

Pre-publication Excerpt from:
Riskonomics: Reducing Risk BY Killing Your Worst Ideas
By
Joseph M. Firestone, Ph.D.
Managing Director, CEO
Knowledge Management Consortium International

Chapter 2 What Knowledge Is

Many who have attended training seminars on Knowledge Management report that their instructors seem to avoid defining knowledge. But that they almost always distinguish between "tacit" and "explicit knowledge" and among "data", "information", "knowledge", and "wisdom". If we're going to be talking about improving the quality of our knowledge, won't we need these distinctions here?

It's not always essential that people define their terms. Definitions are really often just "elevator speeches" anyway, and not attempts to express the "essence" of an idea. Nevertheless, I do think that before we talk about how to improve something we ought to have a reasonably clear idea of what it is we're trying to improve. So, for the next few pages, at least, I want to talk about what "knowledge" is. I'll begin by saying what I mean by "information."

Information is a non-random structure or pattern of relationships within a system, indicating future interactive potentialities. It either originates along with the system, or is acquired or developed by it in the course of its interacting with and responding to its environment, and the problems generated by that interaction. Note that this definition does not require correspondence between information and the environment. Nor does it assert that information is *encoded* in some simple cause-and-effect fashion, but leaves room for emergent information in the context of a system's interaction with the environment.

"Emergent Information" is information *created by the system* that cannot be entirely accounted for as just an effect of the causal interaction of the system's components with the environment. For example, you create beliefs (information in your mind) that, to the best of current knowledge, can't be accounted for by the causal interaction of your brain with your environment. Getting closer to home, the fine Bouillabaisse you may have enjoyed at your favorite French Restaurant was produced from information in the chef's mind that, while influenced by her brain functioning, memory of previous Bouillabaisses, and the recorded recipe she probably consulted at some point in the past, was created by her in the process of making this dish.

Sidebar 2.1: Emergence

Emergence refers to the processes through which autonomous entities self-organize to produce wholes (organizations of process) whose collective properties are both novel, and unpredictable, relative to an analysis of the properties of the self-organizing entities and their interactions. Closely associated with the idea of emergence is the idea of 'downward causation,' that the collective properties of the emergent system, after self-organization, then constrain the behavior of the autonomous entities whose self-organizing interaction created it.

There is a lot of controversy about the idea of emergence. It is an area of exploding research that is interwoven with research on complexity, complex adaptive systems, multi-level hierarchical complexity, and self-organization. When the focus is on emergence itself, a central issue is opposition of the idea of emergence to the position of reductionism in both natural and social science.

Both reductionism and emergentism are related to the idea of multiple hierarchical levels of interaction or reality in systems. The reductionist position asserts that higher levels are causally dependent on lower levels and that, in principle, the emergence of higher system levels, and the processes and properties of these levels are caused

by and can be scientifically explained in terms of the processes, properties, and dynamics obtaining at lower levels. The reductionist position has often characterized the development of Science and is frequently credited with many of its greatest successes. However, in the very many scientific domains where reductionism has been unsuccessful, or is far from even serious pursuit, claims that higher level dynamics can be reduced to interactions and relationships among lower level entities seem little more than the presentation of promissory notes, or even worse, in certain areas (such as the mental, where previous attempts at reduction have failed miserably), the statement of some ancient dogma. The truth is that both reductionism and emergentism are metaphysical research programs that may guide scientific research. But either assumption may prove false in particular problem domains, and neither should be favored in testing competing theories.

I've started talking about systems because that is the relevant context for the idea of information. Information is created by systems, and it is sent (communicated) from system to system in the form of content that may or not be "understood" by the receiving system.

The most important aspect of information, however, is not whether it is complex or simple, or

produced quickly or slowly, or gained or lost over time, or whether there is a great or a small amount of it. All of these are important, but *the most important aspect of information is whether its use in decisions enhances the ability of the system using it to adapt.*

And this ability to adapt, in turn, is most likely to be enhanced if the information itself actually "corresponds" to the reality of the system's environment. Evolution provides such "correspondence" by selecting for those life forms that fit the environmental constraints in which they live. Errors in genetic information are eliminated over time by the environment, along with the organisms that contain them. Learning provides such "correspondence" on a much shorter time scale by providing us with an opportunity to eliminate our errors in information and to create new information that survives our evaluative efforts and our experience.

Sidebar 2.2: Shannon Information

You may have noticed that the ideas of 'information' as a non-random structure, and as 'emergent,' don't fit well with the account of information that, by far, is the most successful, namely Claude Shannon's account. But, it is important to understand that Shannon's model is, and from the beginning was, intended to be an

account of information that focused on only one aspect of the pre-existing notion of 'information' common in ordinary language, specifically the gain in information resulting from a communication of symbols or bits.

Shannon emphasizes that messages have 'meaning', and "that these semantic aspects of communication are irrelevant to the engineering problem." In other words, he is not concerned with, ***nor does his theory treat***, either the structure (pattern of relationships, form, organization, etc.) that is information or its content, including whether a semantic network expressing information is true or false. What he does treat rather, is the problem of designing a communications system that will "operate for each possible selection" of a message.

Thus, Shannon's theory is neither a comprehensive theory of information, nor a theory that measures "truthlikeness". At best, it is a theory about communications that makes use of a measurement model focusing on one aspect of information -- ***how much of it is gained as a consequence of selection of a possible message or event, other things being equal.*** This measure is specified relative to a "before and after" context of change in amount of information gained, and relative to some set of messages and the symbols within them. It does not speak to the question of amount of

information contained in a message based on the internal structure of the message alone, the problems it addresses, and the semantic content of the message.

So it is not a measure of the total amount of information present in a message, but rather a relative measure of the reduction of uncertainty (and gain in information) resulting from selecting the message. Its application requires the specification of selection probabilities of all documents in a finite and supposedly complete comparison set, conditions which cannot be realistically fulfilled in many situations. Also, it doesn't speak to other contexts for measuring information that we may construct.

Efforts to transcend Shannon information are increasingly common. These focus on models and measures concerned with structure and on alternative measures of the basic unit of information, providing alternatives to the 'bit' identified by Shannon and named, famously, by John Tukey. Some of the new efforts apply the notion of statistical complexity. Others are based on R. A. Fisher's earlier measure of information, which focuses on the idea of 'likelihood' rather than probability. Still others develop entirely new alternative measures and theories, enabling easier recognition and specification of the very broad

range of alternatives to Shannon information that may prove useful in new applications.

It is no exaggeration to say that Information Theory is exploding today, as new conceptions broaden its foundations and attempt to treat the problems of structure, pattern, content, and meaning. But it is also true that this explosion has not yet reduced the interpretative, semantic aspect of information to a formal or physical model. Information, in the broad sense including 'meaning', still carries with it a strong element of cultural interpretation. 'Meaning,' in turn, is related to patterns of language usage in speech communities: patterns that in open societies evolve constantly as new 'hypotheses' and conventions about meaning, embedded in new patterns of language use, are proposed, and accepted, by members of those communities.

Thus, the analysis and measurement of information in the broad sense still remains unconquered territory. And mastery of this territory seems related to mastery of the study of self-organization, complex adaptive systems, and emergence: an open-ended and formidable task that will occupy our best minds for many years to come.

Since the most important aspect of information is "correspondence" with reality, the most important

measures of information networks are those that evaluate this correspondence. Thus, *the most important measures we can develop describing information networks are measures that help us to evaluate those networks*, and that brings us to "knowledge." I favor a "unified theory" that states a general viewpoint about knowledge, but which also distinguishes different types of knowledge. Here it is.

Knowledge is a tested, evaluated and surviving structure of information (e.g., DNA instructions, synaptic structures, beliefs, or claims) that is developed by a living system to help itself solve problems and which may help it to adapt.

This is my general viewpoint. It is consistent with my definition of information. And it is consistent with the view that knowledge is something produced by systems in order to help them adapt to environmental challenges.

This definition is a highly generalized view of knowledge. Why is it useful to us in understanding knowledge as it relates to decision models?

The most useful view of knowledge, because it is the most continuous with modern Biology, is the one that identifies it as a general phenomenon characteristic of all living systems from "the

amoeba through Einstein", through organizations, to the international system. If we accept this view we can distinguish three types of knowledge that fit it:

- tested, evaluated, and surviving structures of information in physical systems that may allow them to adapt to their environment (e.g., *genetic* and *synaptic* knowledge composed of biological structures used in developmental and learning processes);
- tested, evaluated, and surviving beliefs and belief predispositions (in minds) about the world (subjective, or non-sharable, *mental* knowledge composed of mental structures used in learning, thinking, and acting); and
- tested, evaluated, and surviving, sharable (objective), linguistic formulations about the world (i.e., claims and meta-claims that are speech- or artifact-based or *cultural* knowledge used in learning, thinking, and acting).

Sidebar 2.3: On Objectivity

'Objectivity' in both description and evaluation can refer, I think, to three aspects of our knowledge.

First, in claiming that our knowledge is objective we may mean that our knowledge claims correspond to whatever they are supposed to correspond to. If the claims are factual, we claim

correspondence to the facts, to what exists, or to what may exist, or to what fails to exist. If the claims are evaluations, we claim correspondence to values, to what should or should not exist, whether the existence reference in such statements is past, present or future. Since we can never be certain that we have attained such correspondence either in the realm of fact or value, this sense of objectivity represents only a regulative ideal. It is something we seek, but can never be sure we have attained. ***And we think that this is all there is to the matter of correspondence and to objectivity in the sense of correspondence.***

Second, 'objectivity' may refer to knowledge that is the product of a particular method, the method of scientific investigation. Objective knowledge then, would refer to knowledge claim networks that have survived testing and evaluation, and the record of their performance emerging from scientific investigation.

Third, ***'objectivity' may refer, simply, to knowledge claims that are sharable and criticizable, and objective knowledge to the combination of sharable and criticizable knowledge claims that have survived our tests, along with the performance record of these claims.*** This sense of objectivity generalizes the second sense of this term. It recognizes that scientific investigation is just a more rigorous

variety of the method of trial and error elimination, or conjecture and refutation, the method we apply to sharable and criticizable knowledge claims in order to grow our knowledge.

"Belief predispositions" are *propensities* to hold beliefs fitting into general classes like values and attitudes. Most of our mental knowledge is predispositional. Our conscious beliefs are only "the tip of the iceberg" existing in a specific situational context and changing with changes in our context. "Claims" are statements we make about the world, or about what's valuable, or what's right, or what's valid or invalid. They can be made in single statements or in networks of statements. "Meta-claims" are claims about claims. They're important because the record of survival of our claims, which distinguishes knowledge from information, is made up of meta-claims.

The original tacit/explicit knowledge distinction was between beliefs that can't be expressed and those that can. So, 'explicit knowledge,' defined this way, wasn't even codified or expressed knowledge. If you use this version of the distinction, you probably don't view cultural knowledge as "knowledge", but only as "information". If that's the case, you still have the problem of distinguishing information that

survives our testing and evaluation from information that doesn't.

On the other hand, many in Knowledge Management think that tacit knowledge is mental knowledge and explicit knowledge is codified knowledge. And if that's the way you put it, then the distinction is the same as the one I've made above, except, of course, as stated in just this way, i.e. tacit vs. explicit, the distinction doesn't make clear the adaptive function of knowledge, whether it is mental or cultural. ***This adaptive aspect is critical in distinguishing "knowledge" from "just information."***

Other important distinctions used in clarifying the nature of knowledge focus on the distinctions among data, information, knowledge, and wisdom. According to the most popular view:

- data is composed of sequences of numbers or letters without context;
- information is data in context;
- knowledge is belief built on data and information in an actionable context; and
- wisdom is knowledge that transcends mere description and makes value choices.

This is the so-called pyramid view relating data, information, knowledge, and wisdom and viewing them as progressive refinements beginning with data. Data is considered to be most plentiful and

the foundation for everything else. Information is somewhat less available, knowledge is much more rarely found, and rarest of all is wisdom.




Figure 2.1: The Pyramid View

I'll start my discussion of these distinctions by asking a rhetorical question. Have you ever seen any data that looked like sequences of letters or numbers that have no context? I haven't. I've worked with data from chemistry, physics, biology, psychology, and gaming experiments. I've also worked with statistical data from attitude and opinion surveys, psychological studies, sociological conflict studies, cross-national conflict, international conflict, value and motivational studies, and I may have forgotten one or two other categories of data I've worked with.

I've never seen any data that conforms to the definition given in the pyramid view asserting that it is without context. In fact, I don't think it exists. I think, instead, that *all data has context and is cultural information*, and that *some of that information*, if it has survived our tests and evaluations, *is cultural knowledge*. So I don't believe in the distinction between data and information just given, or in the process of refinement underlying the pyramid "image".

APPENDIX A
DOMESTIC CONFLICT - RAW DATA MATRIX

Nations	Variables								
	Assassinations	General Strikes	Guerrilla Wars	Major Govt. Crises	Purges	Riots	Demonstrations	Revolutions	No. Killed in Domestic Violence
Afg 55									
Afg 56									
Afg 57									
Afg 58						1	1	1	
Afg 62									
Afg 63									
Afg 64									
Afg 65									
Afg 66							1		
Alb 55									
Alb 56						1			5
Alb 57									
Alb 58									
Alb 62									
Alb 63									
Alb 64									
Alb 65									
Alb 66									
Arg 55	3	0	1	2	3	17	6	3	4279
Arg 56		2	1	1	3	3	2	1	91
Arg 57		2	1	1	2	5			
Arg 58		4	1	1	1	5			
Arg 62		1	1	6	1	1	1	1	9
Arg 63		1	1	1	1	3			36
Arg 64		1	1	1	1	10	7		3
Arg 65		1	2	2	1	10	7		14
Arg 66		2	1	1	1	9	2	1	1
Aus 55									
Aus 56									
Aus 57									
Aus 58									
Aus 62									
Aus 63									
Aus 64									

Figure 2.2: Cross-national Conflict Data: Structured Information Including Context

Instead, I prefer a distinction between two kinds of information: *structured* and *unstructured*. Data is just structured information. It is cultural in nature,

and is only meaningful in the context of theoretical frameworks and preconceptions. Also, I don't believe that data, in the sense of structured information, is "mental" in character. There is no mental information "structured" in the way that data is structured in cultural products.

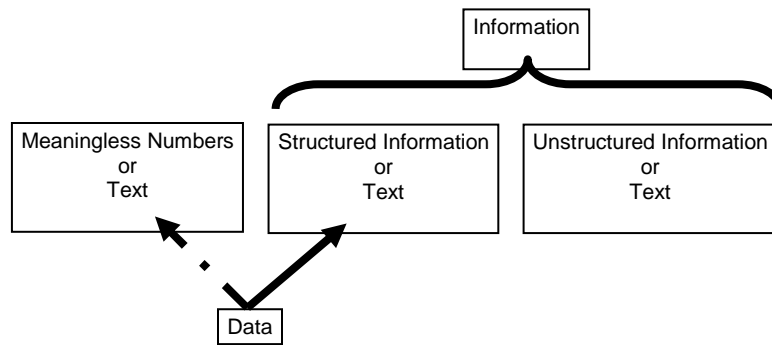


Figure 2.3: Data and Information: The Choice

Nor do I think that the distinction between information that is external and knowledge that is internal and mental is very helpful, either. First, I've already said that there are three types of knowledge and implied that there are three types of information: biological, mental, and cultural, in contrast to the popular idea that there is only belief or mental knowledge. Second, I certainly don't think the distinction between information and knowledge is just that knowledge has a more extensive and actionable context than information. For me, it is not the general idea of more extensive context that is the distinguishing factor. Rather, *the*

context that is specifically relevant is the meta-claim record showing that some information has survived our tests and evaluations and is therefore knowledge, while other information has not survived and is therefore only information.

The distinction I like the best from the pyramid model is the one between knowledge and wisdom. I think that's basically correct if interpreted from the point of view of the other distinctions I've made. *Wisdom is about having quality knowledge in the area of decision models.*

That is, to have wisdom is to have a decision model with as low a risk of error as possible since such a model will take truthlikeness, probabilities, and values all into account. So, I think that when we are learning how to make better quality decision models, and to reduce our risk, we are also learning how to do a better job of producing wisdom.

In considering the views I've just stated, many will find my three-way classification of knowledge a little odd. When they think of knowledge, they always think there must be a "knower", a mind that "knows" the "knowledge". But in the case of the biological and cultural categories of knowledge there are no "knowers", (though 'knowers' make cultural knowledge, and they also have biological knowledge they 'know' in a mindless way) because

no minds are involved. So why call these categories "knowledge"? Why not simply refer to them as adaptive information that has survived our tests and leave the term "knowledge" to the category of mental belief?

This view of things recognizes that all three categories of information I've identified, share with one another the quality of having survived environmental challenges, and also that there are three separate categories at issue here: biological, mental, and cultural. So, apparently, there's nothing in it that disagrees with my analysis, except that it reserves the term "knowledge" for mental belief because "knowledge implies a "knower"". While I think this view is certainly reasonable, I also think there are a number of arguments against it.

First, in common parlance, we don't always use "knowledge" in a way that implies that we always have a knowing subject or mind directly associated with it. I'm sure you've heard of the idea of "body knowledge", where the body, but not the mind, or even the brain, is presumed to have a predisposition of some kind. In addition, we use the word "knowledge" all the time to refer to "bodies of knowledge", "scientific disciplinary knowledge", "books of knowledge", etc., where we are clearly not referring to beliefs in people's

minds, or to physical predispositions, but to *recorded* adaptive information of some kind.

Second, restricting our use of "knowledge" to mental belief precludes any reference to "objective knowledge", since we can't share our mental beliefs directly with others or subject them to shared criticism. Cultural knowledge can be objective in this sense, so one of the consequences of an agreement to use the term "knowledge" as I proposed, is that we can continue to use the term "objective knowledge", rather than having to refer to "objective adaptive cultural information", which would be our alternative if we accepted the proposal to restrict knowledge only to the mental category.

Third, and perhaps most importantly, restricting the word "knowledge" to mental belief doesn't change the practical situation we have to face at all. We still have to distinguish between information that survives our tests and information that does not. We still have to reduce the risk of error in our decision models and we still have to learn how to do this through a better understanding of how we make cultural information that survives.

So my bottom line is: I don't think restricting the use of "knowledge" to mental belief does anything to solve any practical problems that we face. Rather, I think it carries *the danger* that we will

worry about *how we ought to justify our beliefs*, rather than about how we can formulate higher quality decision models that are closer to the truth and carry less risk of error, and therefore can also provide better support for our mental decision models.

What do the three types of knowledge have to do with improving the quality of our decision models and reducing our risks? The first type, biological knowledge, may seem only indirectly relevant to improving decision models in organizations, but improving our decision models is, most immediately, all about mental knowledge and cultural knowledge and their continuous interaction. When we make new knowledge, as we do when we are improving our decision models, we rely on both our previous mental and cultural knowledge. And as we produce new mental knowledge, it both directly contributes to our decisions, and influences any new decision models we are developing, which, as they are being developed, influence any further thinking we may do.

Sidebar 2.4: Is Biological Knowledge Relevant to Organizations?

Biological knowledge may not be directly relevant to the context of making better decision models, but improving an organization's biological

knowledge through better recruiting can provide enhanced individual capability to solve problems, and, in this way, help it make higher quality mental and linguistic models along the way. This isn't just a question of inherited genetic knowledge. Biological knowledge is also stored in synaptic structures and memory. That knowledge includes predispositions to recall content, and procedural knowledge about how to do things, including better generalized knowledge about to solve problems and learn.

In other words, if you recruit more intelligent, creative problem solvers, who know how to collaborate with others to both develop new ideas and test them severely, you'll improve your organization's problem solving capability, adaptiveness, and performance.

Further, when we talk about our decision models, we must recognize that our mental decision models can only be very partial, because we cannot hold the details of complex decision models and various reasonable alternatives in our minds at once. On the other hand, our cultural decision models can be as complex as we need them to be to support our mental models and our decisions.

So, I think the aim of reducing risk by increasing the quality of our mental decision models needs

to be pursued by directly strengthening our cultural decision models and by relying on the continuous interaction between our cultural and mental models to strengthen those mental decision models. This brings us to my next set of topics. These concern the process of making new and better decision models.

