

This is part 1 in a series of articles on the harvest and reuse of expert knowledge. This article focuses on the harvesting process itself. Other articles focus on facilitating the practical application of expert knowledge by non-experts and the use of DCT to significantly expand the value of performance improvements..

Harvesting the Experts' "Secret Sauce" and Closing the Performance Gap

William Seidman, Ph.D.
Michael McCauley

Executive Summary

A Couple of Quick Questions...

Do some of your people consistently and systematically outperform others?

Is there financial or strategic value in raising the performance of the many to the level of the few?

If you answered "yes" to these questions, then gathering and using the "secret sauce" of these top performers, these experts, will be of considerable value to you.

Main Topics

- The top performers have a "secret sauce" that consists of their mental models, subtle cueing mechanisms, action plans, risk management strategies and many other capabilities
- The secret sauce is largely unconscious, but is the real difference between top performers and less effective personnel
- The secret sauce has the following performance improvement impacts
 - 80% reduction in planning time
 - 50% reduction in training time
 - 30-50% reduction in task performance time
- It is easy to identify the experts and motivate them to participate
- The expert secret sauce consists of:
 - An "ecological survey" that frames the situation for the expert
 - Patterned behaviors which are highly efficient actions required to be successful in the situation
- Even for such complex areas as launches of interplanetary vehicles, knowledge harvesting takes a maximum of 30 hours, and the content is ready to "coach" less experienced personnel
- It is essential to have a plan to use the secret sauce in order to create significant value
- Digital Coach Technology is critical to knowledge harvesting and reuse
- Companies using DCT generate millions in additional sales and profits

Harvesting the Experts' "Secret Sauce" and Closing the Performance Gap

A Couple of Quick Questions...

Do some of your people consistently and systematically outperform others?

Is there financial or strategic value in raising the performance of the many to the level of the few?

If your answer is "yes" to these questions, then gathering and using the "secret sauce" of these top performers, these experts, will be of considerable value to you.

Top performers' secret sauce is their true tacit knowledge. It is their mental models, subtle rules of thumb, early problem detection mechanisms, attention allocation processes, and actual work efforts that enable them to consistently outperform others. In addition, it is a deep, passionate commitment to the success of the process that causes the experts to apply their specialized knowledge more intensely and profoundly than non-experts.

Is the experts' secret sauce worth exploring? You bet! Companies that systematically harvest and apply the experts' secret sauce consistently report spectacular performance gains such as:

- 80% reduction in planning time
- 50% reduction in training time
- 30-50% reduction in task performance time

In this article, we are going to help you quickly identify the experts in your organization, explore the methods that they use (their "secret sauce"), effectively harvest this secret sauce, and successfully reuse it to enhance your organization's productivity.

But Who Are the Experts?

Surprise! Everyone already knows who the experts are in any given function or domain. People already have an exceptionally rich, highly consistent, comprehensive mental map of the expertise in the organization. Consequently, instead of having to purchase and develop expensive, ineffective profiling systems, all you have to do is ask a few key managers or senior professionals a simple question: "Who are your experts?"

When we asked for a list of experts in a performance management process at a semi-conductor manufacturing company, the VP of Corporate Human Resources initially said that he couldn't possibly isolate the top performers. However, after a few seconds, he paused and stated, "Actually, these are the top performers," and proceeded to list 10 people. This is typical. We've found that, to an amazing degree, most associated people in the organization will come up with a virtually identical list of names.

Corporate knowledge of these “real” experts has two significant implications for achieving the ultimate goal, an improvement in performance. First, because there is such a high level of agreement on who the experts are, their participation in the process creates instant credibility. Just because they are participating, many others will support the initiative. Second, the content they produce is likely to be more robust and meaningful and will be treated by everyone as more valuable. This quality is critical for acceptance by less effective personnel.

Getting the Experts Involved

Conventional wisdom says that you must provide the experts with some incentive to participate. Nonsense! It is easy to get the experts involved if you tap into their intrinsic motivation and professionalism. You should begin by asking them the following simple questions:

- What percentage of your job is in some way a repetition of previous work? (*Usually the answer is 60-80%.*)
- How do you like doing these repetitive tasks? (*They're "OK."*)
- What do you do with the other 20-40% of your time? (*"Meet with people." "Design things," etc.*)
- How do you like doing these things? (*"These are fun."*)
- Which do you think drives your compensation more, the repetitive things that presumably others can do or the creative things where you really add your expert value? (*"The creative things."*)
- If you could use an approach that allowed you to spend less time doing the repetitive things that you don't really like and more time doing the creative things that you do like, would you be willing to participate in that approach?

Not surprisingly, the experts willingly participate every time.

The Secret Sauce

The structure and content of the experts' secret sauce is actually quite different from the foundations underlying current systems. In particular, experts do not use branching probabilities (the foundation of decision trees, artificial intelligence or expert systems). They also don't rely on generalized searches of databases for appropriate content (the foundation of most enterprise knowledge management systems).

Instead, experts employ a much simpler, more robust and focused model that has only two elements:

- An “ecological survey” that points them to a limited number of core behaviors
- Patterned behaviors based on the survey results

Ecological Surveys

The single most critical aspect of the experts' secret sauce is the categorization process the experts use to organize their work. Typically, experts only require about 90 seconds to assess their environment and categorize any situation into

one of only three categories. Further, they categorize using just one or two key factors and label the categories using very simple language.

For example, safety inspectors for an insurance company enter a customer's premises through the loading dock, observe the trucks parked there, and quickly categorize the customer as having "clean trucks," "dirty trucks" or "other." Using this "clean truck rule," the inspectors can predict with almost complete accuracy the safety conditions they will find at the customer's site.

Similarly, an expert in the transfer of family owned business assets from the founders to the next generation, categorizes his customers' businesses into "high trust," "low trust" and "other." In just 10 seconds, based on the tone of voice (e.g. whispering indicates low trust) and the terminology used by the founder, he can categorize the situation and predict how to manage the transition.

Experts categorize quickly and efficiently using subtle criteria. They are "unconsciously competent" and know from experience both the correct assessment and the implications of the assessment on their follow-up behaviors.

The following questions can help convert the experts' unconscious competence into something more tangible. Just ask the expert:

- Do you organize your work into a few main categories? (Assuming they said: "yes", use some simple, plain language to label their categories)
- How quickly do you know you are in one of the categories? (Usually "very quickly")
- What tells you that you are in one of those categories?

There is a remarkably consistent pattern in the frequency of occurrence of the categories. For each expert, the leading category almost always occurs about 70% of the time, the secondary category 25% of the time, and the remaining category (usually "other," "miscellaneous," or "don't know") only about 5% of the time. To continue the above examples, dirty trucks and low trust environments occurred 70% of the time, clean trucks and high trust environments occurred 25% of the time and some other situation occurred only 5% of the time. To further simplify the analysis, the 25% category is almost always just a derivative of the 70% category. Thus, by mastering the cueing mechanism and associated behaviors of the 70% category, experts become extraordinarily efficient at making quick decisions.

If this categorization process is so robust, why do so many people think their situations are completely unique? We humans appear to be genetically wired to focus more on differences in our environment than on similarities. Many knowledge management systems reflect this bias. They are actually focused on managing the 5% "other" category (which cannot be done cost-effectively in an automated system and which also dilutes the focus on the 70% category) more than the 70% category.

Focusing on differences is a reasonable survival mechanism. However, if we had to uniquely recreate an entire process or function each time we needed it, we would be overwhelmed and would never progress. Instead, we systematically group our work so we can become increasingly efficient at performing it. The top performers even account for variation in their 70% models.

The message is clear. Identify and concentrate on your 70% category. If you master that, then you master the core of your business.

Patterned Behaviors

Based on the 70% category, the experts have a set of precise, patterned behaviors that make them the top performers. These patterned behaviors include (Figure 1):

- A set of mental models about the outcomes of the process
- A strong focus on action
- A set of threat monitoring and management strategies
- An extensive, but highly focused, set of references to supporting materials and resources

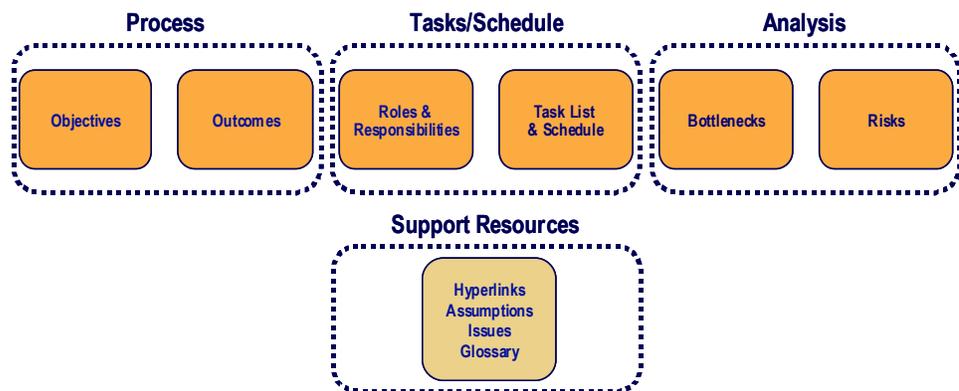


Figure 1: Patterned Behaviors

Mental Models

The single most important element in the experts' patterned behavior is their mental model. However, they don't usually know they have a mental model and it is difficult for the new person to grasp the model because it is both unconscious and defined using highly abstract thinking. Understanding the experts' abstraction is critical to understanding their secret sauce.

Experts consistently conceive of the ultimate objective of their work in more visionary, moralistic, and exciting terms than less effective performers. For example, the top performing managers at a manufacturing company conceptualized the outcomes of a performance management and pay administration system as "refreshing and energizing the organization," while the less effective managers thought of it as doing "pay administration."

Most people would much rather "refresh and energize the organization" than "administer pay!" It is easy to see why the top performers distinguish themselves.

The experts refine their mental models by organizing this energetic objective into groupings of work required to achieve the objective. They further define these groupings using highly compressed vocabulary (what the work grouping "Is") and early indicators of problems situations (what the working grouping "is not").

For example, two of the work groupings for "refreshing and energizing the organization" are:

- "Leveling" the management "ratings" of individuals
- Creating development plans for the leaders

The abstractions supporting the first grouping include the definition of a compressed term, "Leveling" ("Creating common agreement on the criteria used to identify leaders"), and an exclusion (Leveling does not mean "arguing for your people regardless of the criteria"). The mental model is completely, if unconsciously, defined from highest-level objective to detailed definition of concepts.

The expert mental models are highly organized, effective and are the absolute foundation of actions. If you don't understand the mental model, you cannot apply the content in any meaningful way.

Actionable Focus

Based on their mental model, the experts always know who is supposed to do what and when. They are very accountability and action-oriented, carrying remarkably comprehensive and accurate, but generic, maps of roles and responsibilities, task definition, and work flow, as well as others. These maps fully define everything required to achieve the objective.

For example, the safety inspectors mentioned earlier know that there will be a "Director of Safety" who will have theoretical responsibility for the safety program and a Vice-President of Operations who will be the real authority. They will also know that implementing a change in safety-related behaviors requires specific activities such as having the Vice-President of Operations "walk the talk" (which needs to occur immediately) and establishing a better metrics system (which can come later).

These generic maps enable the experts to enter a specific situation and have a complete, applied plan almost instantly by simply substituting the specifics of the situation for the generic. Thus, the generic underlying process, the secret sauce, leads to the creation of an applied project.

Threats to Success

Experts know that success is rarely easy. From harsh experience, they have learned what is most likely to go wrong, how to prevent it from going wrong, how to detect when something is about to go wrong, and what to do about it if the problem occurs.

More specifically, experts have a strong, prioritized list of general threats to the success of the process. They can identify the biggest threats and have a predefined set of actions that both reduce the probability and the impact of the risk. In addition, the expert monitors subtle triggers of the risk condition. For example, a single phrase by a Vice-President of Sales to an expert sales consultant ("We have to get some other people involved.") cues the expert that the company is not supporting the newly implemented program and that an alternative implementation strategy is needed. A critical element of the secret sauce is the ability to quickly detect and manage risks.

Using Existing Resources

Experts are tremendously efficient at taking advantage of all the available resources in the organization, including reference material. However, unlike generalized search engines, the experts use this material in a highly targeted way. They say things like, “Now I show them this article,” or “Now I use the manufacturing release check list.” The experts know when and how to use specific materials.

Often these materials include extremely detailed process documentation. For example, the experts at one computer networking company maintained references to documentation on how to set router switches. Similarly, experts from a biotech company kept a mental list of all the federal forms required for FDA certifications of their production line. The secret sauce includes a detailed list of all of the various supporting references required for success.

The Total Secret Sauce

As may be apparent, the breadth and strength of the top performers secret sauce is quite remarkable. It begins with a quick categorization scheme and ends with tightly coupled linkages to other resources that provide extremely specific “how to’s.” It includes the subtle nuances of understanding found in the mental models, the specific work required to be successful and robust risk management capabilities. Best of all, the experts have integrated all of these things into a single, unified view of the process that can easily be applied to any specific situation.

Harvesting the Secret Sauce

Are you now saying to yourself, “This looks too difficult”? Don’t be alarmed. It is actually quite easy to quickly harvest the experts’ secret sauce. Using “the naïve new person role-play” and emerging “digital coach technology (DCT)”, even such complex knowledge as analysis and decision-making about interplanetary probes takes only about 30 hours to gather. Simpler processes can be harvested in as little as 6-8 hours. Furthermore, the secret sauce is ready for use by less experienced or successful personnel at the end of this time without significant additional work.

The knowledge harvesting process is best done in a workshop-like setting with a DCT application projected onto a screen. The facilitator begins the session by describing herself as the ideal new employee. She is ideal because she is intelligent, competent and motivated. She just doesn’t know how to do it “right” and needs the experts to coach her to perform the function correctly. Experts can’t resist coaching an eager naïve person.

The facilitator then asks “naïve new person questions” prompted by the projected DCT application. These questions are typical of the questions asked when a naïve new person joins an organization. They are virtually universal and evoke a specialized kind of story by the expert called a “naïve new person story.” The facilitator records the expert answers in the DCT application, both as a means of providing immediate feedback for the experts, and to slow the use of abstractions to a comprehensible pace.

Initially, the naïve new person story told by the expert is almost always an “official” version and does not yet contain their secret sauce. Official stories use relatively bureaucratic language and lack vision and energy. Most people will detect these “official” stories immediately and consistently reject them as unrealistic. In contrast, DCT is designed to cause people to question themselves and switch from the official story to the “real” story. Real stories are more colloquial, have a less formal syntax and are bursting with energy, vision and passion. This conversion occurs approximately 22 minutes into the knowledge harvesting process. These “real” stories are the secret sauce.

However, the secret cause in story form is too raw to be easily consumed by less experienced personnel. DCT applications guide the experts to adopt differing perspectives on their expertise, creating cognitive conflicts, which the experts resolve by asking themselves: “What is the *right* thing to do?” Focusing on doing the “right thing,” drives them both to a deeper interpretation of the secret sauce and the conversion of the sauce to a repeatable best practice. When developed this way, the best practice is highly useable by less experienced personnel.

Once the experts become completely engaged in the knowledge harvesting process (this takes about 3 hours), it is actually quite difficult to get them to stop talking. As one expert put it: “This is the best work I have done in five years! I don’t want to stop!” In some instances, the facilitator has had to turn off the projector and get extremely forceful to get the experts to break for lunch or end the day. Expert participation is never a problem.

This knowledge harvesting process generates several deliverables and one great opportunity. The deliverables are:

- A significantly improved expert process
- Documentation of the expert secret sauce
- A DCT database which stores the expert content

The DCT database provides the opportunity for improved organizational performance by enabling the use of the secret sauce by less skilled or successful personnel.

Using the Secret Sauce for Performance Improvement

Do you plan to use the secret sauce you just gathered? This is not a frivolous question. Too many people become overly focused on gathering knowledge with little regard for its use. Yet, the benefit of the secret sauce only occurs when it is used to improve performance.

Unfortunately, it can be difficult to get people to use the secret sauce. In most cases, a naïve person judges the value of new knowledge in only 90 seconds, stopping further use if the content is found wanting. To get past that 90-second barrier, the user must immediately perceive that the content is credible and begin to apply it to their situation. Credibility for the content comes from the “real” stories that have been polished into best practices.

Application of the content occurs when DCT guides the naïve user to emulate the experts’ process for converting the generic to the specific. DCT “coaches” them to become introspective about the application of the best practice to their own

situation in a way that increases the users "ownership" of the best practice, which, in turn, results in behavioral changes consistent with the secret sauce.

Typically, this change occurs in about eight minutes. The naïve users lean forward more, concentrate their focus on the content, increase their respiration, and use more intense language. Effective DCT can drive behavioral change.

Alternatives Approaches

Are there alternatives to focusing on the secret sauce or the approach to gathering it presented here? We don't think so – at least not today! Many things have been tried in the recent past, but few have succeeded.

The most common past approaches have been lengthy interviews (often videotaped), and/or following the experts around with extensive process mapping during the data-gathering period. These labor-intensive approaches have produced thick process binders and lengthy training courses. These have failed for many reasons, including:

- They only capture and display the "official" story (which is rejected by users)
- At best they capture the actionable portions of the expert behavior and consistently miss the critical categorization schemes, mental models and cueing mechanisms
- They are fairly rigid and therefore difficult to apply to real situations
- They are both disruptive of the expert behavior and invasive of their professional domain, and are therefore resisted by the expert
- They are extremely labor intensive and therefore quite expensive, often taking many months of work

For all of these reasons, and many others, most prior efforts to simply "capture knowledge" have not been successful.

Two other, more indirect approach to harvesting and using the experts' secret sauce have been tried as well; the establishment of "communities of practice" and the use of personnel profilers. In communities of practice (CoP), an electronic bulletin board is established in which the experts theoretically come together to share their expertise. When a naïve person needs some critical information, they join a CoP talk session and pose a question to the community. The presumption is that an expert participating in the CoP will respond with at least a portion of his secret sauce.

Personnel profilers are based on a similar core assumption, that the experts are willing to participate. In personnel profilers, each expert has a profile describing their area of expertise. When the naïve person needs assistance, they initiate a query into the profiler, which provides a match between the naïve person's question and the expert's profile, usually using some form of key word search capability. The system then asks the identified expert if they would accept a question on the topic. Assuming the expert is willing to participate, they respond, providing some of their secret sauce to the questioner.

Unfortunately, neither of these approaches has worked. In both cases the experts drop out of the system quickly. They get bored with the low level of content in the CoP or they become inundated with simplistic questions in the personnel

profilers. In both cases, they quickly realize that they are being pulled away from their main functions and that their compensation is likely to suffer. Their initial interest and participation rarely lasts more than a few months.

In contrast, using DCT, the experts' participation is usually limited to a three-day session (or less) that they enjoy, from which they benefit, and that produces superior content.

Benefits of the Secret Sauce

Does this matter? YES, the impacts are remarkable! For example, by applying the DCT processes described here:

- A semi-conductor manufacturing company expects to reduce training time 66%
- A photography company said that they, "Took three days to accomplish what had previously taken six months."
- An audiology consultant has boosted revenues in a string of clinics from the industry average of \$400k per year per clinic to \$1.3M per year per clinic
- A fast food chain projects that it will increase the sales in its less effective restaurants \$2,000 per restaurant per week

Harvesting and reusing the experts' secret sauce does make a difference. Now, go harvest your own experts and begin reaping the benefits!