

Chapter 9

The Power of Homologies

Situations represent the most complete way of understanding our experience of the surrounding world and the human qualities of the world. They also endow experience with durability in relation to which other experiences can acquire meaning... What we normally refer to as reality, believing that it is something fixed and absolute, is always a result of our ability to experience, visualize, and articulate—in other words, to represent so as to participate in the world. There is a point where the interpretation and the way of making come so close to each other that they become fully reciprocal: what we know contributes to what we make, and what is already made contributes substantially to what it is possible to know (Vesely, 2004).

In re combinatorial environments, meaning is characterized by a fluid, shifting, continuous state of becoming. In this form of fleeting context, content is always emergent, arising out of the superimposition and or juxtaposition of a series of "poetic" elements and processes functioning in relation to one another.... Fleeting and shifting qualities of engagement become an experiential focus. During interaction, the user, through direct experience, encounters a series of potential "states" of meaning. We should always view these states as a temporary glimpse at a continuous process of meaning-becoming, motivating the thought and behavioral reaction of the user (Seaman, 1999).

You've got the looks, I've got the brains ... let's make lots of money... You've got the brawn, I've got the brains ... let's make lots of money" (The Pet Shop Boys "Opportunity")

Coherence is not a principle of order and equilibrium, but a principle of interaction and self-creation. Complexification, development and renewal, may provide emergent energy; but their instability can also endanger organizational existence. The becoming of self and the other may lead to constant creation, activity and participation; but will the result have enough structure and order to sustain itself as organization?

Self (and organization) does need a minimum of psychic stability to exist. Too much process leads to a disintegrating, psychotic self that is unable to accept any order. Structure is not the antithesis to process, identity and self. Structure is a prerequisite to emergence, to organization and to coherence. The dialectic has to be reversed, instead of being wary of organization and embracing self and the other; it is necessary to acknowledge that self and other define a (desirable) form of structure. In chaos, there is no self and other. When self and other participate together in activity, there is organizing and organization. Structure and participation are inherent to self and other being able to do things together.

Coherence can be a force for good or evil, creativity or repression, order or chaos. Coherence is order emerging from disorder, purpose coming out of chaos, and statements imposed on noise. Organizational coherence, and accepted norms or values, are closely linked. Coherence is never only a mission statement or statement of intention. Coherence, in effect, is co-evolving, jointly experienced purposefulness. Co-motivation or mutuality of purpose is anything but self-evident. Many organizations are led by collective fantasy, shared capriciousness and group self-delusion. In such organization many shout 'coherence' but there is little joint awareness, action or purpose.

Illusion(s) of coherence can be (almost) as powerful, as persons, events and actions really coming together. Emergent coherence can be false --- in effect, a form of collective insanity. Coherence can be the choice to live a (common) myth, celebrate (collective) delusion and be enclosed in dreams. Coherence, cohering around delusional fantasies, can be supported by very strong networks of affordances, events and meaning, under which lie strong homologies. The desires, fears, hopes and needs that motivate false coherence, can be very strong and intense. People will (sometimes) fight to maintain their illusions. The group believes - or at least its members all (often) attest to believing --- that what they are doing, who they are, and what they are becoming, is uniquely and forcefully coherent. Claims of unicity and of possessing an inimitable will to succeed, are common.

Homologies are shared, perceived underlying models. That the models are perceived as being a common underlay is the only criterion. Truth claims, potency, vision all may be asserted by those who share the underlying perception that the model is in common, but the existence of a homology is not in itself a proclamation of truth. The financial crisis of 2008-2009 is perhaps the most recent and strongest testament of that in-congruence. While nearly everyone involved in the housing bubble shared an underlying homology that housing prices only go up, the power of that seemingly self fulfilling prophecy failed us when the summer storms of 2008 hit. This pattern holds in most "bubbles" where the participants all seem to share an underlying homology which grants them access to new affordances and to the prospect of greater riches. For a while, the commonality of the homology and the willingness of participants to act upon the affordances which that homology makes available creates a strong upward bias. (See the song lyrics at the beginning of this book.) The resonance amongst the participants is

great. The claims of the power of the resulting coherence (synergy, alignment, new way of working etc.) are huge. Self delusion becomes a propelling force. And then, the music stops.

While we were doing our original research, one of the organizations that presented at the Managing the Complex workshops displayed all these aberrations. Its spokespersons claimed a powerful coherence. They were so convincing that several attendees actively sought to invest in the venture. After the fact – see the data below --- the coherence appears entirely delusional. How do we handle real values and false promises? The homologies can be powerful enough, but the ability of the group to pull together and to deliver what it promises, can be close to nonexistent. What is experienced as coherence, can be a collective delusion. An organization can produce coherence in seriality --- that is, each person apart claims coherence, and together they share next to nothing. The individual experience of ‘coherence’ is real enough; the collective experience is close to nonexistent. Myths of collectivity can exist in individual consciousness without there being sharing, without bonds of commonality and without collective action. Ideas of coherence can be very powerful, without the social activity of coherence ever taking place.

As you read this technology based case study, let your mind substitute the arena of housing and the sophisticated product known as “sub-prime option rate mortgages.” The parallels are there to be observed. They too are homologies.

WebMind the Simulacra

The organization examined in this case study had two different names, ‘WebMind’ which was the product to be, and ‘Intelligenesis’ which was the name of the corporation. Webmind ‘burned’ more than twenty million dollars before it was closed down. The myth was supplied by its charismatic founding scientist, Ben Goertzel. In essence he promised to change how artificial intelligence was done. He seemed doggedly determination to lead an AI. (Artificial Intelligence) revolution. WebMind was created due to that determination and WebMind died due to that determination. Along the way, the myth of WebMind (or later Intelligenesis) the company, and WebMind the technology, was very powerful. Both were told as participatory schema, but how realistic, authentic or genuine what they pointed to was, is debatable.

At WebMind, different AI projects competed with one another with internal champions who were barely on speaking terms. The corporate direction was yet another faction. There was a common and very powerful affordance --- everyone professed the conviction that a next generation of genuinely intelligent AI was in the making. The new generation of AI (of which WebMind was to be a part) would not only absorb information – however intelligently. It would be the first generation of proactive AI --- that is, a genuine ‘thinking machine’. The next step in AI would be to develop software that understood what it was doing and could act based on that understanding. Software with such a meta-level

intelligence could direct itself to what was important and could design its own search agenda. Such was the underlying homology (much like the “tools for life” homology at Tripod).

WebMind started in late 1996, when Ben Goertzel decided to start a software company commercializing his design for a revolutionary artificial intelligence (AI) system. Soon after getting started, Goertzel was contacted by Lisa Pazer, an economist and market strategist with a decade of experience on Wall Street, who was researching a novel on complexity. Pazer had stumbled upon an article Goertzel had written and posted on his website on a “worldwide brain.” The two realized they had a common fascination with the enormous potential of Goertzel's revolutionary approach to artificial intelligence. According to Pazer:

I got really excited, and thought it would make a great movie. When he told me about his software, I saw these great application ideas, and it seemed easier to start a company than write a good screenplay.

Together with several other founders, Pazer and Goertzel incorporated Intelligenesis in August 1997 to commercialize the “WebMind” artificial intelligence software. The two had complementary skills, as one of the first programmers hired commented: “Lisa was the brawn, Ben was the brains.”

Intelligenesis was conceived and marketed in the initial business plan as a product company that would introduce powerful new and important capabilities into the software market. Ben was confident that he could solve the age-old problem of natural language understanding in the first summer of operation. He was, in fact, convinced he *had* solved it mathematically, and that the simple part was in the coding. When it became clear that the natural language “revolution” was nowhere near attainable despite the months the development team labored at it, the “brawn” railed against the “brains” to produce *any* working software – even if it meant settling for conventional means of producing useful incremental AI-based products that didn’t qualify as “true” machine intelligence.

One very public example of this tension within the company: The far-flung Intelligenesis staff communicated extensively (almost exclusively in the beginning) via e-mail. The “tech list” was the ongoing thread of technology discussion throughout the company. For several hours one day, Lisa watched an ongoing philosophical discussion regarding AI *morality*. Would Webmind be “good” or “evil”? It was an invigorating, even fascinating discussion involving much of the technical staff, and this type of interaction is, within reason, of course, generally positive for team unity and emotional involvement in the product and company. However, clever and insightful commentary became as near a metric as any that existed for measuring accomplishment and productivity amongst the development staff since contributions in this manner were encouraged by the attention and plaudits from Ben and other technical management, and Ben saw no need to instill a corporate culture based on real

achievement towards the company's product goals. Finally, Lisa shut the discussion down with a terse statement: *What do you say we actually produce intelligent software and then worry about whether it's moral or not?* In short, "the brains" fantasized about changing the world, and "the brawn" had to meet payroll.

The stated long-term mission of Intelligenesis was to build the world's first thinking machine, an elusive goal long contemplated by many scientists, but considered by others impossible to attain. Goertzel, however, had developed a radically different approach to artificial intelligence. By creating an evolving series of "nodes" and "links" representing concepts, the Java-based WebMind was designed to emulate the processes of the human mind. The claim was that WebMind's natural language capabilities allowed it to understand what information meant in the dual context of user and environment. If true, WebMind could then tap into the estimated 90% of the world's stored business information residing in text of some form. Information that previously could only be analyzed as fast as decision-makers could read it. WebMind would provide a fast and effective way for users to extract higher quality information from lower quality information. Goertzel's description:

The WebMind AI Engine project is, in many important ways, the most ambitious initiative in the history of the AI discipline. Unlike most researchers and engineers working in the AI field, we in the AI Development Division of WebMind Inc. are actually making a serious attempt to create a truly intelligent computer program, in the short term. We have detailed software designs, and detailed engineering and testing and teaching plans, and a highly competent team of roughly 45 scientists and software engineers and testers executing these plans. We've been at it since mid-1998, and we estimate that within 1-3 years from the time I'm writing this (March 2001), we will complete the creation of a program that can hold highly intelligent (though not necessarily fully human-like) English conversations, talking to us about its own creative discoveries and ideas regarding the digital data that is its world.

The enthusiasm of WebMind employees regarding the promise of artificial intelligence cannot be overstated. The programming staff which Goertzel recruited was first rate and highly devoted to their tasks. They truly "believed" in the potential of AI and that they were changing the world. Goertzel had charisma. When he talked about AI or his theories of AI, both novice and expert fell under his spell. Such personalities are rare in the near genius populated world of AI computer programming, making Goertzel's charm even more persuasive.

WebMind was not merely a research project; it was a stand alone start-up firm. As such it had to generate revenue and build up a business. Goertzel's AI theories did not deal with the business aspect. The AI research side of WebMind and the "business" side were little integrated. The claim to have a

unique networked solution to the notorious shortfalls of AI created all the excitement. Many software designers and AI specialists were very eager to be involved in such an effort. The homology of networked intelligent computing and its many potential uses was enormously powerful. The pull of the 'business' was far less attractive. In the eyes of the developers, the 'business' existed to support the research, not the other way around. 'WebMind' referred in the minds of the researchers to a technology, not to a company. Because that technology was 'revolutionary' its commercialization would eventually prove very profitable. How or why the new technology would produce profits was always "Someone else's concern."

Alongside the development of the core WebMind architecture, a secondary project focused on smart search applications of WebMind. A demonstration version of the smart search function was developed. Tests of this early version of "smart search" met with mixed results -- it worked either brilliantly or not at all. These difficulties, together with the fact that the market at the time included five major search engines -- which pundits were predicting would consolidate should have discouraged WebMind's business side from further actively pursuing sales of the search application. However, Intelligenesis' pursuit of search and its associated technology infrastructure dominated the bulk of the company's technical and research resources throughout the duration of the company.

The company's veneer of market-driven AI development was as good as abandoned by technical management and the new CEO, Andy Siciliano, in late 1999. By then, the search market was well saturated, and the only thing companies in the space searched for with little success was a revenue model. The key problem was that "search" was a "soft" peripheral function within corporations, not mission critical by any means, and demanded the adoption of new processes and ways of working. Moreover, the value drivers in that space had little to do with the only differentiator ("intelligence") that the Webmind technology purported to offer.

As stated by Andy Siciliano and Ben Goertzel, the reason Intelligenesis remained in pursuit of search was to provide a tangible product focus for technology development, period. Certainly, the "technology constituency" still paid lip service to the "market need for intelligent search", but they *blatantly* ignored all market evidence to the contrary - even when it came not from within the firm, but from potential integration partners. For example, in one of the too few meetings held between Webmind and potential search technology integration partners, a knowledge management application developer listed the key drivers in the search market. Scalability, speed, and the ability to search many different file formats topped the list. "Intelligence" wasn't mentioned. This assessment of search market value drivers was exactly in line with the market intelligence provided to Ben Goertzel and Andy Siciliano time and time again. When the potential partners were asked "What about intelligence?" and the response was "oh, yeah, sure, intelligence would be good", the "market-be-damned technology constituency" actually walked away feeling validated.

The applications that WebMind now turned to seem to share a common fatal flaw --- they were supposed to do things that cannot be done in the way proposed. For example, WebMind made a public relations splash with announcements about making a stock market prediction application of the WebMind software. The business management thought they had a 'winner' --- with a market prediction tool they could demonstrate WebMind's value. With most other applications, the distance between bottom line profit and technology is hard to make, market prediction makes the link as close as possible. The initial objectives for the market prediction application were ambitious enough, as the senior business manager put it: "we want a piece of all the economic action on the globe." Their artificial brain would pick better investments than Wall Street's finest. WebMind's AI would scan thousands of newswires and market reports to pick out companies that are likely to do well. The business model called for the system to only be made available to Wall Street traders. Tests of a prototype supposedly showed a success rate of 90% in predicting the Dow Jones index while Wall Street's best traders had a success rate of 63%. The claim was that WebMind would answer questions like: "Which stocks will outperform the market in the next six months?" or "What company looks the best bet for tripling the size of my bankroll?" With a variant called "newscruncher" WebMind promised to scan the Internet and help owners decide which shares to buy and sell based on sentiments in chatrooms and message boards. Newscruncher was distinct from the 'market predictor' because an initial attempt to put them together had led to incoherent results

WebMind's spokesperson reported that the system had been tested for several months, using an American news service that holds two years of financial data, and that it had performed well: "We are already ahead of the human traders, but we should be. There is no way a human trader can process all the data we look at." This was marketing hyperbole, indeed. Ben Goertzel was adamant about Intelligenesis not becoming a financial forecasting firm, so the vision for the technology had to incorporate Ben's promise of true natural language and the Webmind architecture as a whole. The market prediction technology, despite marketing verbiage, never incorporated any of the WebMind architecture, or to the extent it did, it didn't work and the system was rebuilt as a standalone without any of the Webmind core technology.

WebMind's user interface promised a real revolution in computing. Users of the Webmind system were meant to be able to ask it questions in normal English. Natural-language processing developed by WebMind/Intelligenesis allows the computer to work out what the user wants. WebMind/Intelligenesis's claims made by senior management:

You can be as specific as you want, asking 'What shares should I buy today?' As the system can apply its intelligence to any problem that has large amounts of data, it will answer almost any question. The system can give trading houses a big advantage. A commercial system should be ready in a few months. It may be used later for other

applications, including news filtering, and even as a Net search engine. Webmind will work with whatever data you give it and will learn about what you want from it as you use the system. The possibilities are endless, but obviously predicting movements in share prices is where we can make the money, so that is what we are going to do first.

WebMind was a total simulacra. Its business management had convinced itself that WebMind could produce whatever predictions they wanted. After the fact it seems amazing that the WebMind insight into consciousness --- that intelligence is a matter of complex networking and its myriad of emergent interactions --- was thrown overboard the minute applications had to be conceptualized. If complex organization is a process of emergent networking, then it is in principle not predictable in the form that WebMind management claimed. The only possible conclusion is that WebMind's senior management never understood the implications of the claims they were themselves making! When they made aggregation level shifts from computing to social organization and back again, they became muddled and confused.

Up for Discussion

WebMind was represented at one of the Managing the Complex workshops, in 1999. By then, programmers had been hired and an office opened. Intelligenesis -- as Webmind was then called -- had thirty some employees and was actively engaged in raising its first venture capital. The group discussing WebMind saw the company as a badly functioning network of relationships in danger of self destructing.

Management reacted to this assessment by claiming that there was an evolutionary process going on wherein the environment was pulling Intelligenesis in many directions and its network of employees was pulling it in still many others. The thirty programmers were already at work while senior management was frantically trying to organize the whole thing to move ahead and follow threads and narrow down threads, and to run the company. Management claimed that it was so busy with crisis management that it had become hard to communicate, to take the time. The first thing that goes, is taking the time to communicate what you're doing. Goertzel was moving off in one direction and the business management claimed to move off in another.

An executive from Sun Microsystems joined Intelligenesis as COO, yet he was virtually immobilized in his efforts to financially manage the company. Before Andy Siciliano came on board, company expenditures were strictly aligned with the resources in the bank. Afterwards, it was a virtual free for all. Expensive musical instruments, video games, vcrs, and pets began appearing in senior technical manager's offices, and unopened cartons of new books lined the hallways. One of the COO's first attempts at budgeting involved classifying the 100 or so software programmers into research, application development, etc., the goal, of course, being to reign in the non-productive parts of the company - the research. But Ben, who from the beginning, was canny enough to frame the speculative, experimental work of many of his

programmers as software development rather than research, thwarted the effort as best he could. In the end, Ben would claim that most of the money was being spent on application development, but such spending included the large number of programmers who were working on Webmind-related code and the new AI.

Inevitably, every few months, a new name would appear on the payroll, and the COO, his hands in the air, would confront Ben and Andy. In some late night conversation, it seemed, Ben had lobbied for a new employee - often a friend with little or no software experience, and got the approval from Andy without either of them going through the process the COO so painstakingly tried to implement. One day, the father of the Brazilian office head appeared on the payroll. The office head was one of Ben's "star" computer scientists. The father was apparently in poor health, and the scientist would be able to spend less time working if he had to spend so much time taking care of his father. The effort to remove this fellow from the payroll was almost comic, and in the end, the father, with no relevant skills whatsoever, remained an employee until the day the company closed down.

Neither management nor scientists contradicted themselves in their told stories or admitted to uncertainty in their approach to the future. They were utterly consistent in their roles: Goertzel was the AI genius and the managers were there to finance and exploit his genius. Management never assumed its responsibility for co-defining a future direction for WebMind. Purpose was shoved off onto the abstract idea of 'emergence' --- WebMind's goals were to be emergent. Management thus avoided responsibility for goal setting and for structuring organizational intentionality or purposiveness. Companies and organizations can have serendipitous events -- that is, the embrace of emergent possibilities by a prepared organization. Serendipity is not only emergent; it includes preparation --- that is, opening oneself for certain possibilities (and not others) and being attuned for some (and not other) affordances. WebMind's management was coherent in its avoiding of all these tasks, which contributed to WebMind's incoherence as an organization.

Goertzel's self-assessment:

My management style is highly related to chaos. I'd say right now our organization is almost entirely programmers and that I have a complexity philosophy of managing people. Because WebMind has so many aspects, I've looked for the most brilliant people with the most ideas of their own -- that relate to what I do. And I hire them. I try to let them do what they've always dreamed of doing within a WebMind context. So far I've worked by really looking for exceptional people. The people who I've hired have influenced the direction which WebMind is going. The applications come from the business people I've partnered with. This has worked so far. What this has meant is that because everyone's doing what they always wanted to do all their lives, we don't have

to pay them as much as we would otherwise. , And hands on management, is really not required. So the management mainly consists of collective thinking, working on ideas. But, as far as the work stuff goes, it hasn't been necessary so far to 'manage'. To what extent this kind of vibe can be scaled up, I'm really not sure. If we need to hire fifty people -- it won't be that nice. Then all of a sudden I'm not hiring brilliant AI nerds anymore. I'm hiring people who are good at building software for some application or other. This is something you need to do in a very modular way. You want to say, I'm going to hire someone who I have faith in and he'll be good at managing these people. Then you let him help recruit and manage the people.

The 1999 workshop was not unique in its critical assessment of WebMind. A senior AI expert was hired by WebMind in the fall of 1998 and quit within six months. His assessment:

Goertzel, is a rank amateur data analyst with apparently no experience with real data. His books largely consist of sloppy recombinations of faddish trends, with no mathematical backing, and with no effort to deal with real data. His design for Webmind has duplicated his sloppy pattern with great accuracy. His original design is amazingly flawed. The simple fact is that Ben is so incompetent that he could not produce even a simple search engine design until forcefully pushed into acknowledging the problems with the original design, {He succeeded in} concealing from management the fact that the original design is a sham. However, this process took four months with many false starts by me, and although I have seen Ben scramble in order to make a last-minute attempt to put the implementation in place to enable the simplest of information retrieval searches, I have not seen any sort of process changes that will allow the company to survive by working-around the current design bottleneck of Ben. The current approach Ben is taking appears to be quite the product of a drug experience, and will require significant redirection and grounding in reality, requiring immense effort by someone willing to pay attention to the psycho-babble being produced in lieu of design.

No one in WebMind was willing to hear such a critique (full disclosure: including one of your authors who at the time was the CFO). The WebMind insiders saw their problem as technology versus business, or research idealism versus hard nosed reality. Business management found it attractive to play the victim and to blame Goertzel. Business management never critically examined --- or let be examined --- the technology development. Nor did it develop a coherent vision of the potential practical applications of WebMind. Even if WebMind had been sound science, management had little or no demonstrable ability to make use of it. There was little or no real leadership coming from management and – after the fact – many claim very powerful opportunistic direction provided by Goertzel.

Throughout 2000, arguing about strategy continued. WebMind hired a consulting firm to help bring focus. More than \$1 million later all that resulted were bills. Top management summarized their position: "The mission of this company is to build applications using general intelligence as a base. Consequently, many of our resources will be assigned to this. If you believe that WebMind will never work or that the general intelligence will never be commercialized either alone or with other layers, then you are working for the wrong firm." WebMind marketed a product to analyze Internet chatrooms. It marketed its 'superior categorization' technology, but to a marketplace unwilling to pay for categorization. It bought a services company and attempted to reorganize around services plus categorization. Staffing was increased by more than thirty people. It tried to market its technology through a spin-off called collectivelysharper, but without proven results nothing came of it. It tried to sell itself to a German newspaper as a 'superior search and categorization tool' --- never mind that the system did not know German. The deal fell through in recriminating arguments over the viability of the technology and over money. Webmind died. The money was gone, but the arguing and the accusations remained.

There were obviously warnings about the AI but management did not want to consider them. WebMind was a technology development, driven organization; with a senior management that had insufficient knowledge of the technology to be effective. Goertzel probably played divide and rule politics with his developers --- in so far as he was in control, only his favorites got the facilities and time they wanted. There were enough scientists along for the ride that the incriminations only came after the bankruptcy. WebMind used the homology of 'complexity' as an excuse for disorder and weak management.

Thereafter

Intelligensis never became credible to outsiders (other than one major outside investor brought in with the new president – who in turn became disillusioned rapidly, Richard Li). Development was always behind schedule. Unlike some technology start-ups, in which the creator of the technology hands over its development to other managers and moves into a research role, Goertzel remained actively involved in development. Though some very high power managers were hired, Goertzel held the scientists together and could pretty much do what he wanted. Since people had invested in the core technology more than in the (potential) spin offs, Goertzel's position was very strong. Nothing really came of the spin offs.

Goertzel claimed to notice an extreme difference in reaction between people closely involved in engineering the AI Engine and others. Some not close to the AI Engine were annoyed about all the time, money and effort spent trying to build real AI. The people working on the AI Engine have a different

feeling. They could feel the success, the momentum, the power in the work that had been done. As for his original business partner; he emailed:

Of course I'll miss you a lot "in principle", but in practice, we really haven't worked happily together in a long time. I suspect that we'll have a better relationship, and more interesting conversations, AFTER you leave the company.... What I'll miss most is the hypothetical world in which we really ran the company together harmoniously and synergistically. But, there are lots of hypothetical worlds I miss, by this age I've learned to live with it pretty comfortably...

She answered:

Do you remember my quote in our first WSJ article? I said something like 'Ben loves ideas for the ideas themselves. I love ideas for their usefulness.' Well, what a canny foreshadowing that was!! The fitness function for a business is profitability, and the fitness function for technology IS its usefulness/commercial viability. While indeed Webmind may one day prove to be useful beyond those 'incremental' technologies out there, I was dead on in my insistence that, once it became obvious you weren't going to develop anything productizeable anytime soon, that we should allocate resources for the creation of more simple technologies that would have allowed us to generate revenues and keep up with the firms in our space. Not only did I see the side of the mountain we were heading towards, but I knew exactly what it looked like. Ben, you've done this firm and our investors a great disservice for being so pigheaded about this – if Webmind, Inc. goes down, and the odds of this are greater than 50%, it will be because of this. Had (everyone) listened to me, we'd have been a real contender in the marketplace TODAY, and would have had far more runway to eventually succeed in developing Webmind. Moreover, by learning right away how to build discrete functionalities rather than having as your key goal the coding of an abstract mathematical process, you'd have had a better understanding of how to build Webmind in the first place - you didn't quite understand how difficult transforming mathematical equations into real computer processes would be. You do understand that we've wasted millions of dollars heading down dead ends – all completely amateurish mistakes that even I recognized over two years ago. Regardless of the genius in your Webmind conception, no experienced technologist would ever have taken some of the paths we did. Certainly, the creation of a totally new technology involves making mistakes, but they were compounded by the violation of very simple rules regarding software teams, process, and productization. You simply didn't know (or couldn't

acknowledge for the sake of ego) what you didn't know. And even when you finally got a grip on how massive the complexity of the Webmind project was, you consistently either underestimated what it would take for us to produce anything commercializable, or you simply misrepresented it ... so (management) wouldn't take some of your marbles away. Sometimes I even suspect, Ben, that on some level you knew me to be correct, but marginalized me anyway so that you could plow on ahead, uninterrupted.

And, some ex-employees piped in:

To the someone entertaining thoughts of acquiring Webmind tech. whatever you end up doing, you may be well advised to treat it as a buyer beware situation. Get full demos in real time. Get full details on the demo setup. Ensure you understand the conditions of the training and testing sets. Ensure disclosure on how comparisons are performed. Be aware tests on limited data or over limited periods of time are simple to rig. And be aware, subterfuge is often how the uneducated become vocal supporters. Pay attention and note no supporters have come forward on these pages. There is ample evidence suggesting there are none. And if at all possible, and with any situation where you have the weak hand because you do not understand the depths of the potential subterfuge, hire a due diligence firm.

There is no evidence of any previous Webmind investors having taken any of these obvious steps. Read all papers and books written by the principal. You may be surprised at how often claims and reality differ. You have been warned. Good luck.

The wise have gotten out of the way. No one who lived the experience will invest or spend time with him. They are all working hard to ensure there is no encore performance, with a great deal of success.

Thanks Ben you asshole! Oh, yes, I put up with the thousands of daily e-mails about how you're a "great leader". Yeah, there are issues. Like how to keep a serial scam artist from striking again!

Must fully concur, Intelligenesis/Webmind was a true rat pack. It sucked in a lot of Ben's friends (former!), and several very good researchers, who then were only used for their names to raise money ('If they work here we must be good').

WebMind did not respond to criticism of any sort. Ben's technology is vacant, pitiful, poorly designed, poorly articulated, and poorly thought out.

WebMind demonstrates a complete misunderstanding of the scope and the science, and it presents nothing new not found in already operational research efforts. And yes, the implementation works just about as well as Ben describes it, i.e. not at all, it is an unclever noise generator. Ben always misrepresents the excitement those of us working on the engine felt - it was all we could do to keep from snickering when Ben entered the room, and especially when he was trying to lecture about his newest 'insight'.

Oh yes, and to those ridiculous prospective investors, others will concur, the only technology that ever worked was not developed nor owned by Webmind, it was either borrowed or licensed! So find out from where and take your money there! Geez.

Believe it, no situation this screwed up could exist without an enabling body and for a while at least, Ben did have management eating off of his plate.

Goertzel has been making wild claims followed by grand excuses for years. He claims an injustice in a torrent of criticism, waaah. Isn't the truth elsewhere, in published lies, in sloppy error ridden books, in a host of badly written drivel of his own, in a consensus he cannot abide criticism, and now in this situation, where he is clearly being punished by a community for wasting its time, for providing no product and no substance while dragging along, to their great detriment, a large number of caring individuals?

Enough becomes enough, especially of this guy.

In retrospect, WebMind was perhaps just another e-business hype. Goertzel wrote it an obituary:

Rationally speaking, bankruptcy was a clear and present danger right from the start. The seed money we got was only enough to last us a few months, and we had no idea where the next round was coming from. But we never really considered we might run out of money. ... I was the ringleader. I'd recruited the others to help me out with the not inconsiderable task that I had set for myself: creating and commercializing the world's first truly intelligent computer program. Creating a thinking machine, and then commercializing it? Well, fine, but how are you going to make money while you're creating the damn thing?

The global brain wasn't something any one person could create, not even me; I knew that. It was something that had to emerge. There were things you could to seed its

emergence, encourage its crystallization. One of these things, I believed (and still do believe), was the creation of real AI. By creating a computer program that could read all the information on the Internet, understand much of it, and place new information back on the Internet, enhancing the environment that it lived in, one could cause a phase transition in the development of the Net. One could jolt the Net into a different state of being – effectively causing the emergence of a new organism, a new form of intelligence and life.

In fact, one of the partners of our first VC firm used to raise money for us like this. He'd tell people: "This is the best cocktail party investment you'll ever see. Think about it this way. If these guys succeed in building a thinking machine and predicting the financial markets, you'll be made incredibly rich, and you can say you got in on the ground floor of something remarkable. If they fail, well, at least you can tell people you were in on a really interesting swindle. Either way, you'll have something to talk about!" People ... had the intuition that we were really clever, and would do something really cool. As one investor said at the time, "Will the software work? Well, that depends.... What does 'work' mean? I'm sure it will do something interesting." Of course, this faith placed in me and my team by strangers was flattering. I felt it was largely justified. We really did have a better idea about how to make computers think. ...

This is a small example of what happens in the business world when money is cheap and experimentation is plentiful. A lot of things get tried, and some succeed, and some fail. The things that succeed get to grow and transform themselves into different forms, and combine with each other through mergers and acquisitions. It's really nothing more – or less -- than evolution. ... [But] the point comes where you have to fight for survival. The environment shows no mercy.

One thing that hurt us pretty badly was our failure to come up with a sufficiently sexy spin-off product, leveraging some of the abilities of the AI Engine codebase in a widely sellable, easily marketable way. There's so much wonderful, intelligent code there, and we on the AI side never did a very good job of making it clear to the business folks exactly what this code could do. We needed some kick-ass visionary marketers, and we should have hired them even if it meant getting rid of a few of my super-brilliant mad scientist software engineers.

I do regret very much that we failed to make Webmind Inc. a profitable business. Partly this was due to my mistakes, partly due to the mistakes of others (I wasn't running the company, and my views were by no means always adhered to ...), and partly it was just

plain bad luck. I also regret not making a real AI as we set out to do, of course. One thing is for sure, we made a hell of a lot of progress in *understanding* how to make a real AI, which we are now incorporating into our ongoing work. It was an expensive way to learn the lessons about AI that we learned ... (though frankly AI research was nowhere near a majority of the firm's expenditures) ... but building a real AI is a hard problem ...

So what next? Well, companies are, in a way, more complicated than individual living organisms. If a human dies, its arms and legs can't get up and walk away and find new sources of sustenance. On the other hand, if a company dies, it is quite possible for its parts to live on – to find ways to sustain themselves (thus becoming whole organisms), or to graft onto other organisms that are better able to support them.

The unity of the company as a whole is gone, but in truth the company was somewhat fragmented anyway. What's happening now is that each group within the company that had its own "natural unity," is going its own way and trying to become its own company. Time will tell how many of these attempted spin-offs succeed. Evolution's work is never done. Ultimately most important is not money but the intersection of ideas and people who grasp them. This is what always really changes the world – for the better or for the worse. Money may come and money may go, but a team of people with a common idea can do more tremendous things than any amount of cash. ... The ideas generated by the company will live on, and it looks like many of the teams formed within the company will live on too, each one associated with a certain cluster of ideas.

To which his business partner reacted:

Ben's main disconnect on the business side was that he confused marketing with sales. Ben, clearly, was a great salesman and he blamed "the business side" for not being able to sell Webmind software. I explained endlessly that only investors bought ideas, not software procurement folks. There was no Webmind software. Again, the only usable components in our arsenal was the open source categorization software we were packaging and a rather simple "reasoning" algorithm that one of the senior scientists had invented before he came to Webmind. And in marketing these components, we were up against companies with similar core technologies that had already been built into full-blown applications and were proven in the marketplace. He had great difficulty with the concept that marketing involved identifying a core competency within the company that gave us some type of advantage over competing products, identifying the market to which this competency would have the greatest value, and then packaging it accordingly.

It wasn't about the sales pitch or marketing prose, it was about offering value. I spent similarly exhausting sessions with Andy Siciliano trying to explain basic marketing concepts. Customers don't care about technology, they care about solving a problem. One of the key problems when Webmind Search 1.0 was theoretically finally "working", it wasn't appreciably more intelligent than other competing systems, but it was much slower and less scalable. What would possibly make a potential customer buy Webmind search as opposed to the other systems, assuming we found one to whom intelligence was a key value driver? His response: "The promise of having the most intelligent search engine in the world in the future. Customers would want to "get in on the ground floor" with Webmind so that they'd have the best search engine down the road."

At the end of the day, it is little wonder that the "business side" was so badly maligned in the context of Intelligenesis/Webmind. The technology vision ruled the company and was entirely incompatible with any business-minded endeavors. Any credible attempt to make Intelligenesis/Webmind a business would have necessarily demanded the complete abandonment of the "webmind vision". In fact, a telling conversation occurred when the COO was interviewing a potential employee in late 1999. The interviewee asked what would happen if Ben Goertzel left the firm. The COO thought for a minute and then replied, "Actually, it would probably be a good thing."

While this account adequately describes the mayhem that existed at Webmind in its insular little environment, it fails to recognize *why* the company had no interest in truly being a company at all. The simple fact was that when Andy Siciliano came on as CEO in mid-1999, Intelligenesis no longer had the need for market discipline. It no longer needed to even attempt to be a viable enterprise to remain alive. Andy had deep pockets (of the \$20 million the company burned through, more than half of it came from Andy), and enough people willing to invest in him personally without worrying much about the business model.

Indeed, Andy not only "drank the Kool-Aid", he got drunk on it. He spent his time learning about the technology rather than the technology business (the business folks brought on later were even savvy enough to purposely avoid Ben and learn as little as possible about WebMind's links and node structures). Those amongst the long-time senior business staff who railed against the endless research expenditures and resource-starved application teams were marginalized. The newer ones didn't bother to fight the wave – they just laughed about how the inmates were running the asylum behind closed

doors. It was clear to all the business professionals and consultants involved with the company that Ben should be “put in the closet” and a real software professional should be made CTO, and Ben agreed as well – but only when the time was right, which of course, never happened. The board was unable to intervene because, quite simply, Andy was paying the bills.

Even Danny Hillis, the Thinking Machines founder who “vetted” Ben’s technology plans on a conceptual level and found them to his own way of thinking in the very early days of the venture conceded that the best course of the action for the company was to jettison the research entirely (“Just give it away if you can find someone willing to support it”) and focus on immediate commercial applications. Ben waved off Danny’s assessment by saying that Danny was simply jealous that he had not succeeded at “real AI” and didn’t want to see anyone else to either.

By the time Webmind’s door’s closed in April 2001, I had never seen any “Webmind” software in operation. Certainly, there was a common algorithm in use here and there in the simplified demos we used on various occasions, but the “core architecture” never produced any usable useful commercial functionality. The “initial” version of the “psynet core” supposedly formed the platform for the first version of the market prediction application, but the market prediction developers complained that it added only complexity, technical difficulties, and no useful functionality, and ultimately, after a near-rebellion, they were permitted to rebuild their system as a standalone. About a year later, a “demo version” of “Webmind” was installed on my laptop for a search demonstration in London. That was to be the first time I had ever seen “Webmind” in action, and when the system performed rather impressively in front of a major consulting company, I was almost moved to tears. Later, in my hotel room, I literally *gushed* to the tech staff mailing list about how wonderful it was for me to finally see “webmind”, how well it worked, how great they all were, etc. The programmer who put it together regretfully wrote me back that all he did was rig a couple of algorithms together in half a day, that this wasn’t, as Ben had told me, “Webmind” at all.

The only original *working* software that the company ever produced was the market prediction technology, and this was a relatively simple construct that had no relation to the Webmind architecture the company was forever trying to build. Even the categorization software the company later marketed was based on WEKA, an invention out of the University of Waikaito. So we spouted the “Webmind vision” repeatedly in our sales efforts, and then had to concoct an explanation to customers as to why they

also needed to license some public domain open source software along with our “product”.

And one observer noted:

An organism must learn to adapt to its environment or die, and our situation in the very early days was much like every other start-up that stumbles and falls down as it learns to walk. The difference was that with Andy’s funding spigot open wide, Webmind never needed to learn how to walk at all. We just hopped into his car and drove off a cliff.

Resonance is not a guarantee of business success.

Ben Goertzel’s vision for AI resonated strongly with those who wanted to believe. WebMind’s simulacra brought in lots of money. Goertzel was able to assemble an enormous research team of very talented people from the States, to New Zealand, to Brazil, to Russia. He reigned – however briefly -- over a scientific empire. WebMind could exist because everyone involved – the researchers, the investors, its own management --- colluded to let it exist. They all preferred the simulacra of Goertzel as scientific genius and WebMind as brilliant science, to the alternative of Goertzel as a mythomaniac and WebMind’s management as delusional. Shared resonance is a form of belief and belief can be ‘false’. At the time of the bankruptcy a common comment was: “They drank the kool aid.”

As WebMind shows, very strong belief can be just as false as very weak beliefs. The blaming mechanisms whereby Goertzel and his followers and the management and its followers each blamed the other party for the failure of WebMind continued long after the company had ceased to exist. In fact, Goertzel has shown more self-criticism than have the managers. After all, what is the chance, that a total AI genius would be walking around as a total outsider unrecognized by the research community, though he has a PhD and worked at various universities? If he’s so good, wouldn’t his peers have noticed it? Part of the Goertzel mystique was the ‘something for almost nothing’ idea. The investors were supposedly getting genius at bargain basement prices. Wasn’t this why they did not ask enough questions?

Business is based on belief --- on trust, anticipation and convictions. Coherence is one of organization’s more powerful beliefs. Goertzel and WebMind, changing the world through AI, and the prospect of making money, all served as charismatic models around which stories could be told and upon which homologies founded. WebMind was held together by willpower, ego, self-delusion and greed. WebMind never succeeded in creating the enabling environment, which could afford shared coherence.

‘Goertzel’ (the man, the idea, the phantom) was a shared simulacra --- he made himself, or let himself be made, into an icon for ‘genius’ as a focus for ‘success’. WebMind’s ideology/homology was ‘self-

organization' --- the Being or ontology of the organization would create its success. This Being was the simulacra --- real self-organization requires constant attention to the affordances. WebMind preached networking as the principle of consciousness and intelligence, but it did not network. It was never a part of the research community or of the IT corporate world. It was a stand alone outsider, enclosed in its own myths. It did not have an open questioning culture but a repressive one very much controlled by Goertzel and a few others. The researchers were not allowed to follow their own initiatives or hunches. There was no rigorous critical discussion of the AI work. WebMind had the secrecy and direction of a corporation and the inefficiency and lack of control of the research world. In effect it had the worst of both worlds.

There were those whose self-definition and focus was money based, those whose self focus and definition was based on Goertzel's supposed scientific charisma, and those who relied instead on the notion that they could help to change the world. Each carried their own meanings into a shared homology built upon Ben and his AI. The notion of a shared community was a cohesive force that only acted to draw each of the three varying perspectives more into self-enclosed isolation. Each formed its own community, fairly incommensurate with the others. When anyone was challenged, an underlying homology of identity or necessity consistent, coherent, but retrospectively false – for Goertzel of 'emergence' and 'self-organization', and for the managers of the necessity of 'how business is done'— was called forth.

We observe the following:

- 1) There was an ever increasing binary opposition between those who viewed WebMind's activities as research oriented and those dedicated to making WebMind a viable business. As WebMind's existence became ever more precarious, each group defended its internal coherence by opposing the other group more forcefully. It was a fairly standard case of 'scape-goating'. If coherence was to be experienced as a unity across this opposition, such unity was possible only when the firm was faced with "evil" outside forces – the pressure to create a working demo, the pressure to attract more funds, the need to divert efforts of most of the organization to deliver something to a real client, etc.
- 2) The story of WebMind --- told both to each other and to the world --- suffered from an ever increasing disconnect from the actual activities occurring within and about the firm. The software was "said" to be working (but was it?), the business plan claimed sales efforts (but did they exist?), the research effort was described as on the verge of a major breakthrough (though breakdown was more accurate). If coherence was to be experienced as emergent, then the language used to tell the Webmind story, would have had to be brought into a better fit with the experiences and situations which made up that story.

- 3) The “selves” at WebMind found little “group” cohesion. Instead the selves were pushed into attending ever more so to them-selves. The communities of practice at WebMind were allowed to become fiefs and like medieval fiefdoms they resisted the needs and demands of a unifying force. What self-organization there was, was temporary and usually forced by the imposition of an external force. While the ‘software’ was supposed to epitomize self-organization, it only worked when attended to rather forcefully by an external meddler. The emergent consciousness at WebMind was of an organization in trouble, of software that did not work and of a desperate hope for that next infusion of cash would “give us the chance to prove everyone wrong.”

Towards the end, WebMind had deteriorated into a whole series of gloms. These were personal stories of fights and dreams, conflicts and intentions, none of which came together anymore in any shared coherent stories. Webmind echoes the findings of Heracleous and Barrett:

A view of organizations as constituted of fragmented, competing and, less often, complementary discourses; [with] clashes among stakeholder groups over contested terrain, [where what] fragile agreement and cooperation [that existed] at the communicative level was based on potentially conflicting deep structures that could assert themselves in different ways under different contextual conditions. ...discursive fragmentation, leading to conflicting actions, stakeholder groups talking past each other, and the lack of common ground on which to base a dialogue.

What was lacking was a dialogic homology. Experiential coherence entails emergence being socially shared. Coherence is not an instantaneous superficial experience of situational sameness, followed by confusion and chaos. It implies a shared meaning and signifying apparatus of some rigor and sustainability. Coherence may be temporary, but it is not entirely fleeting. The tradition of Habermas --- operationalized, for instance by Deetz --- also stresses the organizational necessity of dialogue. (Habermas, 1984; Deetz, 1986) In that tradition, the lack of coercion is the key factor – dialogue can only take place if the participants feel that they are ‘free’ to think and express themselves.

In WebMind, there was far more collusion than coercion. WebMind is not a story of too many or too few constraints, but one of the wrong constraints. Everyone tried to accelerate the research and product development, when the deceleration of the simulacrum and myths of performativity was what was needed. Trying ever faster to get to the goal was self-destructive. WebMind needed to (re-)consider itself and its intentionality, if it was to survive. In networked intelligence, it is not the part of the brain

that reacts the most intensely, that conveys the subtlety of meaning. Weak signals provide the nuances, the contextual details and the finer aspects of meaning. WebMind got all the weak signals wrong. Instead of networking to its environment, it avoided it; instead to dialoguing with difference, it channeled all interaction into a few primitive conflicts.

WebMind's coherence was too tight --- it denied the value of loosely bound systems, of vague zones of possibility and of multi-interpretable significances. WebMind's sense-making was not loose enough --- it did not provide enough room for weak signals. The organization was locked into its conflicts and oppositions, without the space needed for emergence. The affordances were very powerful -- what WebMind claimed to do was something the scientific and business communities were shouting for. What a 'thinking machine' could really think; or how AI could serve humanity, was never really considered. Goertzel's AI is a Frankenstein waiting to happen --- and, as it should be, the monster ate up the organization.

The irony in all this is that ten years later the technology is proving itself to have value – absent the company of which it was a part. The major news/data organizations (Reuters, Dow Jones, Bloomberg) have products which make use of the underlying WebMind idea and patent infringement actions have begun.

Substitute the Housing Bubble

Part of the skill in being a real estate developer is predicting future demand. The next few stories about housing bubble are of developers who mistook the emergent boom for a new rule. "Housing prices can only go up." As the real estate boom expanded during the bubble years, developers and home buyers believed that residential golf resorts were a sure-fire bet. Many buyers looked to buy properties that they could flip for a quick profit. Others were lured by stunning views, club services and security. The aggressive building of new resort courses continued from the mid-1990s into the 2000s, contributing to an increasing glut of inventory that finally found no market.

From the 1970s to the 1990s, Al Hoffman's WCI built thousands of homes in the Tampa Bay area. It was all very nice, but not necessarily the Ritz. "Housing is pretty basic and it's very unglamorous," Hoffman said in 1995. At the height of the housing boom, speculators around the country flocked to Florida. With its balmy weather and miles of beachfront, the Sunshine State was a natural lure for real-estate buyers. As WCI's chief executive, Hoffman expanded the business by cutting deals to build championship golf courses, luxury resort hotels, houses and condominiums. Hoffman focused the company on affluent baby boomers, who economists predicted would be a key driver of development in Florida for decades to come. "You can't stop it," Hoffman told The Washington Post in 2003, describing the development march from southwest Florida toward the Everglades. "There's no power on Earth that can stop it!"

But it stopped. With eerie results. One of WCI's projects, Tuscany Reserve in Naples Florida saw more than \$200 million spent on landscaping and development. At the end fewer than 25 homes were ever sold. Similar fates await other projects. Off the turnpike in central Florida, hidden behind stucco walls, sits a sprawling Tuscan-style clubhouse on a hill overlooking a string of lakes, a golf course and green fields. This 1,900-acre property, called Bella Collina, was designed to hold 800 homes. Today, only 48 houses dot the landscape, and just three are occupied. The clubhouse, though open, is eerily quiet, and a promised swimming pool and equestrian center have yet to be built. Bella Collina, the brainchild of Edward Robert Ginn III, (Bobby Ginn) looks like a ghost town. So does Tesoro, another resort opened by Mr. Ginn near Port St. Lucie, where just 150 houses sit on 900 lots. And the Conservatory in Palm Coast, also from Mr. Ginn, is even more barren: 335 out of 340 lots are empty. During the real estate boom Ginn grew the Ginn Resorts company to 13 properties. No one developed a resort better than Bobby Ginn, and when times were better there was no better place to buy property.

Ginn had done this before. In 1986 Ginn was quickly making a name for himself, though not for the right reasons. He owned eight of 20 golf courses (among other properties) in Hilton Head Island, S.C. and his company was failing. Ginn cut a deal to acquire the two oldest and biggest developers in Hilton Head, S.C., only to see the venture collapse under a pile of debt amid the nation's savings-and-loan meltdown. Ginn failed to the point that employee paychecks were bouncing and the residents of Hilton Head drove around with bumper stickers that read: "Honk if Bobby Owes You."

Ginn started his new company, Ginn Resorts, based in Celebration, Florida in 1998. Credit Suisse arranged a total of \$675 million in loans to Ginn in 2006, according to credit rating company Standard & Poor's. According to S&P, Ginn used the money to develop five properties: Tesoro, a gated community in Port St. Lucie, Florida, with two golf courses, one of which was designed by Arnold Palmer; Quail West in Naples, Florida; Hammock Beach River Club in Palm Coast, Florida; Laurelmor, a community under development in Boone, North Carolina; and Ginn sur Mer, a resort on Grand Bahama Island, the Bahamas. Mr. Ginn's development style was unusual: he didn't build clubhouses or other services until a large number of lots were sold. "Typically, developers start with a hotel or amenities because skeptical buyers want to see things," Robert Gidel, a former president of the Ginn Companies told the New York Times. "But here the market was so strong, and people wanted to believe."

Ginn's timing was as bad this time as it was in the 1980s. Sales at the developments "have been severely affected by ongoing economic pressures and the drastic downturn in the real estate market," Ginn said in a court filing. Mr. Ginn says his remaining properties will eventually pay off. "My belief is that when the depression ends, there will be a pent-up demand for happiness," he told the New York Times. "Sometime between 2035 or 2040, Florida will double in size." Mr. Ginn is "ready to sell properties in trophy locations" when the market turns around. "If you can't sell," he said, "you die."

The golf resort stories reveal, just as with WebMind, the power of a resonant homology to bring forth action. There is nothing in that argument which suggest that the action which so emerges will be good, successful, or productive. Resonance provides a collective power to shared homology. It fuels the coherence which those share that resonance perceive about a given situation. It allows for affordances to be grasped which might otherwise have been overlooked. But the ultimate success of the coherent actions will not rest upon either the resonance or their coherence. Powerful tools do not ensure the skill of an amateur craftsman. Great skills do not ensure commercial success.

Constrained Ambiguity Powers Homology

The beauty of cognitive homology lies in its power to generate multiple words, alignments and meanings, appropriate to a specific situation and aligned with its identity. Managers frequently discard distinctions, introduce new distinctions, use old distinctions in new situations, put new words into new contexts, and use distinctions in a metaphorical sense. Their thinking is a complex process, in which images, concepts and schemata are always struggling to adapt to each other. The pieces of the conceptual puzzle, form networks of relations. Changes in a particular time or place ripple throughout the web. Meaningful stories are not made up of isolated words. They must evoke deeply held values and images. To offer up isolated words, is to evoke a shallow stream of water, in a hot desert. Whatever value there is dries up quickly. The multiple ambiguous reactions to homology are constrained by situation and context. Indeed, it is situation and context which allow the “model” to both make predictions and to help shape the possibility space.

Events are local, partial and unique --- events define specific moments of activity and change. For an event to persist, it has to be named and communicated. Events are a product of emergent change, which is situated and temporary. The actions that reveal new possibilities and alternative ways of being are often very short in duration. Such events (often) create something new - a form of being that is different from what previously existed. Emergence is what makes the organization more than the sum of its parts. It is the recognition of new items, when scale or levels (for instance between persons, groups, organizations and societies) change, or perspectives shift. Emergence makes organization exciting, creative and powerful --- the sum of the parts is much more potent, than the elements alone.

Shared understanding only occurs occasionally. Much of the time, relationship is impersonal, perfunctory and bureaucratic. Rule-bound behavior gets us most of the time through our daily activities. The social environment provides necessary structure. Concrete situations of work and organization are dialogical and communicative. Communication takes place in the 'in-between' of organizational existence ---- in the spaces where people meet, talk and participate. Interaction follows from undecidability and indefiniteness. Because identity, as well as goals and strategy, are always uncertain, shifting and (somewhat) indeterminate, people always have reasons to talk. Contingency --- the

indefiniteness of organization; makes for dialogue. Dialogue is not a product of organizational design or structure, or of unilateral control, or of rules and authority. An organization either finds dialogic cohesion or (eventually) it can no longer function. For a person, total breakdown in interaction leads to schizophrenia; for an organization it leads to crisis and/or to being dismantled.

Dialogic resonance, on the interpersonal scale, manifests itself as felt understanding between 'self' and 'other'. On the organizational scale, it reveals itself in the relationships between the persons, and between the collectivity and its environment. Organizational resonance requires the social 'in-between' of individual openness and discussion. But it also requires the exploration of boundaries and possibilities.

To these questions we now turn.