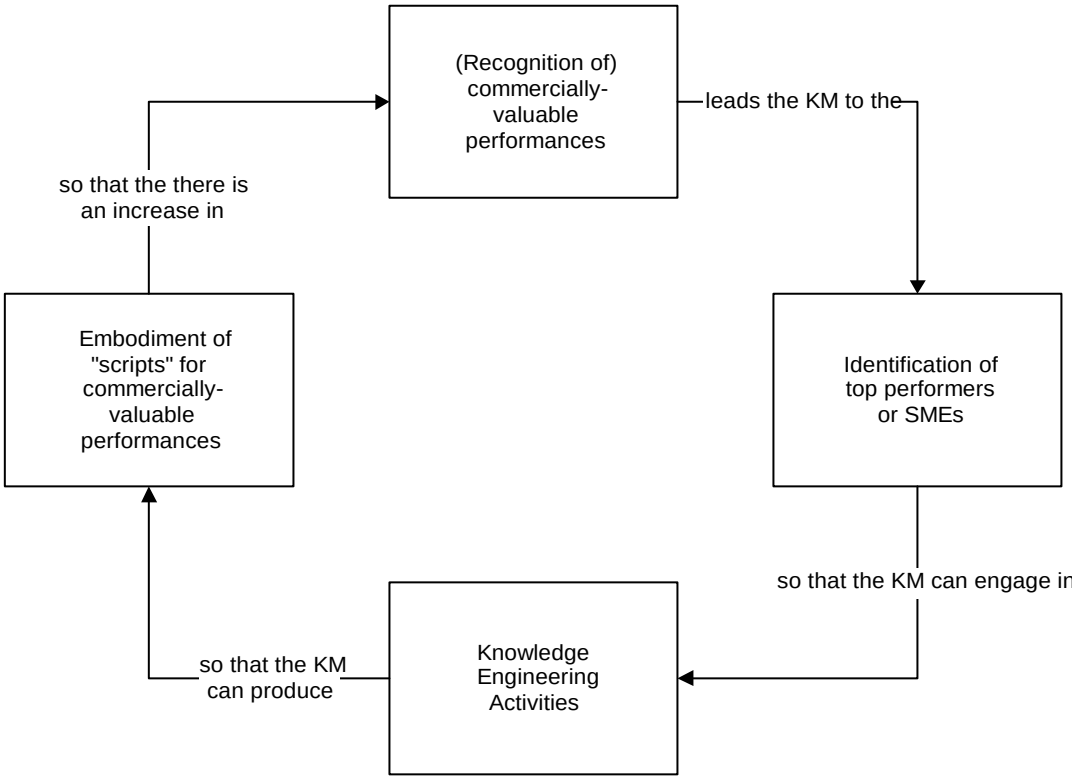

2. The knowledge manager identifies the performances of this individual or group of individuals as desirable for reproduction and scaling: she intends for all employees of that role type to behave as the top performers behave, and to produce approximations of the behavior of the SMEs.

3. Recognizing that she cannot reproduce the behavior of these top performers merely by observing the behavior[1], the knowledge manager meets with the top performers to engage in a **knowledge engineering** session: to "extract" the SMEs' knowledge in such as to embody

[1] The scientism of some knowledge managers, who believe that non-interactive observation of commercially-desirable performance on their part is sufficient for them to grasp the complexities of that behavior, never ceases to amaze me. Knowledge managers get their hands dirty because they must; there are no successful arm-chair knowledge managers.

and disseminate that knowledge in the interest of producing **an approximation[2] of the SMEs' performances** with other individuals.

4. The knowledge engineering sessions are inconclusive, in that the knowledge manager cannot extract any coherent version of the SMEs' practice from the discussions. Even a raw, unedited transcription of the knowledge engineering sessions does not provide the means to reproduce the SMEs behavior.

5. Subsequent attempts to embody the "knowledge" retrieved during the sessions, whether the embodiment is in software, in formal education, or in any other form, fail and do not produce even partial approximations of the desirable behavior of the SMEs.

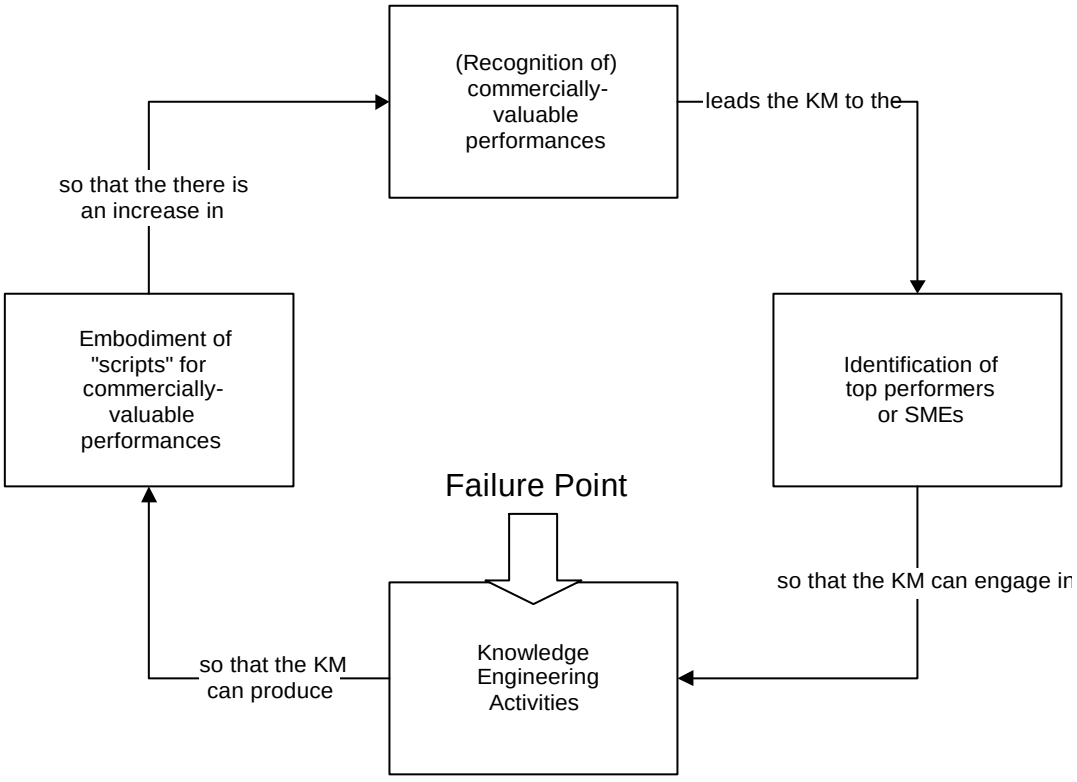

This situation is quite common among practicing knowledge managers, and, quite often in this situation, the knowledge manager concludes that the SMEs are engaged in wilful acts of obstruction: that they are trying not to share what they know. This is in large part the origin of the so-called "tacit knowledge problem" in knowledge management, also referred to as the "knowledge sharing problem" and characterized by the plaintive question, "How do we get people to share what the know?"

[2] The word approximation is critical in this context. Part of a human being's performance is inextricably linked to the capabilities, skills, aptitudes, flexibility, experience and psychology of that individual, and cannot be reproduced by any available technique. Anyone who has ever interviewed SMEs in a group is aware of this phenomenon, yet knowledge managers do not seem to acknowledge explicitly that they are in the business of reproducing approximations: that is, make-shifts.

It is probably not the case that acts of wilful obstruction are occurring. What is more likely is that the knowledge manager has not brought the SMEs into the knowledge construction process in a way that matches:

- how these SMEs believe knowledge is acquired, assimilated and used: **their knowledge center (KC)**
- how these SMEs view models of knowledge construction, assimilation and use that are different from their own:: their **knowledge flexibility (KF)**

Responsibility for the failure of the knowledge engineering activity rests, in these cases, not with wilfully-obstructionist SMEs but with insufficiently-skilled knowledge managers. Given that knowledge engineering exercises, most of them verbal and between KMs and SMEs, are the bread and butter of knowledge management, one of the key questions about these exercises is: why do they fail?

Knowledge Centers: How Do We Know What We Know?

Survival, at the small-e evolutionary level, is based on an organism's ability to adapt to its surroundings effectively. In the case of sentient self-reflecting organisms like human beings, the primary means by which the organism adapts successfully to its environment is through the acquisition of knowledge about successful and unsuccessful adaptation techniques. Thus, the premium on knowledge acquisition is deeply-wired into our evolutionary inheritance. Everyone is acquiring knowledge all the time.

However, individuals do not share a single model of how knowledge is (and how it ought to be) acquired. Most frequently, the variances in the models held by individuals are expressed in terms of false binary oppositions: book-learning versus common sense, theory versus practice, education versus experience. Like all binary oppositions, the two terms mark the poles of a spectrum[3], not an either/or choice, and reflect the fact that the real world is always already a mix of all the positions on the specturm. The place at which an individual locates herself on the spectrum between the two binary oppositions represents her knowledge center: what she believes to be (a) the foundation of her knowledge and (b) the best normative way to develop or acquire knowledge in general. In other words, every sentient creature has a theory-in-action for how we know what we know.

Aristotle recognized a three part spectrum of knowledge that serves well to stylize this problem of knowledge centers. The model is described in the diagram below.

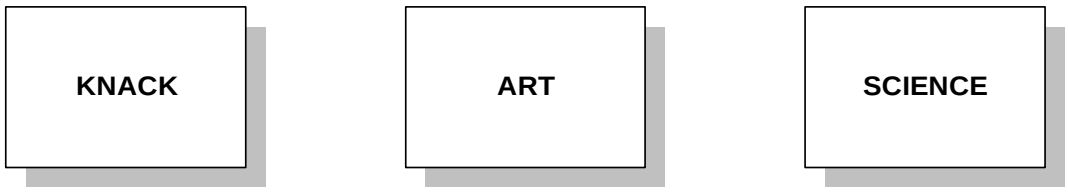

[3] Often, that spectrum is a Moebius strip, in that the two extremes are actually closely related to one another: two faces of the same coin. Fascism and Soviet Communism mark ends of a certain political spectrum, but share many common characteristics and are really more closely related to one another that either is related to, say, representative democracy, which sits in the middle of the spectrum the ends of which are marked by fascism and communism respectively.

A knack, is of course, the ability to do something well and repeatably: to perform consistently at a high level. As knowledge managers, this is where we always begin the process of managing knowledge: with the recognition that some one or some group has a knack that we need to reproduce and scale. Knacks, Aristotle would have argued, are non-reproducible: they are characteristics of individuals, and born of experience, insight, epiphanies, and instinct. They can be taught, if at all, only by long cycles of apprenticeship and corrective tutelage: by pair-bonding a student to a teacher who works and in so working provides a set of experiences whereby the student, over a long period of time, acquires the knack to whatever extent the student is able.

Through this process of pair-bonding and apprenticeship, some knacks become, by the actions of time, arts: heuristic, recipe-based knowledge that can be shared, but which is very imprecise, and retains a large measure of idiosyncracy. For Plato, rhetoric was clearly an art, as was dramaturgy and, to some extent, poesy.

Some arts become, again by the action of time and diffusion (and the intrusion of the philosopher or scientist), sciences: bodies of hard, largely invariant rules that if applied by any one with a modicum of training and experience, can reliable reproduce substantially the same results.

We can recognize, in this spectrum or typology, three fundamental knowledge centers, which I will term case, pattern and rule.

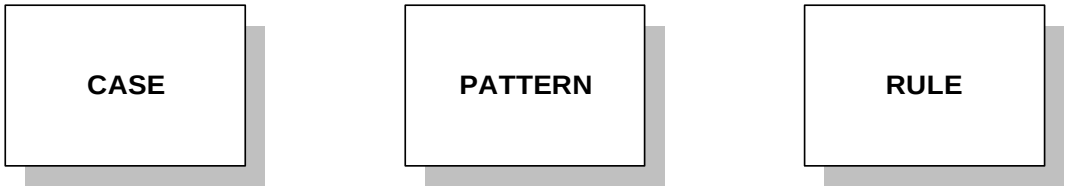

The case-based knowledge center is one characterized by the primacy of instinct, personal experience, anecdote, epiphany and insight in the acquisition of knowledge: I know what I know because (a) I have “a feel” or “a knack” for this kind of thing, (b) I have been doing this for a long period of time, (c) I heard that so-and-so did it this way and/or (d) as a result of either my knack or my experience, I have had several rare and important insights into the matter at hand that most people have not had. The case-based knowledge center places a premium on “practice” and completely devalues “theory” (the opposite end of the false binary opposition practice versus theory); similarly, case-based knowledge asserts that common sense is better than book-learning, and that no really useful knowledge can be acquired without the passing of significant amounts of time.

The pattern-based knowledge center is one characterized by a premium on the “recipes” – loose, fuzzy collections of probabilities that produce approximations of an ideal – that can be constructed from examining a set of cases. In other words, where the case-based knowledge center suggests, in the extreme, that one cannot know anything about a case before its appearance in the field of endeavor, the pattern-based knowledge center suggests that one can know a little or a lot about a case before its appearance in the field of endeavor if one knows about previous cases “like” the one in the offing: the key question is “how do we know that one case is like another? What are the key variables that indicate ‘likeness’?” As a trivial example, the purely case-based knowledge center would argue that one could not know how to drive Interstate 95 from New York to Florida successfully if one had not driven it already; the pattern-based knowledge center would argue that I-95 could be competently driven by anyone who had driven an US interstate, or perhaps any interstate or inter-country motorway anywhere in the world, but that one or more previous traversals of Interstate 95 would certainly increase the competency of a driver. Pattern-based knowledge centers recognize the value of experience, intuition, and epiphany, but also recognize that significant recipe-ish heuristic knowledge can be derived from little or no direct experience by means of inductive analytical techniques, that reflection on a large body of cases by a competent observer can produce heuristics that can be

applied to future cases in the interest of pre-judgement, or efficiencies, and that these heuristics can be learned and applied by someone with no practical experience whose subsequent behavior will approximate that of someone who "has the knack".

The rule-based knowledge center is one characterized by a premium on models, typologies and invariant rules. To return to our trivial example, the rule-based knowledge center would maintain that anyone (regardless of experience) can navigate Interstate 95 successfully simply by passing a driver's test: learning a set of invariant rules including drive on the right, don't pass on the right, leave the interstate at numbered exits, obey the speed limit, etc. The rules-based knowledge center, like its opposite number the case-based knowledge center, is very normative: theory is superior to practice, and as such book-learning is a far more efficient method of imbibing knowledge than is experience. The rules-based knowledge center recognizes the pattern-based knowledge center as pseudo-science, and recognizes recipes and heuristics as rules-in-the-making that need only the gaze of a practitioner trained in "science" to reduce the fuzziness of heuristic to the clarity of rule.

Thus we can see the spectrum recapitulating the two faces of logic, as indicated in the diagram below.

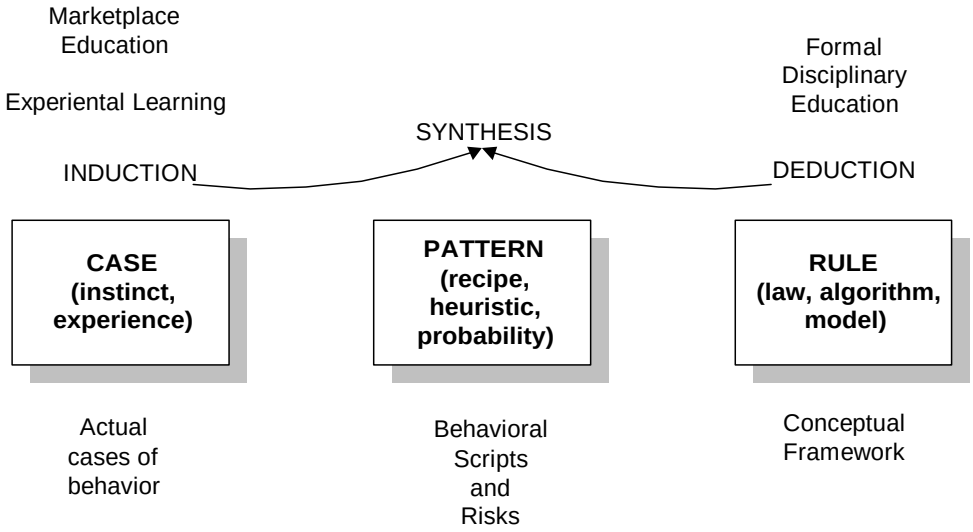

Now, proponents of each knowledge center will, presented with such a model, always argue that their knowledge center makes provision for the other two, and is therefore the superior model. But the point of the model **is not** to make invidious discriminations between the three models in an attempt to set up one knowledge center as superior to the other, **since any discipline** – from jazz piano to quantum mechanics – **is at any given time a mix of case, pattern and rule**. All three types of knowledge are essential to top-flight performance. The point of the model is rather to point out, to practicing knowledge managers, that any given SME with whom the KM is likely to work has a deeply-seated knowledge center that is, at root, one of the three above – the SME that recognizes the value of all three models of knowledge acquisition is rare indeed. For most SMEs, their knowledge center is their theory-in-action of how knowledge is acquired (their descriptive theory) and how it ought to be acquired (their normative theory), and approaches to an SME based on a different knowledge center are bound to fail from the outset because the SME cannot (a) understand and/or (b) accept any program or process based on a theory-in-action of knowledge which is different from their own.

As an example, consider an exercise I undertook with sales people in a high-technology firm recently. I identified the top performers, using a combination of hard metrics (how well the

performers did against their quota and in league tables against one another) and soft metrics (the propensity of their customers to rebuy, and to repurchase in larger volumes). I then began to interview the top performers, stating as the purpose of the interviews the "formulation of principles of good account management." Almost to a person, the sales professionals interviewed responded by saying, "Well, to understand how I sell, you'd have to understand my entire professional career…" or "I can't explain why I do what I do – it's something I've picked up along the way…" or, more confrontationally, "what you're doing won't work. What I do works only for me."

None of the SMEs was wilfully obstructing my inquiry. They were in fact responding to an inappropriate positioning of the knowledge center of my inquiry. Instead of saying, "I'd like to understand how you've learned to sell the way you do, and what you believe the secrets of effective account management are…", a very case-based approach, I adopted a rule-based position – the formulation of principles of good account management – that was as far away from the sales professionals' case-based knowledge center as it is possible to be. We had, as the philosophers like to say, incommensurate vocabularies, and could not even discuss the topic. Sales professionals gravitate to sales in part because they have overwhelmingly case-based knowledge centers: specifics, detail, practicality, the here-and-now are highly valuable to them.

Knowledge Flexibility: Are There Other Ways Of Knowing?

In general (and I speaking here in the context of the West and its case-heavy pedagogical methods), knowledge managers will find that the case-based knowledge center dominates the commercial sphere: around 80% of SMEs are likely to deploy a case-based knowledge center in their lives and work, while 18% are likely to exhibit rules-based knowledge center behavior and 2% are less are likely to be pattern-oriented.

Even though any given discipline is, as it matures, likely to move from cases to patterns to rules (with an uneven, even bumpy progression), the practitioners of the discipline do not move with the discipline itself, by and large. People's fundamental knowledge centers – their theories-in-action – are set early in their lives and do not fundamentally change. Experiences that people have that reinforce their sense that their knowledge center is proper are assimilated as proof of the correctness of their theory-in-action, and those experiences that do not reinforce that sense are discarded as exogenous.

Occasionally, the KM will encounter someone who exhibits all the characteristics of a particular knowledge center, but who is comfortable with other ways of knowledge: for example, a case-based SME who can readily describe the rules that shape her performances, or a pattern-oriented SME who can define what she thinks the emergent rules in her patterns are, as well as the specific characteristics of specific cases that do not fit the pattern she uses. In other words, some people – a minority – appear to have high degrees of knowledge flexibility (KF) with respect to the knowledge centers they are asked to work in: they can talk cases, talk patterns or talk rules as the situation demands. Ultimately, a person's knowledge flexibility index (KFI) answers the question "are there other valid ways of knowing other than the one you habitually employ?" A person with a high KFI is effectively answering, "Yes," while a person with low KFI – a rules-based thinker who can see no value in exceptional cases (edge-casing) for example – is effectively answering, "No."

Employing KC and KFI Models For Effective Knowledge Engineering

KMs Must Know Their Own Minds First

Practicing knowledge managers, too, have their knowledge centers and knowledge flexibility indices. It's safe to say that no successful KM comes from the hard-line case-based knowledge center (since, in that position, it would be impossible for the KM to do her job, since by definition

she could not scale knowledge), and it's equally safe to say that successful knowledge managers have high KFIs, the first task of a knowledge manager is to understand the blindnesses and insights of her own knowledge center. Over-emphasis on anecdotes, practicality, how-to's, reified lessons learned and a distaste for theory indicate a KM with a case-based orientation. An unwillingness to get one's hands dirty with the details of particular SMEs' performances, as well as the sense that "when the model's done, the work is done," indicates a rules-based orientation. What the KM must do, at the outset, is work on her own flexibility, so that she can understand the real and unique value that cases, patterns and rules bring to every performance and to every discipline. Otherwise, KMs find, in any situation, what they themselves bring to that situation: their own biases.

Assessing The KC and KFI Of Subject Matter Experts

The best way to measure[4] an SME's KC and KFI is to ask direct questions about (a) their models of how knowledge is acquired and assimilated and (b) their beliefs about the values of other models of knowledge acquisition. For example, given the task of defining "best practices in marketing collateral development" and a defined set of "best practitioners", a KM should begin the knowledge engineering project with a set of pre-interviews, using a script that looks like this:

1. ***You have been identified as a top performer in the area of marketing collateral development. How did you master this skill*?** Allow the SME to self-diagnose and identify "why" she has mastery. Where she locates mastery – in her experience or life-history, in a book she read, in a process she follows, in a knack she has – tells the KM what the KM needs to know about the SME's knowledge center, and possibly about the SME's KFI (if for example the SME identifies multiple sources of mastery, including case, pattern and rule sources).

2. ***If you were going to go about reproducing your mastery in other people, what would those other people need to have in the way of skills and experience and knowledge, and how would you go about teaching them to do what you do*?** This question is open-ended with a reason; the KM needs to see where the SME focuses (on prerequisites, which indicate case-orientation, or on teaching method, which indicates rules-orientation), as well as what specific prerequisites and teaching techniques or methods the SME suggests.

3. ***I am attempting to [state the purpose of the knowledge engineering exercise appropriately]. What guidance would you give me, and what approaches do you think I should take?*** This question, also open-ended, allows the SME free space in which to deploy their knowledge center while offering the KM guidance and advice.

Obviously, "I don't know" and "I've never thought about it" answers, and passive-aggressive behavior in general, are indications that the SME should be weeded out of the SME pool immediately. More generally, the way the SME responds to the questions – with short answers ort with detailed digressions – indicates something about their KFI.

What you will end up with, at the end of these pre-interviews, is a good sense of the KC and KFI of each prospective SME, which you can characterize using the model below:

---

[4] I am sure it is possible to come up with a standard survey instrument that quantifies KC and KFI; I am unsure that instrument will have any more value for knowledge managers than would a good pre-interview chat with the SMEs in question.

Knowledge Center	Knowledge Flexibility Index	Tendency	Type
CASE-BASED KNOWLEDGE CENTER	LOW	Universal Law	**A**
	MODERATE	Cases	**B**
	HIGH	Patterns	**C**
PATTERN-BASED KNOWLEDGE CENTER	LOW	Cases	**D**
	MODERATE	Patterns	**E**
	HIGH	Rules	**F**
RULE-BASED KNOWLEDGE CENTER	LOW	Patterns	**G**
	MODERATE	Rules	**H**
	HIGH	Edge-Case Experimentation	**I**

Positioning The Knowledge Engineering Exercises With SMEs

After you have done pre-interviews, assessed SMEs' KCs and KFIs, and are ready to conduct the knowledge engineering exercises proper, the KCs and KFIs of your SMEs help you position the knowledge engineering exercise itself, as indicated in the table below:

If the SME is	With a low KFI, then	With a moderate KFI, then	With a high KFI, then
Case-oriented	Position the exercise as case documentation (have the SME pick several cases) with the KM acting merely as a scribe	Position the exercise as an attempt to extract from the SME's experience some loose "recipes" for the SMEs peers	Explain precisely what you are doing and what you hope to accomplish
Pattern-oriented	Position the exercise as an attempt to extract from the SME's experience some loose "recipes"	Position the exercise as an attempt to produce first-pass "rules of behavior" for someone doing with	Explain precisely what you are doing and what you hope to accomplish

	for the SMEs peers	SMEs job without their experience and "sense of the job"	
Rule-oriented	Position the exercise as an attempt to produce first-pass "rules of behavior" for someone doing with SMEs job without their experience and "sense of the job"	Position the exercise as an attempt to both define "rules of behavior" and to call out edge-cases that provide illustrations of exceptions to those rules	Explain precisely what you are doing and what you hope to accomplish

Of course, this implies that the KM is interviewing the SMEs on a one-on-one basis, which seems to be the best technique for initial data-gathering.

In short, the general rules for conducting knowledge engineering sessions are:

1. Know your own KC bias and make sure you don't contaminate the session from the outset.

2. Talk cases with case-oriented SMEs, patterns with pattern-oriented SMEs, and rules with rule-oriented SMEs, particularly when they exihibit low KFIs.

3. Try to move the SMEs "right" one position in the KC spectrum. Test "patterns" with case-oriented SMEs ("would that experience suggest a pattern like …"); test rules with pattern-oriented SMEs ("So, is it safe to say that in the vast majority of situations, one should do…"); and test edge-cases (exogenous cases) with rule-oriented SMEs (well, what about this situation, which seems to run counter to your rule…).

KCs And KFIs Help You Validate The Results Of Knowledge Engineering Activities

KCs and KFIs also help KMs validate the synthesized results of knowledge engineering activities, by providing KMs with a "validation" model to be used with the original SME set, as indicated in the chart below.

If the SME is	**With a low KFI, then**	**With a moderate KFI, then**	**With a high KFI, then**
Case-oriented	Do not use the SME to validate findings	Have them test a set of performance patterns you generate against their "experience base"	Have them review your findings as you have embodied them
Pattern-oriented	Do not use the SME to validate findings	Have them test a set of performance rules you generate against their patterns	Have them review your findings as you have embodied them
Rule-oriented	Do not use the SME to validate findings	Have them test a set of performance edge-cases you provide them with against the rules	Have them review your findings as you have embodied them

Using each category of SME to validate the "practice logic" in a form that is not their own allows you to leverage the innate talents of each category of SME: the case-oriented SMEs can identify

quickly all the reasons why the patterns you have generated are incomplete; the pattern-oriented SMEs can poke holes in your rules; and the rule-oriented SMEs can anaylze edge-cases far better than the KM could.

Knowledge Engineering: The Process

We thus arrive at a process model for conducting successful knowledge engineering sessions in in the context of a knowledge management program that looks something like this:

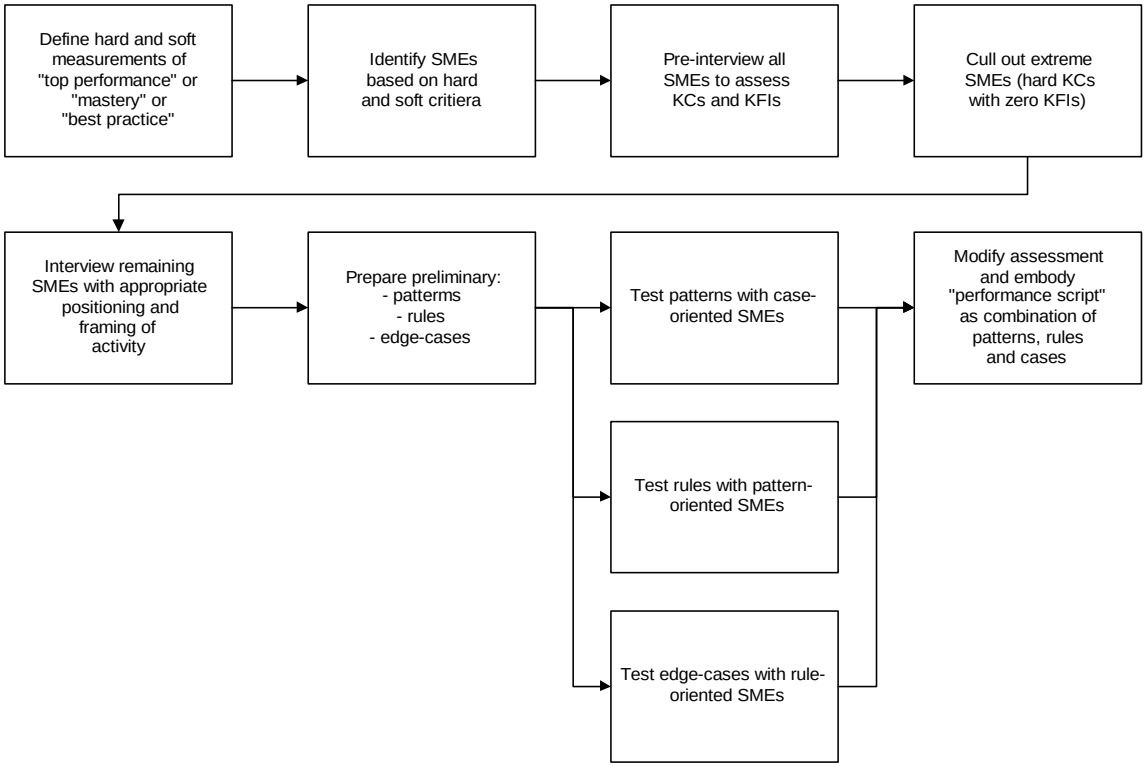

Coda

Knowledge engineering failed us once, in part because of the overly simplistic models we brought to the interactions with the subject matter expert, and in part because we failed to recognize that the knowledge engineer had as many biases and blindspots as the SMEs she was interviewing.

As is the case with most interpersonal work, successful knowledge managers struggling to scale and distribute “best practices” need to have:

- A working model and theory of their work that they are not afraid to modify
- A clear sense of theirown blindnesses and insights
- The ability to tailor their programs to the needs, models and vocabularies of the people with whom they are working.

Without these things, KM will simply be knowledge engineering redux, with predictable results.