

A version of this appears as:

"Leadership and Influence: The Manager as Coach, Nanny and Artificial DNA" in J. Clippinger (ed) THE BIOLOGY OF BUSINESS: DE-CODING THE NATURAL LAWS OF ENTERPRISE (Jossey-Bass: San Francisco,1999) p. 47-66

CONTROL & INTERVENTION IN COMPLEX ADAPTIVE SYSTEMS: FROM BIOLOGY TO BIOGEN

Andy Clark

*Philosophy/Neuroscience/Psychology Program
Department of Philosophy
Washington University
St. Louis, MO 63130
e-mail: andy@twinearth.wustl.edu*

draft of chapter for John Clippinger (ed) THE BIOLOGY OF BUSINESS

0. ABSTRACT/INTRODUCTION.

Markets, companies and various forms of business organizations may all (we have argued) be usefully viewed through the lens of CAS -- the theory of complex adaptive systems. In this chapter, I address one fundamental issue that confronts both the theoretician and the business manager: the nature and opportunities for control and intervention in complex adaptive regimes. The problem is obvious enough. A complex adaptive system, as we have defined it, is soft assembled and largely self-organizing. This means that it is the emergent product of multiple, and often very heterogeneous, interacting factors and forces, and that the crucial interactions are not controlled and orchestrated by an overseeing executive, detailed program, or any other source of strict hierarchical control. There is thus a pressing problem -- are such systems strictly out of control and beyond the reach of useful governance? I shall argue that they are not. Such systems, though initially unfamiliar, can nonetheless be led, influenced and enabled in a variety of ways.

I begin then, by clarifying the nature of the problem. What is it about complex adaptive systems that makes control and intervention so difficult? I then (section 2) pursue a biological example: the role of genetic factors in influencing mature form. This example yields some lessons that can be applied to business organizations (section 3). Section 4 pursues a speculative extension of these lessons, applying them to the task of understanding the role of tagging and explicit models in certain kinds of advanced system. There is a brief conclusion.

1. AGGREGATION, SOFT ASSEMBLY AND CONTROL.

Aggregation is defined (chapter x) as a process in which multiple (often heterogeneous) elements interact so as to yield a distinctive collective effect, and in which the relevant interactions are not controlled and orchestrated by any distinct overseeing element. As an example, consider the way the individual, amoeba-like cells of the cellular slime-mold cluster together (when food is short) to form a mobile slug-like mass (the pseudoplasmodium). This process does not, as it was once believed, involve leader cells which issue a chemical call to aggregate. Instead, when food is low *each* cell releases a chemical (cyclic AMP) which attracts other cells into local mini-clusters. These mini-clusters then constitute denser concentrations of cyclic AMP, which in turn attract more cells. The ensuing process of positive feedback (in which mini-clusters promote larger and larger clusterings) yields the pseudoplasmodium via a process of self-organization in which the higher-level pattern (in this case, the slug) emerges from the interactions of multiple elements without the benefit of any Atop-down@ control.¹

Closely related to this notion of self-organized aggregation is the process dubbed (Thelen & Smith (1994)) soft assembly. Consider an infant learning to reach. The developmental psychologists Thelen & Smith argue² that individual infants confront very different problems in learning to reach. For a very active infant (one who thrashes about continually), the problem is to dampen the thrashing and to convert it into directed reaching. For a more passive infant, the main problem may be generating the initial burst of energy needed to overcome gravity and

¹ For a fuller account, see Ashworth & Dee (1975). See also Resnick (1994), p. 51.

² Thelen & Smith (1994), Ch. 9.

generate torque. In each case, however, the problem that confronts the brain and the central nervous system is how best to massage and exploit the existing dynamics of the system so as to increase the likelihood of successful reaching. To do so is to seek a solution that takes into account multiple factors and pre-existing sources of order within the system. It has thus been argued³ that the job of the central nervous system is not generate a rich inner plan, or explicit model of how to reach. Instead, it is to learn how to modulate a few simple parameters in a bio-mechanical system of spring-like muscles and masses. Modulating limb stiffness, for example, may interact with intrinsic spring-like dynamics so as to yield a brief oscillation whose resting point is upon a desired target. The controller (the central nervous system) is not a locus of detailed blueprints for success. It is rather a source of subtle prompts, nudges and minor modifications that act upon an ordered biomechanical system, situated in an environmental context. The job of the controller is to modulate parameters (stiffness, etc.) that will then interact with the rich dynamics of the target-system-in-context so as to yield a desired outcome.

Processes of aggregation and soft assembly thus suggest a more subtle -- and powerful -- role for the wider environment itself. Instead of seeing the environment as simply a source of problems and an arena in which problem-solving processes are played out, it becomes necessary to view aspects of the environment as equal partners in extended, soft-assembled, problem solving. Consider the expert bar-tender. Faced with multiple drink orders in a noisy and crowded environment, the expert mixes and dispenses drinks with amazing skill and accuracy. But what is the basis of this expert performance? Does it all stem from finely-tuned memory and motor skills? By no means. In controlled psychological experiments⁴ comparing novice and expert bartenders, it becomes clear that expert skill involves a delicate interplay between internal and environmental factors. The experts select and array distinctively shaped glasses at the time of ordering. They use these persistent cues so as to help recall and sequence the specific orders. Expert performance plummets in tests involving uniform glassware, whereas novice

³ Polit & Bizzi (1978), Jordan et al (1994), Thelen & Smith (1994).

⁴ See Beach (1988), Kirlik (in press).

performances, though worse overall, are unaffected by any such manipulation. The expert has learned to sculpt and exploit the working environment in ways which transform and simplify the task that confronts the biological brain. Aspects of the working environment thus form part of what I will call our *cognitive scaffolding*: the external sources of order and influence that enable embodied biological brains to solve complex problems.

This idea of scaffolding has its roots in the work of Lev Vygotsky, a Soviet psychologist of the early part of the 20th century. Vygotsky stressed the way our interactions with external structures (including other agents, words, and texts) could reconfigure our problem-solving activity. My use of the notion of scaffolding is broader still, encompassing all kinds of external aid and support including simple environmental structuring (as in the bartender example) and the more diffuse and complex kinds of support provided by larger-scale linguistic, social, and economic factors. These factors likewise sculpt, simplify and orchestrate the activities of multiple biological brains. Much of the distinctive power and scope of advanced human reason is rooted, it seems clear, in these larger collective efforts. But the parallel runs deeper. For certain larger-scale organizations (such as corporations and businesses) may themselves be usefully seen as adaptive entities that must respond rapidly and flexibly to changing environmental pressures, and may themselves (as we shall later see) interact closely with a variety of extra-systemic factors and forces. Multiple types of systems and actors thus exhibit the special, but pervasive, kind of natural organization that allows them to adapt fluidly to changing circumstances, and to piggyback on a variety of extra-systemic props and aids. All such cases raise pressing problems of control and guidance: how, given such diffuse, multi-factorial, self-organizing and soft-assembled organizational forms, can we (or nature) intervene so as to push such systems in specific directions? We have already seen some hints of a solution in our discussion of the way the central nervous system might promote successful reaching. In the next section, I look at the clearest natural example of successful intervention in complex regimes: the role of DNA in the production of biological form.

2. GENES, COACHES AND NANNIES.⁵

According to a certain traditional view, genetic material constitutes a kind of detailed recipe for the form of an adult organism. To know the complete sequence of an animal's DNA would (if the recipe idea were true) allow you to predict the adult form on just a few minimal assumptions (such as the presence of sunlight, sufficient food, vitamins, etc.). In this vein, a participant in the Human Genome project once asserted that the chromosomes in the fertilized egg constitute:

the complete set of instructions for development, determining the timing and details of the formation of the heart, the central nervous system, the immune system, and every other organ and tissue required for life

Delisi (1988) p. 488.

It is now increasingly clear, however, that such radical instructionism wildly overstates the role of the genes. What the genes actually encode looks to be a more minimal instruction set whose effects depend heavily upon a variety of non-genetic sources⁶ of form and order. These non-genetic elements do not merely feed a pre-specified process. Instead, they play an active role in determining the form and characteristics of the mature organism.

A simple example⁷ is sex determination in Mississippi alligators. In this case, the chromosomes in the fertilized egg do not even select the sex of the offspring! Instead, the ambient temperature of the rotting vegetation in which the eggs are laid fixes the baby alligator's gender (26-30° C: all female; 31-33° C: mixed; 34-36° C: all male). Instead of regulating sex directly via the genes (as in humans), the alligators thus exploit a kind of external lottery that appears to serve them equally well.

The case of chick imprinting provides a slightly more complex example.⁸ Imprinting occurs when newly hatched chickens become rapidly attached to the first mobile object they see

⁵ For a fuller discussion see Clark & Wheeler (1997).

⁶ See e.g., Goodwin (1994), Maturana & Varela (1987), Kelso (1995), Keijzer (1997).

⁷ See Goodwin (1994), p. 38.

⁸ See Elman et al (1996), Johnson (1997).

-- normally, the mother hen. But it is not a simple, pre-wired response. Instead, imprinting is the emergent effect of the interaction of multiple processes, some genetic and some environmental. A neural system first disposes the chick to prefer stimuli that resemble the head/neck configuration of a bird. This causes the chick to expose itself to lots of perceptual inputs from the mother hen. A second neural system, constrained to operate only for mobile stimuli of a certain size, then kicks in, allowing the chick to learn to recognize the highly attended object from multiple angles, etc. It is the combination of these two genetically determined factors, along with the loop out into the hen-attending behavior of the chick, that finally yields the translation-invariant representation of the mother hen.

Moving up to the human case, it has been estimated that a full molecular specification of a human body would require 5×10^{28} bits of information, whereas the human genome contains only about 10^5 bits (Callow (1976), reported in Elman et al (1996), p. 319). The entire human genome would not even support a detailed plan of the human brain. As a result, nature is clearly forced to rely on a variety of less direct methods: the provision of well-chosen prompts and nudges applied to a system (the developing organism) that has its own intrinsic properties and dynamics, and that is embedded in a complex environment rich in natural order. Mature form can thus only be explained and understood by looking at the complex interactions that take place between the genetic materials, the developing organism, and the wider environment.

The biologist Brian Goodwin uses the term morphogenetic fields (Goodwin (1994), p. 88-92) to describe the aspects of natural (cytoplasmic and environmental) order that, in close concert with various genetic prompts and nudges, yield the distinctive forms of plants and animals. He shows, for example (p. 105-119) that the arrangement of leaves on plant-stems (phyllotaxis) is not directly genetically determined but instead results from the interaction of stress patterns (during growth) and the way the emergence of one leaf alters the patterns of cellulose resistance that determine the probable location of the next leaf (and so on). The most common pattern in nature involves leaves emerging in spirals at a rotation of 137.5 degrees. Goodwin deploys a variety of modeling studies to show that given typical initial conditions and

values for stress, resistance and rate of leaf-formation, the 137.5 degree pattern will indeed predominate, with other (less typical) natural patterns arising from other (less typical) initial conditions. In this example, then, one salient and robust feature of the mature plant is controlled not by direct genetic specification but by the complex interactions between what *is* genetically specified (some aspects of the initial conditions) and a morphogenetic field of mechanical stresses and compensations.

The lesson, it seems, is that DNA *never* constitutes a full, detailed blueprint of the structure of any complex life form; instead, the DNA must gear its own strategic nudges and influences so as to make the most of a variety of natural sources of order (the morphogenetic fields, the normal ecological environment of the chick, the temperature of the alligator-egg-concealing vegetation, etc.). Indeed, some quite striking features of biological form (such as plant phyllotaxis) look to result almost entirely from the interplay between various non-genetic sources of order and constraint.

But genes *do* matter. Genes are, indeed, Nature's way of exploiting and channeling all those complex, heterogenous and non-linearly interactive sources of natural order. The trick is to see that genes work by nudging the system along at crucial points and by actively creating the conditions under which extra-genetic sources of order and form can be maximally exploited. We are thus invited to think of genes not as full and detailed⁹ instruction sets, but as catalysts (Elman et al (1996), p. 351), seeds (Goodwin (1995), p. 16) or modifiers (op cit, p. 144). Catalysts, to pursue one metaphor, are individually inert. Yet placed in the proper context, a catalyst can have a dramatic effect: it can initiate and maintain a chemical reaction that would otherwise not occur (Elman et al, op cit, p. 351-352). The active contribution of the genetic elements can thus be characterized along two (related) dimensions. The genes create and maintain the local conditions under which self-organization, aggregation, and soft-assembly can occur. And they actively

⁹ There is, indeed, a debate concerning the appropriateness of conceiving the genes as codes or instruction-sets *at all*. My own view is that it is perfectly proper to view the genetic contribution as a kind of minimal instruction set: a recipe that succeeds only given the reliable presence of multiple sources of extra-genetic order and control. See Clark (in press), Clark & Wheeler (1997), for full discussions./

nudge or influence the unfolding and self-organizing system at crucial points in its development. To borrow one last example from Elman et al (1996):

The *catalase* gene produces an enzyme which allows reactions to occur at rates greater than 1,000,000 per second. The gene does not define the conditions for reaction. These conditions [are]... specified by the laws of biochemistry. [] What genes do is to harness those laws by ensuring that critical components are present at the right time, and then nudging the reaction forward (p. 352).

Genes thus act as nanny and as coach: the nanny (ideally) creates and maintains an environment in which a child's own curiosity and intelligence will lead her to learn and flourish; the coach looks out for difficulties and potential sticking-points and tries various ploys and tricks to push the player onwards. But neither one can simply force their charge to conform to a detailed blueprint of behavior and action. Similarly, the neural controllers of motor action do not specify the activity of each individual muscle and tendon -- instead, they work with the synergetic biomechanics of the body and with the constraints and opportunities offered by the local environment. The common principle that unites all these cases is the need to reconceive control not as the top-down imposition of pre-specified form but as the selective application of relatively simple forces to a highly complex self-organizing system.

A major thesis of the present work, is, of course, that markets, companies and various forms of business organizations may all be seen as instances of complex adaptive systems. What, then, are the implications for understanding business organization and strategies of intervention?

3. NEW ORGANIZATIONAL FORMS.

Consider the case of the biotechnology industry. The sociologist and theoretical economist Walter Powell discerns, in the biotechnology market, the emergence of a new logic of organizing:¹⁰ one that depends critically on (what we can now see as) principles of aggregation, self-organization and soft-assembly. The key features¹¹ of the biotechnology case are:

¹⁰ Powell (1996), p. 197.

¹¹ This list is an attempt to distill the main lessons of Powell (1997), somewhat influenced by the related analysis of Stark (1996).

- 1) The exploitation of open corporate architectures capable of making maximal use of external resources.
- 2) The use of minimal hierarchical control structures both within individual companies and between interdependent ventures.
- 3) Acceptance of waste and redundancy as the natural cost of a continuing search for productive collaborations with other ventures.
- 4) The development of complementary, symbiotic or co-specialized skills and assets as a stable and productive mode of extended organization.

Let's unpack all this in a little more detail. The biotechnology industry operates in a high-technology, research intensive, high-uncertainty market. The typical products are new, biologically engineered drugs, vaccines and medicines: for example, Activase and Neupogen -- major new drugs for the treatment of heart attack and kidney failure (developed by Genentech and Amgen respectively).¹² The development of such drugs is time-consuming, expensive, unreliable and highly knowledge-intensive. By comparison, traditional pharmaceutical companies operated in a much more stable and less knowledge-intensive regime. The focus was on direct competition to discover drugs that would help with a few major ailments and the research procedure was to choose a demographically promising disease, define a model of it, and then to simply take compounds off the chemist's shelves and screen for efficacy (Powell (1996), p. 204). The biotech companies depart from this methodology by probing deeper into the causal and genetic mechanisms and trying to directly engineer designer drugs. What seems to have emerged in the last decade, however, is a kind of new, symbiotic organizational form in which the big pharmaceutical companies (Glaxco, Roche, Ciba-Geigy) outsource research activities to bio-tech companies (which in turn work in close synchronization with each other and with university groups), and in which the big pharmaceutical companies shoulder the burden of 1) sustaining the extended clinical trials needed to get a new drug on the market and 2) the complex business of promoting and marketing drugs world-wide.

¹² See Powell & Brantley (1996).

What we here confront, I suggest, is a kind of environmentally scaffolded, soft-assembled solution to a daunting business problem (and opportunity). To display the full flavor of such a soft-assembled organizational mode, consider Biogen -- a biotechnology company with 410 employees and at least 37 close ties to outside groups spanning the big pharmaceutical companies, university departments and hospitals. Powell reports that more members of the core Biogen workforce work on products outside the firm than within it (op cit, p. 200). Or consider the recent development of an animal model for Alzheimer=s research courtesy of a >symbiotic= collaboration between 34 scientists, two biotech companies, a pharmaceutical company, a federal laboratory and a university department (Powell (1996), p. 205; Nature, February 9, 1995). Examples could be multiplied, and Powell and his colleagues have gathered a formidable array of data over an eight year period. But the general lesson is exemplified by Biogen -- the company that had more ongoing activity outside it's formal boundaries than inside (op cit, p. 206).

Biogen displays all four of the key features sketched earlier. First, it is by no means a self-contained drug-discovery and marketing operation. Instead, crucial functions are off loaded to external ventures -- much as genes exploit local morphogenetic fields to press maximal benefit from minimal internal encoding. Second, this easy dovetailing with external sources of order is positively encouraged by a loose internal control structure. Biogen is organized around shifting project teams and individual researchers are granted a great deal of freedom, pursuing their favored lines of research with merit pay geared to publication. Such researchers are encouraged to seek out and form external links as required for their own specific projects. (A Harvard Business School report calls this managed chaos.) Moreover, no strict hierarchy prevails among the various collaborating groups. What forms is a complex, multiparty web in which it is exceedingly difficult to pinpoint the center or starting point (op cit, p. 209). Third, Biogen is tolerant of the inevitable episodes of overlap (several teams pursuing closely related projects) and failure. Indeed, Stark (1996) stresses the way such organizational forms foster innovative search at the inevitable expense of on-the-spot optimality and efficiency. By maintaining an open

organization tolerant of multiple, overlapping, often dead-end, efforts, a company increases the likelihood that when some new environmental opportunity arises it will be more closely positioned to benefit (Stark's example: it is (short-term) efficient to seek food where you have found it before. But it is a mistake to lock in that efficiency by suppressing the capacity to explore new territories, as the large-scale distribution of rewards may one day change). Expensive internal diversity is thus one key to long-term adaptability. Putting all this together we derive the fourth feature (really a summary of the other three): Biogen has developed a set of internal expertise that is carefully adapted to allow it to participate in multiple, extended dovetailed ventures. Considered in isolation, Biogen is not so much an impoverished unit, as an intrinsically interactive one; one that is highly specialized so as to contribute specific skills and capacities within the context of a wider soft-assembled business network.

The immediate task of management, then is to create the internal conditions (sketched above) that allow the discovery and maintenance of multiple collaborative endeavors. The manager's task is to promote these collaborative links, to actively encourage internal diversity, to tolerate overlap, redundancy, and failure, and generally to act as a coach and nanny, as artificial DNA, catalyzing reactions and actively creating the conditions necessary to press maximal benefit from multiple sources of external order and opportunity. The role of management is not to create detailed blueprints for performance or change. Rather, it is to create the conditions for exploration, de-centralized adaptation, and environmental exploitation.

4. MODELS, TAGS AND EXPLICIT CONTROL.

We have, however, left something out. For the genes have dovetailed their contributions, by trial and error over evolutionary time, to the reliable extra-genetic features of a stable, normal ecological backdrop. But the dovetailing and externalization we are witnessing in various high-technology markets and endeavors (see also the discussion of the Internet as a self-organizing domain in chapter x) is driven by a backdrop of continual change and instability. Nature's solution to the need to achieve rapid, cheap dovetailing with multiple, changing, unstable environment is, surely, the biological (and especially the human) brain: those brains, that is, most

strongly marked by cortical plasticity and flexible learning. One distinctive feature of advanced human cognition is our capacity to promote successful environmental dovetailing both by basic learning and by a variety of more reflective tricks and ploys. But given an image of the business world as an interacting nexus of fuzzily-bounded, soft-assembled coalitions, it is initially unclear exactly what kind of productive role we might assign to more reflective practices, such as the creation of market models, the use of tags, symbols and labels¹³, and the issuing of explicit instructions. Yet properly conceived, such practices can greatly enhance the efficiency and flexibility of the newly emergent organizational forms.

Let's begin by considering the fundamental and under-appreciated process of tagging: the process, that is of providing a label for some salient object, feature, or process. Tags (for products, teams, etc.) clearly support easy recognition and simplify communication. But they also (more importantly) facilitate model-building, thought and problem-solving. A striking example is provided by recent work on pan troglodytes -- a breed of small chimpanzee. Thompson, Oden and Boysen (in press) investigate the ability of the chimps to solve puzzles that involve the detection of higher-order relations i.e., relations-between-relations. Spotting a situation in which a simple relation (such as the relation of *sameness* between two identical cups) obtains is a first order task well within the reach of chimps and other animals. A higher-order task would be, for example, to match one pair of identical items (the cups) to a different pair of identical items (such as two identical shoes). Or to match a pair that satisfies the relation different (e.g., a shoe and a padlock) to another pair that satisfies the same relation (e.g., a banana and a cup). The higher order task is *not* to match real world objects but to match the abstract relations that obtain between such objects.

Thompson, Oden & Boysen show that the higher order task is tractable only to chimps that have been trained to associate concrete tags or labels (specifically, plastic shapes) with relational features, so that e.g., a plastic circle tags sameness while a plastic square tags difference. The

¹³ It was Daniel Dennett who first alerted me to the importance of tags and labels. See e.g., Dennett (1994) (1995).

authors conclude that it is practice of tagging and labeling itself that opens up new problem-solving horizons to the tag-trained chimps.

My speculation (Clark (in press)) is that the experience of tagging provides the chimps with something closely akin to a new perceptual modality. Like a tuned perceptual modality (vision, touch, etc.) tagging depicts certain features of our world as salient: it makes them pop-out rather like concrete, basic objects. Perceptual tuning involves learning to treat complex sensory patterns as specifying objects: objects that can then be attended to in ways that reveal patterns among them (e.g., the relations-between-relations in the chimp case). Experience with the tags allows the chimps to treat the sameness of two items as a simple object (the complex relation is replaced by the simple tag). This effectively reduces the higher order task (spotting relations-between-relations) to a lower order one (spotting relations *between tags*). Such a process obviously iterates: a being sensitive to several higher-order relations can, if need be, treat them all (courtesy of tags or labels) as new simple objects and extend its search space accordingly. Human science is, I believe, a testimony to the power of this process of iterative labeling and search.

In sum, tags play a dual role. They evidently aid communication and interpersonal understanding. But they also contribute to the fundamental processes of thought and reason by providing new domains over which mental operations may be defined (a more vigorous computational analysis of this idea is provided in Clark & Thornton (1997)).

The cognitive benefits of tagging need not, however, be confined to individual thought. When a group, team or company confronts a problem, it constitutes an extended cognitive system within which public tags and labels may again play a fundamental (i.e., not merely communicative) role. To bring this idea into alignment with the image of the company as a complex adaptive system, we need to stress, once again, the particular way in which tags and labels can contribute to the problem-solving process. And once more, there is a partial parallel with the genetic case. For the tags and labels (like DNA) do *not* participate in detailed plans and

instructions that rigidly guide systemic activity. Instead they act in a variety of more subtle, modulatory and catalytic ways. In particular, tagging and labeling can:

- 1) help create and maintain adaptively valuable diversity
- 2) help to create and stabilize the multiple external alliances characteristic of the new organizational forms
- 3) participate in processes of *diffuse control*
- 4) promote the kind of systemic self-understanding that underpins effective (albeit indirect) intervention and control

Regarding the creation and maintenance of adaptively valuable diversity, consider Stark's (1996) anecdote concerning the roadside rabbit breeder whose sign reads Pets and Meat. This is a case in which a single resource is tagged twice, allowing it to participate in two very different economic regimes. Similar tag-driven diversity, Stark notes, is common in the emerging markets in Eastern Europe. Here, a company's survival can depend on its capacity to account for itself according to multiple standards. For example, a company may be eligible for central support only if it employs mainly local workers. But it may also be assessed, by outside investors, in terms of net worth and market share. Companies operating in economies subject to rapid change and reorganization do well to maintain a variety of identities and self-justifications, switching from one to another as opportunity and politics dictate. Multiple, evolving overlapping and sometimes cross-classifying tagging of products, resources and personnel looks messy and inefficient from a kind of neat, engineering and design perspective. But it may be the economic equivalent of the ubiquitous biological strategy of multiple overlapping functionality: the use, for example, of a single CPG (central pattern generator) to generate quite different motor outputs during different operating regimes (see e.g., Selverston et al (1997), p. 114).

Tagging can also help create and stabilize the kinds of powerful inter-group alliances highlighted in section 3. The role of tagging here is to keep the collaborative elements in the sweet spot between (too much) disorder and (too rigid) order. If the collaborations are becoming too complex, costly and inefficient, it may be fruitful to allow only a few highly successful units

to actively seek and create new collaborative ties. Such units will be tagged, within the company and for the benefit of outside interests, as effective collaborators. The use of tags to thus stabilize, specific interactive regimes is somewhat parallel, I suspect, to the role of the tags in the chimp experiments described earlier. In each case the process of tagging creates a simple, easy to identify and manipulate, entity that can subsequently participate in increasingly complex regimes (as when, for example, a university department and a project team at a Biotech company are labeled as an effective extended unit and recruited at a single stroke as an element in a new collaborative venture). But tags, labels, and guidelines can equally well play an important *destabilizing* role in cases where a system strays from the sweet spot in the other direction -- the direction of too much order. For example, if the number of successful external ties is low and stagnant, it may help to create new project teams and to issue new labels and guidelines, encouraging experimentation and rewarding *attempted* collaboration as much as successful joint endeavors. In all these cases, tags and explicit guidelines function not to directly instruct the system (the company) but to modify the background conditions against which processes of diversification and self-organization operate according to their own intrinsic dynamics. (See also the related discussion of Internet dynamics in chapter x above). The use of tags, labels and guidelines thus constitute potent methods of *diffuse control*: rather like the way the release of a certain chemical or neurotransmitter may bathe a whole population of cells, temporarily altering their global behavior.

Finally, a good set (or sets) of tags and labels are the essential pre-requisite of explicit systemic self-understanding and self-analysis. The extent to which such self-understanding is possible within a complex adaptive system is currently unclear, since the simplifications and distortions inevitably imposed may well be intolerable in the context of a system of multiple, non-linearly interactive, parts. Yet, it seems clear that some attempt at self-modeling is often fruitful and can play a role in successful intervention. Moreover, if we are to apply new computational tools to the understanding of market behavior, it is necessary to adopt some kind of description language in which a company (for example) is represented as a structured entity

comprising a set of features and relations and assessable (as fit/unfit, successful/unsuccessful) according to some measure. Without such a rendition it is impossible, for example, to use genetic algorithms to help explore the space of possible organizational forms.

At this point, however, the lacuna in our story looms large. For the problem of the origination (discovery, creation) of useful sets of tags and labels is surely one of the major unsolved problems in contemporary cognitive science (see e.g., Dennett (1994) (1995), Clark & Thornton (1997), Deacon (1997)). The survival of a company in an evolving, volatile market may depend crucially on its adopting the right tag for its products (as when telephone companies began to identify themselves more broadly with information and telecommunication, or when entertainment companies are reborn as media and communications giants). But the process by which a good tag is discovered or created remains annoyingly ill-understood. As a result, there is currently no substitute for human intuition and creativity as the prime originator of effective tags and labels. We can, however, add a kind of complex adaptive systems *caution*. The best tags and labels will emerge as successful modulators of the ground-level interactions that already characterize the working of the system. It is not likely that tags and labels simply imposed by high-level management will be properly geared to effectively modulate the flow of activity. A better strategy is for management to closely observe actual working practices and to pay special attention to the vocabularies and quasi-formal groupings that emerge from these base-level interactions. The issuing of new tags and labels should be inspired by -- not imposed upon -- these emergent organization forms. (Compare, again, the case of the Internet (chapter x)): here, recent attempts by government to fix the number and ownerships of domains looks like a case in which the complex adaptive systems caution is not being heeded, and top-down imposition threatens to undermine (rather than to sculpt and streamline) effective emergent practice).

In sum, tags and labels are powerful (but still somewhat ill-understood) tools for the analysis and control of complex adaptive systems. Contrary to traditional top-down models, such resources cannot simply be seen as imposing order on some relatively plastic base. Instead, the tags too constitute subtle forces that nudge systemic self-organization in a certain direction. They

are not mere shorthand for existing dynamics. But they are not elements of effective blueprints either. Better, perhaps, to see them as seeds or anchor points around which new organizational forms can soft-assemble; cheap but effective parts of the environmental scaffolding that supports complex problem-solving and flexible response.

5. CONCLUSIONS: MANAGEMENT AS ARTIFICIAL DNA.

Complex adaptive systems cannot be efficiently controlled by overseeing executives, detailed plans, or other sources of strict top-down control. Instead, the role of management must be to create the conditions for flexible, self-organizing response and to modulate the unfolding of intrinsic dynamics by the judicious application of simple forces and nudges and by the use of a variety of diffuse influences. A concrete example of such strategies in nature is the way DNA modulates (but does not fully describe) developmental processes. The biotechnology industry, we have argued, displays a surprisingly similar organizational form: one in which successful companies exploit a wide variety of external links and collaborations and in which the role of management is primarily to maintain the company in a sweet spot between fixed and inflexible organization forms and counter-productive anything goes states of anarchy and decay. We have also begun the hard task of reconceiving the role of tags, labels and explicit models in the context of such an adaptive systems perspective. Such resources cannot play the traditional role of directly imposing form and order. But they can, we suggested, play a powerful indirect role by acting as the seeds or anchor points around which new pockets of organization soft-assemble.

Central to our more theoretical picture, then, is an image of individual companies as (increasingly) dependent on a process of *dovetailing* to an evolving set of external sources of order and opportunity. Dovetailing occurs when (as in the biotechnology case) the best shot at market success involves keying a given company's expertise and resources very closely to what might be termed a *collaboration niche*. In such a niche, a company contributes an essential, sometimes unique, but always insufficient, set of skills and knowledge to a project that spans multiple companies and institutions. An individual company is thus rather like the expert bartender from section 2: successful performance depends not on the bartender's internal

resources alone, but on the way the bartender adapts those resources to exploit a variety of powerful external sources of order and problem-solving. As a result, the boundaries between our core systems (companies, biological organisms) and the wider environment become more flexible, less clear-cut. This blurring of boundaries has practical, moral, and legal consequences. It affects our conception of the very nature of individuals and of businesses, and raises deep questions concerning accountability,¹⁴ reward and the basic units of moral and economic significance.

¹⁴ For example, it may be unclear whether it is the biotechnology firm, or its loose academic partners, or the temporary union of both, that should be held accountable for some costly mistake or fatal error. Moreover, given the picture we develop, we cannot assume that a relevant boundary, once established, will continue to fall in the same place. Instead, we may confront a kind of fluctuating nexus of organization and order.

REFERENCES.

- Ashworth, J., & Dee, J. (1975). *The Biology of Slime Molds*. London: Edward Arnold.
- Beach (1988).
- Calow, P. (1976). *Biological Machines*. London: Arnold.
- Clark, A. (Ed.). (in press). *Magic Words: How Language Augments Human Computation*. Cambridge: Cambridge University Press.
- Clark, A., & Thornton, C. (1997). Trading Spaces: Connectionism and the Limits of Learning. *Behavioral and Brain Sciences*.
- Clark, A., & Wheeler, M. (1997).
- Deacon, T. (1997). *The Symbolic Species*. New York: Norton.
- Delisi (1988). The Human Genome Project. *American Scientist*, 76, 488-493.
- Dennett, D. (1994). Learning and labeling: commentary on A. Clark & A. Karmiloff-Smith. *Mind & Language*, 8, 540-548.
- Dennett, D. (1995). *Darwin's Dangerous Idea*. New York: Simon & Schuster.
- Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (Eds.). (1996).
- Goodwin, B. (1995). *How the Leopard Changed its Spots*. London: Phoenix.
- Johnson, M., & Bolhuis, J. (199x). Imprinting Predispositions and Filial Preference in the Chick. In R. Andrew (Ed.), *Neural and Behavioral Plasticity* (pp. 133-156). Oxford: Oxford University Press.
- Jordan, M., Flash, T., & Arnon, Y. (1994). A Model of the Learning of Arm Trajectories from Spatial Deviations. *Journal of Cognitive Neuroscience*, 6(4), 359-376.
- Keijzer (1997). *The Generation of Behavior* (Doctoral Dissertation, Leiden University, Netherlands).
- Kelso, S. (1995). *Dynamic patterns*. Cambridge, MA: MIT Press.
- Kirlik, A. (Ed.). (in press). Everyday Life Environments. In W. Bechtel & G. Graham (Eds.), *A Companion to Cognitive Science*. Oxford.
- Maturana, & Varela (1987). *The Tree of Knowledge* (New Science Library).
- Polit, A., & Bizzi, E. (1978). Processes controlling arm movements in monkeys. *Science*, 201, 1235-1237.
- Powell, W. (1996). Inter-organizational Collaboration in the Biotechnology Industry. *Journal of Institutional and Theoretical Economics*, 152(1), 197-225.
- Resnick, M. (1994). *Turtles, Termites and Traffic Jams: Explorations in Massively Parallel Microworlds*. Cambridge, MA: MIT Press.
- Selverston, A., Panchin, Y., Arshavsky, Y., & Orlovsky, G. (1997). Shared Features of Invertebrate Central Pattern Generators. In P. Stein, S. Grillner, A. Selverston, & D. Stuart (Eds.), *Neurons, Networks, and Motor Behavior* (pp. 105-118). Cambridge, MA: MIT Press.
- Stark, D. (1996). Heterarchy: Asset Ambiguity, Organizational Innovation and the Postsocialist Firm. Paper presented at the Annual Meeting of the American Sociological Association, New York.
- Thelen, E., & Smith, L. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.