

“Deep” Knowledge Management and Sustainability

By Mark W. McElroy

How can we get commodity industries (and businesses in general) to behave in more sustainable ways? One prevailing, if not dominant, theory of practice is to highlight the behaviors we object to (e.g., pollution, habitat destruction, etc.); reveal the mental models that lie behind them; and somehow in the process get managers in the offending companies to acknowledge their errors and change their ways. In practice, however, this approach has not been very successful. Is it possible, then, that another ‘theory of practice’ might be more effective?

Consider, for example, that business behaviors are really nothing more than *business knowledge in use*. What if instead of trying to change people’s operating behaviors, we were to focus on changing the separate and *different* behaviors that account for knowledge production and integration – not just knowledge *use*? Indeed, we *can* differentiate between business- and knowledge-processing behaviors. In fact, enhancing knowledge processing is precisely the idea that lies behind second-generation (aka, ‘the new’) knowledge management (KM). Here, however, there’s more to the story. If we also take steps to encourage and support a high degree of sensitivity for ecological considerations whenever we engage in knowledge processing, what follows is a new theory of practice for achieving sustainability called “*Deep*” *Knowledge Management* – an extension of second-generation KM inspired by Deep Ecology.

Deep Ecology

One of the more interesting and sophisticated products of the environmental movement in the 1960s was a branch of philosophy called *Deep Ecology*. Described by George Sessions in his 1995 anthology, *Deep Ecology for the 21st Century* (Sessions, 1995), Deep Ecology’s “main concern [was and is] to bring about a major paradigm shift – a shift in perception, values, and lifestyles – as a basis for redirecting the ecologically destructive path of modern industrial growth societies” (Sessions: ix).

Sessions points out that the Deep Ecology movement was also partly inspired by the work of Rachel Carson, in whose well-known book, *Silent Spring* (Carson, 1972), some of the first environmentally-motivated attacks on anthropocentrism first appeared. Carson challenged the direction and goals of Western society, including the presumed right to dominate and manage the earth. She wrote that “the ‘control of nature’ is a phrase conceived in arrogance, born of the Neanderthal age of biology and philosophy, when it was supposed that nature exists for the convenience of man” (Carson: 297).

The key distinction between *shallow* and *deep* approaches to ecology was first articulated in 1972 by the Norwegian philosopher, Arne Naess,¹ who is still today

“Deep” Knowledge Management

widely regarded as the founder of Deep Ecology. Shallow Ecology, according to Naess, was (and is) marked by a “[f]ight against pollution and resource depletion” (Naess, 1973), especially with “the health and affluence of people in the [first world] developed countries” (Ibid.) in mind. Shallow Ecology is therefore not only anthropocentric, but is marginally ethnocentric, as well. Affluent societies, where Shallow Ecology reigns supreme, are prone to externalizing both the sources and costs of their pollution by exporting them to less developed, third-world countries.

Deep Ecology, by contrast, is *ecocentric* in its orientation, not just *homo-* or *anthropocentric*. It includes concern for the welfare of humans, but is not confined to it. Naess and his supporters, then and now, sought to position Deep Ecology as an ecological philosophy, or *ecosophy*,² that rejected anthropocentrism in the broad conduct of human affairs and replaced it with ecocentrism. To Deep Ecologists, there are no nation states, boundaries, or ethnic strains of humanity that call for favored status, nor are there first-, second-, or third-world countries. There is only Gaia, the comprehensive biosphere, and all life within it, the holistic health of which should be our chief concern.

To better appreciate the thinking behind Deep Ecology, it is worth taking a moment to review its key principles (i.e., its ‘platform’), first formulated by Sessions and Naess in 1984, then published by Naess in 1986. Here they are (Naess, 1986):

1. The well-being and flourishing of human and non-human life on Earth have value in themselves (synonyms: intrinsic value, inherent worth). These values are independent of the usefulness of the non-human world for human purposes.
2. Richness and diversity of life forms contribute to the realization of these values are also values in themselves.
3. Humans have no right to reduce this richness and diversity except to satisfy vital needs.
4. The flourishing of human life and cultures is compatible with a substantially smaller human population. The flourishing of non-human life *requires* a smaller human population.
5. Present human interference with the non-human world is excessive, and the situation is rapidly worsening.
6. Policies must therefore be changed. These policies affect basic economic, technological, and ideological structures. The resulting state of affairs will be deeply different from the present.
7. The ideological change will be mainly that of appreciating life quality (dwelling in situations of inherent value) rather than adhering to an increasingly higher standard of living. There will be a profound awareness of the difference between bigness and greatness.
8. Those who subscribe to the foregoing points have an obligation directly or indirectly to try to implement the necessary changes.

“Deep” Knowledge Management

As for how the Deep Ecology platform should inform human behavior, Sessions wrote, “Finally, it is claimed by Arne Naess that the Deep Ecology movement should be thought of as being characterized by the *deep questioning process* [emphasis added], the Deep Ecology platform, by the needs for humans to identify with the non-human, and by non-violent environmental activism, which stems from having a ‘total view’ ” (Sessions: xxii).

Deep Questioning and the Self

In an interview published in 1982, Naess elaborated on the role of deep questioning for himself and other adherents to Deep Ecology, as follows (Bodian):

“The essence of deep ecology is to ask deeper questions. The adjective ‘deep’ stresses that we ask why and how, where others do not. For instance, ecology as a science does not ask what kind of a society would be the best for maintaining a particular ecosystem – that is considered a question for value theory, for politics, for ethics. As long as ecologists keep narrowly to their science, they do not ask such questions. What we need today is a tremendous expansion of ecological thinking in what I call *ecosophy*. *Sophy* comes from the Greek term *sophia*, ‘wisdom,’ which relates to ethics, norms, rules, and practice. *Ecosophy*, or deep ecology, then, involves a shift from science to wisdom.

“For example, we need to ask questions like, Why do we think that economic growth and high levels of consumption are so important? The conventional answer would be to point to the economic consequences of not having economic growth. But in deep ecology, we ask whether the present society fulfills basic human needs like love and security and access to nature, and, in so doing, we question our society’s underlying assumptions. We ask which society, which education, which form of religion is beneficial for all life on the planet as a whole, and then we ask further what we need to do in order to make the necessary changes. We are not limited to a scientific approach; we have an obligation to verbalize a total view.”

The obligation that Naess speaks of here is echoed in the last plank of the Deep Ecology platform: “...an obligation directly or indirectly to try to implement the necessary changes.”

In addition to deep questioning, Deep Ecology rests heavily on another important idea: Self-realization. But this is not the ‘self’ we normally associate with ourselves as individuals. Rather, the Self (with a capital ‘S’) in Deep Ecology is the totality of life. In Naess’s own *ecosophy*, which he calls *Ecosophy T*, he

“Deep” Knowledge Management

explains this orientation towards the larger self in the following way (Naess, 1986):

“Ecosophy T has only one ultimate norm: ‘Self-realization!’ I do not use this expression in any narrow, individualistic sense. I want to give it an expanded meaning based on the distinction between a large comprehensive Self and narrow egoistic self as conceived of in certain Eastern traditions of *atman*. This large comprehensive Self (with a capital ‘S’) embraces all the life forms on the planet (and elsewhere?) together with their individual selves (jivas).”

Reduced to imperatives, the essence of Naess’s ecosophy is: “Maximize (long-range, universal) Self-realization!”; “Maximize symbiosis!”; and “Maximize (long-range, universal) diversity!” (Ibid.). For Arne Naess and many other Deep Ecologists, the *end* is Self-realization and deep questioning is the *means*. It is the encouragement of deep questioning, then, that ultimately underlies Deep Ecology, “a social movement,” Sessions tells us, “whose purpose is to bring about the ‘profound social change’ that is necessary to resolve the environmental crisis” (Sessions: 7).

Deep Epistemology

Turning to another source on what it means to rigorously question our knowledge and our assumptions, we find the great twentieth-century philosopher Karl Popper, whose notion of the ‘Open Society’ (described in his book by the same title) was nothing if not a treatise on the importance of deep questioning, and of openness and transparency in human social systems. Indeed for Popper, life was all about asking questions in the context of solving problems. He stressed, therefore, the importance of creating and maintaining the political conditions in which open questioning could occur, questioning without fear of reprisal or the dark hand of coercion. We produce new knowledge, he believed, as a consequence of the questioning process, and it is crucial to our survival, then, that we make open questioning possible in the normal conduct of human affairs. In so doing, he said, we make it possible for ourselves *to kill our worst ideas before they kill us*.³

A similar ethic can be found in the work of another great twentieth-century philosopher, Charles S. Peirce. The originator of American Pragmatism, Peirce once wrote, “Do not block the way of inquiry” (Hartshorne and Weiss: 135-140) – our survival depends on it. He and many other pragmatists, such as William James and John Dewey, stressed the importance of inquiry – open inquiry – to the development of useful (and reliable) knowledge.

While Popper, Peirce, James, Dewey, and many others all stressed the importance of inquiry, questioning, and openness in related processes *for humans*, it wasn’t until Naess, Sessions, and the Deep Ecology movement came

“Deep” Knowledge Management

along in the early 1970s that the ethics of inquiry and ‘deep questioning’ were extended to other species, to ecology, and to the biosphere in general. Upon reflection, then, we could say that these events triggered the birth of a new *Deep Epistemology*. Why? Because value theories always play an important role in the production of new knowledge; and for the first time, thanks to Deep Ecology, the theory of knowledge processing by humans was effectively amended to take ecological considerations explicitly into account.

That Deep Ecology sees openness in questioning and inquiry as the path to sustainability is clear. But will adherence to this ethic necessarily lead to the outcomes we seek? Will enhancements in organizational learning and inquiry, for example, necessarily lead to sustainability in business outcomes? Perhaps not. Nevertheless, we can also ask ourselves if sustainability in business outcomes and human behavior can occur without them. Can we make the kinds of changes required to stem the tide of humanity’s destructive influence on earth *in the absence of* strong organizational learning in our affairs? I don’t think so.

Let us conclude this part of our discussion, then, by agreeing that openness in the processes of inquiry and questioning in human social systems is at least a necessary (but insufficient) condition for ecological stability on earth, sustainability in business, and the production of new knowledge that is at least closer to the truth. We may not always get it right, but we will never get it right if we constrain our capacity to continually ask the deep questions in an atmosphere of openness and trust, free from the threat of reprisal or the dark hand of coercion.

Knowledge Processing

By all accounts, businesses (especially transnational ones) are the dominant human institution on earth, the influence of which is acutely felt both by humans and non-humans, alike. Wealth, poverty, corruption, war, famine, species loss, habitat destruction, biospheric pollution, ozone layer depletion, desertification, global warming, endocrine and genetic disruptions – all of this and more is either mostly caused or reinforced in one way or another by business and the imprint of commercial human affairs. I won’t attempt to make that case any further here. There’s plenty of that found elsewhere. Let us simply regard this claim as a premise for current purposes.

Now comes knowledge – human knowledge. What did Popper mean by the idea that we must make it possible *to kill our worst ideas before they kill us*? Popper knew, as we all do, that our behaviors, actions, and practices are knowledge-based and that our knowledge is what we enact in practice. We express it in our behaviors. Thus, bad practice is bad knowledge in use. Bad knowledge in use, then, if left unchecked over time, can undermine our very survival – left unchecked for too long, and it can kill us. It is critically important, therefore, that

“Deep” Knowledge Management

we make it possible for ourselves to continuously subject our ideas to open criticism and to ‘kill’ the worst of them, lest they, instead, undo *us* in the end.

Popper also believed that our knowledge is fallible – all of it. There is no certainty of truth in humanity, no ‘justified true belief.’ We cannot prove truth, but we *can* avoid falsity. And as we do, he argued, we can continually eliminate the errors in our thinking, and thereby move steadily closer to the truth. Popper felt that the very fallibility of our knowledge makes it imperative that we maintain conditions of openness in our knowledge processing affairs (i.e., the processes we enact to produce and integrate our knowledge). Our ability to discover and eliminate errors in our thinking depends on it. We must, therefore, maintain the conditions needed for continuous deep questioning (per Naess) and inquiry (Peirce) – in other words, we must create and maintain the conditions required for open conjecture and refutation.

But what about business? What about openness in knowledge processing in the conduct of human economic, commercial, and organizational affairs? Is there a case to be made for the ‘Open Enterprise’? (McElroy, 2003: 20-24; Firestone and McElroy, 2003: chapter 10). Before considering this question further, let me first try to clarify the all-important distinction between *business processing*, *knowledge processing*, and *knowledge management* in organizations (see Figure 1).

Turning first to the middle of the trichotomy,⁴ *knowledge processing* behaviors are epistemic in scope; they include such things as learning, innovation, inquiry; the process of testing and evaluating ideas; subjecting them (ideas) to conjecture and refutation; sharing them; and other forms of knowledge production and integration. The outcomes of knowledge processing are the *beliefs* and *knowledge claims* we then practice or enact, such as business strategies, organizational models, business processes, and so forth.

Business processing behavior, by contrast, is operational in scope and involves none of that. It consists only of knowledge *use*, most often in the form of value-chain activities and commercial transactions between business, their suppliers, their regulators, and their customers .

Knowledge management, then, can be seen as a category of behavior that sits ‘above’ knowledge processing, and whose intent, thereby, is to oversee and ensure the quality of our epistemic endeavors – hence, my designation of KM as the ‘meta-epistemic’ level of behavior. The purpose of knowledge management is to enhance organizational knowledge processing.

Business processing typically differs from the other two behaviors in that our operating routines are often normative and proceed in accordance with existing knowledge. From the perspective of organizational learning theory, this can be seen as ‘single-loop learning’ (Argyris, 1991).

“Deep” Knowledge Management

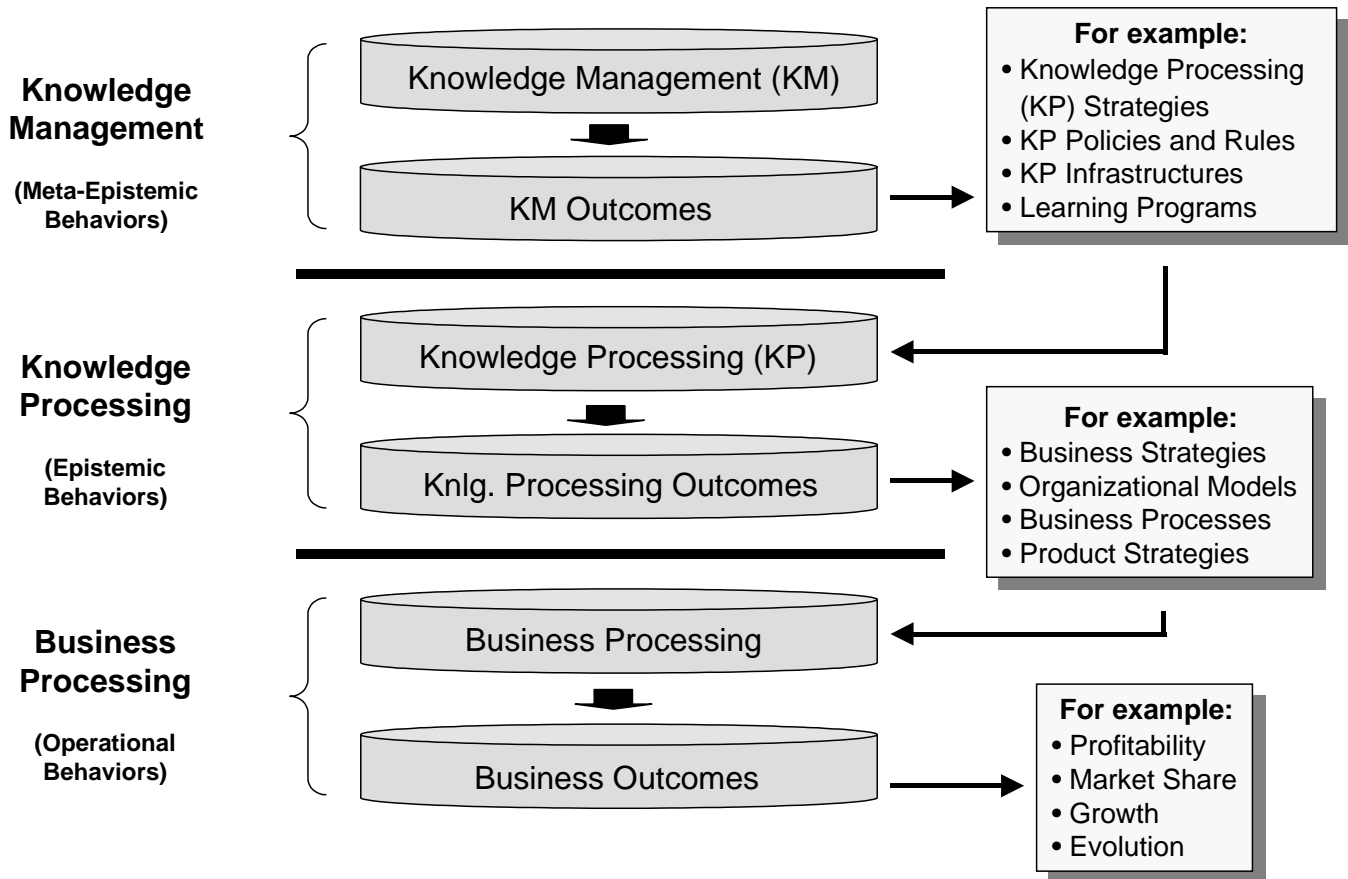


Figure 1 - The Three Levels of Behavior in a Firm

Knowledge processing behaviors, on the other hand, may or may not be normative (i.e., they may be emergent). Indeed, epistemic behaviors are the processes we spontaneously enact in order to create and integrate our knowledge. They are therefore akin to what organizational theorists refer to as ‘double-loop learning’ (Ibid.). When we engage in knowledge processing or epistemic behaviors, we effectively suspend our business processing behaviors and put them aside. We only return to them after new knowledge has been produced and integrated. In some cases this takes years; in others, milliseconds.

Armed with a metaphorical set of tri-focals, then, we can discern in organizations three kinds of behavior: (1) behaviors related to ‘getting the work done,’ (2) behaviors related to learning, and (3) behaviors related to enhancing our *capacity to learn and adapt*. The first ones are *business processing* behaviors (operational); the second ones are *knowledge processing* behaviors (epistemic); and the third ones are *knowledge management* behaviors (meta-epistemic).

“Deep” Knowledge Management

Now comes the challenging question of how, if at all, it may be possible in practice to separate epistemic behaviors from operational ones. In other words, can we have different rules for each so that organizational learning can occur independently of organizational management? Can we, for example, have a ‘closed’ command-and-control management system in a company in which the knowledge processing routines are ‘open’? Or must all knowledge processing behaviors necessarily be carried out by managers in power? This, of course, is a knowledge management problem because it transcends the epistemic and operational levels and is concerned, instead, with the relationship between them.

For the sake of argument, let us assume that the two (*knowledge* processing and *business* processing) can be separated, and that it is possible to have a firm that is ‘managed’ in every conventional sense of the term, while its knowledge processing behaviors are not. Let us further assume that this kind of enterprise (i.e., an *Open Enterprise*) is desirable because of its ability to, (a) support a broad range of management styles and, (b) enhance an organization’s ability to learn and adapt. This, then, would be an enterprise implementation of Popper’s ‘Open Society’ and Peirce’s notion of unfettered inquiry – precisely the kind of enterprise that Naess’s vision of ‘deep questioning’ implies. Why? Because any chance we have “for redirecting the ecologically destructive path of modern industrial growth societies” (Sessions: ix) is clearly dependent upon the degree to which openness in knowledge processing can be found in the dominant human institution on earth we call ‘business.’

“Deep” Knowledge Management

Meanwhile on an entirely different front, a new body of thought and practice in the not-so-new field of knowledge management has emerged under the moniker of ‘second-generation KM.’ (McElroy, 1999: 86-88). Why not-so-new? In the early 1900s, management guru Frederick Taylor made a legend of himself by advocating a style of management he called ‘Scientific Management’ (Taylor, 1947). Key to Taylor’s approach was the systematic capture, codification, and repetition of ‘best practices,’ an approach to management that later returned in the form of business process re-engineering in the 1990s, and then later still as ‘KM.’ Taylor never referred to his methodology as ‘knowledge management,’ but most of what passes for KM today is unmistakably Tayloristic. Indeed, Taylor could easily be described as the father of knowledge management. But that was one hundred years ago!

Today, KM has evolved into two distinct bodies of practice: first-generation KM and second-generation KM. First-generation KM is Tayloristic. It begins with the assumption that valuable knowledge *already exists*, and it falls to managers, therefore (*knowledge managers*), to capture, codify, and distribute it for widespread imitation and reuse.

“Deep” Knowledge Management

Second-generation KM takes a decidedly different stance. First, it is fundamentally Popperian. There is no ‘justified true belief’⁵ or infallible knowledge, nor does knowledge ‘already exist.’ People produce it through an ongoing social process of conjecture and refutation, triggered by the periodic detection of problems – epistemic problems, that is, not operational ones. In other words, second-generation KM concerns itself mainly with the management of knowledge processing and its outcomes, not business processing. Its chief concern is the health and well being of the organizational capacity to engage in deep questioning, problem-solving, and error elimination. Thus, its primary goal is to achieve high-performance knowledge processing – sustainably so.

By contrast, in first-generation KM there is no conception of knowledge processing as something which is separate from business processing, and there is no consideration of things like epistemic problems, questioning (deep or otherwise), conjecture and refutation, the fallibility of knowledge, sustainability, or any such thing. There is only recognition of existing knowledge (e.g., ‘best practices’) followed closely by the need to capture, codify, and use it. Life, for first-generation thinkers, is just business processing punctuated every now and then with a need for information.

There is, however, one other thing of primal concern to practitioners of first-generation KM: *business strategy* and the need (no, the *imperative*) to support it. Conventional KM marches to the tune of business strategy almost without fail. Strategy’s fulfillment through the implementation of KM is of singular, paramount importance to first-generation KM, and therein lies one of its major weaknesses. Why is this a weakness? Because most business strategies are *shallow* in the Naessian sense of the term thanks to, (a) their anthropocentricity and, (b) the extent to which they discourage – even eschew – deep questioning. The former is the product of the latter.

Knowledge management strategies and interventions that subordinate themselves to business strategies in this way are, therefore, themselves shallow. Their sole purpose is to round up and make ‘sharable’ whatever information resources the strategy du jour might require. Never mind the general quality of knowledge processing in the firm, or the degree to which knowledge in use is, or is not, open to thoughtful criticism or refutation by stakeholders. And never mind how well, or not, current knowledge processing routines are actually serving the adaptive needs of the firm; or whether the interests of the communities in which the business operates are adequately represented in knowledge processing – be they human communities or otherwise. And never mind, as well, the *content* of strategy or what its impact on society or the environment might be.

No, the conventional practice of KM concerns itself with none of that – it is, by this account, extraordinarily shallow, irresponsibly so.

“Deep” Knowledge Management

“Deep” KM, on the other hand, is second-generation KM with a bias – an ecocentric one. Like second-generation KM in general, “Deep” KM concerns itself with the health and well-being of knowledge processing (i.e., the social and organizational capacity to learn and adapt), and so it strives to achieve and maintain conditions of openness, accordingly (i.e., the Open Enterprise).

But “Deep” KM places a premium not only on human performance in related systems (businesses, industries, and economies), but also on minimizing the impact of human activities on non-human environments, as well. To do this, “Deep” KM aims to create and maintain the conditions required for open learning and adaptation; and for producing outcomes, as well, that do not conflict with the “richness and diversity” of other non-human life forms and habitats. “Deep” KM is ecocentric, not anthropocentric.

What might this mean for us in practice?

First, “Deep” KM holds to the principle that strong organizational learning systems will not necessarily lead to sustainability and ecocentrically-stable outcomes, but that no such outcomes are possible without them. Here we can equate ‘strong organizational learning systems’ with ‘strong knowledge processing systems,’ and the *belief* – but not the certainty – that sustainable, ecocentric outcomes will *most likely* emerge in organizations whose knowledge processing environments are open. Indeed, hope springs eternal in “Deep” KM. Why so?

According to Deep Ecology, people are fundamentally endowed with an affinity for (and a love of life), an idea expressed in a corner of Deep Ecology called *ecopsychology*⁶ and also by Edward O. Wilson’s notion of *biophilia* (Wilson, 1986). But that is a subject for another day. Suffice it to say, for now, that the practice of “Deep” KM (like Deep Ecology) is predicated on a belief that sustainable outcomes are *more likely to emerge* in an Open Enterprise, and that KM’s mission, therefore, should be to create and maintain the conditions needed for openness in knowledge processing.

Next, neutral or *Shallow KM* tends to beget anthropocentric outcomes, not only because of its failure to promote deep questioning, but because of the criteria it supports (value-laden ones) in the evaluation of knowledge claims. Knowledge production rules that, themselves, are laced with anthropocentric overtones (or merely the absence of ecocentric ones) will, not surprisingly, give rise to anthropocentric ideas and outcomes more often than not. If, on the other hand, the validity of our knowledge claims are, at least in part, determined from an understanding of what their ecological implications might be (costs and benefits), we could fairly expect that more of our business processing behaviors (i.e., our *validated knowledge in use*) would be ecocentrically aligned.

“Deep” Knowledge Management

Organizations have the capacity to choose their *knowledge claim evaluation criteria*. They therefore have the capacity to change *the rules for making rules*, including the option of holding all of their knowledge claims to a standard of ecological integrity. This is the bias inherent in “Deep” KM.

Finally, “Deep” KM is inclusive of all stakeholders in an enterprise. By contrast, Shallow KM is, in a sense, *Orwellian* in that it reflects the view that some knowledge processors (people) are more equal than others (Orwell, 1946) (i.e., business strategies produced by senior managers are always more valid and should only be produced by them). “Deep” KM makes no such presuppositions. A stakeholder is a stakeholder, and more importantly, all knowledge is fallible and ought to be continuously subjected to open criticism and refutation – a tenet that applies even to Deep Ecology, itself.

But none of this need challenge or undermine the authority of conventional managers. The casebook of firms in which knowledge in use is freely debated, even as current-day managers continue to employ classical command-and-control styles of leadership, is growing (Bragdon and Karash, 2002). It is – read, *is* – possible to separate *business processing management* from *knowledge processing management* in a copacetic way, and second-generation KM (aka, ‘*the new KM*’) is trying to show us how (McElroy, 2000; Firestone and McElroy, 2003).

Sustainable Innovation

Once we see that our business processing behaviors are unsustainable, it’s a short hop from there to concluding that the knowledge processing systems that produce them are, too, unsustainable. Any approach to knowledge production or innovation that consistently gives rise to ideas hatched in the absence of ecological considerations does its dependents a grave disservice. Indeed, we do *ourselves* a disservice by continuing to rely on learning and innovation regimes that consistently produce such prodigious volumes of false, high-cost ideas, while failing to inform us of the damage done and the need for change along the way.

Nonetheless, this is a fair description of the learning and innovation process found in most corporations. Indeed, our prevailing knowledge processing systems are arguably dysfunctional, not only because of the (false) factual knowledge they turn out, but also because of the (false) ethical knowledge they create. Much of our problem of dysfunctionality could be solved by subjecting knowledge claims to legitimacy or integrity evaluations rather than just to factual tests (Firestone, 1974). Value claims are just as important as knowledge claims, if not more so.

Closed knowledge processing systems in which the majority of stakeholders in a firm are excluded from the process, and which routinely make it possible to validate knowledge that *devalues* consideration of environmental and ecological

“Deep” Knowledge Management

outcomes, are inherently unsustainable. Moreover, the management practices that help to create, fine tune, and hold them in place are unsustainable, as well.

We can say, then, that there are three levels of sustainability to consider (again, see figure 1). First is sustainability in operational or business processing behaviors; next is sustainability in epistemic or knowledge processing behaviors; and third is sustainability in meta-epistemic or knowledge management behaviors. Thus, any knowledge management practices that have the effect of holding unsustainable knowledge processing behaviors in place are themselves unsustainable; and any knowledge processing behaviors that have the effect of holding unsustainable business processing behaviors in place are also unsustainable. What we need is a prescription for sustainable innovation – a *knowledge processing model* that would actually *serve*, not *subvert*, our interests.

I believe that the best measure of sustainability in the conduct of human affairs is taken by gauging our impact on the ecology around us. This, of course, is not a new idea. But I also believe that the solution to our problems, ironically, has less to do with the need for change in our operational behaviors than it does with the need for change in our epistemic behaviors. Our operational behaviors must also change, of course, but the *deeper* change required, I believe, is in the manner in which we produce and integrate our knowledge. Why? Because until such time as we fix our dysfunctional knowledge processing routines, we should not expect their unwanted progeny (i.e., our worst ideas and the bad practices that follow) to appear any less often.

Here, as Popper would put it, our goal in knowledge processing should be to maximize our capacity to *eliminate errors*. My only contention, then, is that modern corporations do not perform this task well enough – *error reduction*, that is. Their epistemic behaviors, therefore, are not sustainable, nor are their downstream operational ones. Indeed, the dysfunction of the former accounts for the tragedy of the latter. These are my claims – no more, no less.

Conclusion

In examining the complex range of activities found in human social systems, I have made the crucial distinction between *business processing* behaviors (operational), *knowledge processing* behaviors (epistemic), and *knowledge management* behaviors (meta-epistemic). Most attempts to cure the unwanted social and ecological effects of business behaviors in society, I claim, have, in fact, been aimed at only one of these levels – the operational one. This is misguided. At best, constructive change on that front can only produce temporary or short-term benefits. Why? Because no amount of well-intended reform at the level of business processing behaviors can change the nature of the underlying knowledge processing system from which bad ideas and their downstream effects spring forth. Change the epistemic system *first*, I say, and

“Deep” Knowledge Management

more lasting and beneficial change at the level of operational behaviors will follow.

This is a politically charged issue, for sure – I know that. The kind of change I hope for will not be easy. It is, however, necessary. There are *limits to growth* in the conduct of human social and economic affairs (Meadows et al, 1972). The rational side of every one of us knows that. What we must do, then, is to make it possible for the rational voice (the ecocentric one) to have a seat at the table in business. To do this, we need to reinvent the formal learning process in organizations – or, rather, to *restore* it – in such a way that knowledge processing becomes more open, inclusive of all stakeholders, and utterly rooted in a respect for the fallibility of knowledge. Fix our knowledge processing systems first, I say, and change in our operational behaviors will surely follow.

This, then, is the central, driving motivation of “*Deep*” *Knowledge Management*, a contrarian view of what it will take to achieve sustainability in business and in the conduct of human affairs. It is a vision largely inspired by Deep Ecology, but which adds the notion of aiming our interventions not so much at our operational or business processing behaviors, but at our epistemic or knowledge processing ones, instead. There, and only there, will we find the leverage needed to rejigger the course of humanity, and to reverse and mitigate our impact on other species, their habitats, ourselves, and on Gaia itself – the *Deep Ecological ‘Self,’* that is.

* * * * *

Notes

1. Arne Naess first used the term ‘Deep Ecology’ in a speech he gave in 1972 at a Third World Futures conference in Bucharest.
2. The term ‘ecosophy’ was coined by Arne Naess.
3. This phrase is a derivative of what Karl Popper wrote in his book, *Objective Knowledge* (Oxford, England: Oxford University Press, 1972), p. 122: “Scientists try to eliminate their false theories, they try to let them die in their stead. The believer – whether animal or man – perishes with his false beliefs.”
4. This model was jointly developed by Mark W. McElroy and Joseph M. Firestone in connection with their work at the Knowledge Management Consortium International (KMCI). It is sometimes referred to as ‘The New Knowledge Management Reference Model.’
5. The idea of ‘justified true belief’ as a definition for knowledge was first formulated by Plato. It assumes that objective truths exist; that humans can believe in them; but that justification, or evidence, must also be present in order for true beliefs to be seen as knowledge. Its weaknesses, from a Popperian perspective, are that it uncritically accepts the validity of such evidence and/or suffers from Foundationalism, according to which some truths can be rationally justified as unequivocally true, upon which such

“Deep” Knowledge Management

certain truths all other claims may reliably be tested and determined to be either true or false.

6. See, for example, T. Roszak, M. E. Gomes, and A. D. Kanner [Editors], *Ecopsychology* (San Francisco, CA: Sierra Club Books, 1995).

References

- Argyris, C. “Teaching Smart People How To Learn,” *Harvard Business Review* (May-June, 1991): 99 – 109.
- Bodian, A. “Simple in Means, Rich in Ends: A Conversation with Arne Naess,” *Ten Directions* (Institute for Transcultural Studies, Zen Center of Los Angeles; Summer/Fall 1982).
- Bragdon, J. and Karash, R. “Living Asset Stewardship: How Organizational Learning Leads to Exceptional Market Returns,” *Reflections* (Volume 4, Number 1, Fall, 2002): 59 - 68.
- Carson, R. *Silent Spring* (Boston, MA: Houghton Mifflin, 1972).
- Firestone, J. M. *From the Philosophy of Fact-Value Dichotomy to the Philosophy of Organized Experience* (Binghamton, NY: unpublished, 1974).
- Firestone, J. M. and McElroy, M. W. *Key Issues in the New Knowledge Management* (Boston: Butterworth-Heinemann, forthcoming, 2003).
- Hartshorne, C. and Weiss, P. [Editors], *Collected Papers of Charles Saunders Peirce* (Cambridge, MA: Harvard University Press, 1931 – 1958).
- McElroy, M. W. “The Second Generation of Knowledge Management,” *Knowledge Management* (October, 1999).
- McElroy, M. W. “The New Knowledge Management,” *Knowledge & Innovation: Journal of the KMCI* (Vol. 1 - No. 1, October, 2000): 43 – 67. An on-line journal found at www.kmci.org.
- McElroy, M. W. *The New Knowledge Management – Complexity, Learning, and Sustainable Innovation* (Boston: Butterworth-Heinemann, 2003).
- Meadows, D. H., Meadows, D. L., Randers, J. and Behrens, W. W. III, *The Limits to Growth* (Washington, D.C.: Potomac Associates, 1972).
- Naess, A. “The Shallow and the Deep, Long-Range Ecology Movements – A Summary,” *Inquiry* (Vol. 16, 1973).
- Naess, A. “The Deep Ecological Movement – Some Philosophical Aspects,” *Philosophical Inquiry* (Vol. 8, nos. 1-2, 1986).
- Orwell, G. *Animal Farm* (New York: Harcourt Brace & Co., 1946).
- Popper, K. R. *The Open Society and Its Enemies* (London, England: Reprinted by Routledge, 1998).
- Sessions, G. [Editor]. *Deep Ecology For The 21st Century* (Boston, MA: Shambala, 1995).
- Taylor, F. W. *Scientific Management* (New York: Harper & Brothers Publishers, 1947).
- Wilson, E. O. *Biophilia* (Boston, MA: Harvard University Press, 1986).

“Deep” Knowledge Management

About The Author

Mark W. McElroy is the Director and Chief Sustainability Officer of the Center for Sustainable Innovation (www.sustainableinnovation.org) He can be reached by e-mail at mmcelroy@vermontel.net.

11/02