

From population to organization thinking*

David Lane, Robert Maxfield, Dwight Read, Sander van der Leeuw

Abstract

This chapter begins by reviewing the Darwinian account of biological innovation, which is based on what Ernst Mayr calls “population thinking” and posits two kinds of key mechanisms underlying the innovation process, variation and selection. The chapter then argues that the increasingly popular tendency to adapt this account to provide the foundations for a theory of human sociocultural innovation is ill-advised. Human sociocultural organizations are self-reflexive and self-modifying, through negotiation processes that can lead to transformations in organizational structure and functionality, including the essential activities of recruitment, differentiation and coordination. Innovation in these organizations is accomplished through processes of organizational transformation, and to understand how these work, “organization thinking” rather than “population thinking” is required. The fundamental questions that organization thinking addresses include the following: What is social organization? How are particular social organizations constructed, maintained, and transformed? What kinds of functionality do social organizations support, and how do they create new functionality? In addressing these questions, the chapter describes a bootstrapping dynamic, whereby organizations generate new functionality, which is instantiated in activities that in turn generate new organizations.

1. Introduction

Our species is still very young by biological time scales, and it is too early to know if we represent the cutting edge of a biological success story, like cockroaches or dinosaurs, or a brilliant but ultimately failed and short-lived experiment in niche construction and destruction. In the mere 200,000 or so years of *Homo sapiens*’ story, and in particular in the approximately 50,000 years since we began to accrue the accoutrements of culture like language, art and multi-component artifacts, members of our species have populated a vast extent of the earth’s surface and exploited for our own purposes an ever-increasing share of the planet’s biologically utilizable solar energy. In the last few centuries, we have ravaged the stock of bioprocessed solar energy accumulated over millions of years, transformed minerals extracted from below the earth’s surface into a huge variety of forms and new materials that satisfy what we regard as our needs, and increasingly concentrated the human population in urban spaces to which nearly all the raw materials necessary for human survival have to be imported from elsewhere. At the individual level, our first *Homo sapiens* ancestors managed to keep themselves going on the 100-300 watts their bodies were able to generate, assisted in their quest for survival by the handful of artifacts they knew how to make and use; in contrast, current residents of New York City mobilize on average about 10,000 watts to propel them through their daily rounds of activity, and the shops in their city offer them a choice of something like 10^{10}

* Chapter 1 of Complexity Perspectives on Innovation and Social Change, edited by Lane, van der Leeuw, Pumain and West (Springer, 2009)

different kinds of artifacts to help them accomplish whatever it is they might feel inclined to do.¹

How have we managed to accomplish so much so fast? The main premise of this chapter is that we have done it through a new modality of innovation, through which human beings generate new *artifacts* that they embed in new *collective activities*, which are in turn supported by new *organizations* and sustained by new *values*. Over time, this new innovation modality gave rise to a positive feedback dynamic, which explains how we have generated so many transformations in our selves, our societies, our culture and our environment.

What is this new innovation modality? We begin by describing something it is *not*. Much of modern biology is based upon Darwin's theory of biological novelty, which analyzes the processes through which species come into being and are transformed, by means of mechanisms of heritable variation and selection. Given the tremendous scientific success of this theory, it is not surprising that many authors have sought to adapt it to other contexts. In particular, it is becoming increasingly fashionable to construct theories of innovation in human society and culture on a Darwinian foundation. We shall argue that this move is mistaken.

To be clear about the claim we are making, we need to define precisely what counts for us as a Darwinian foundation. An evolutionary theory seeks to understand a phenomenon by describing the processes that brought that phenomenon into being and that generate the transformations it successively undergoes.² A *Darwinian account* is a special kind evolutionary theory that, like the theory Darwin set out in the *Origin of Species*, rests on two foundational bases:

- it is characterized by what Ernst Mayr (1959, 1991) called *population thinking*; and
- it analyzes the transformation processes with which it is concerned in terms of two *fundamental* and *distinct* classes of mechanisms: variation and selection.

We discuss these two concepts, in the context in which Darwin introduced them, in the second section of the chapter, and we explain why they were such a great success in this context. We then describe some issues that must be resolved before they can be successfully applied in other contexts.

In the third section, we examine some features of human sociocultural innovation that distinguish it from biological evolution. We argue that these features are not consistent with the foundations underlying a Darwinian account. In particular, the distinction

¹ See Chapter 12. The number of artifacts refers to the number of different SKU's – that is, product labels used for distinct bar codes – on offer in New York, as estimated and reported in Beinhocker (2006). This is of course a rather crude measure of “kinds of artifacts”.

² In contrast, essentialist theories seek to explain phenomena by isolating their (unchangeable) essences and classifying all other aspects of the phenomena either as deducible consequences of these essences or inessential epiphenomena. For example, a noted neoclassical economist once told one of us, that “if it isn't about optimization and rationality, it isn't economics”. Essentialist theories tend to be associated with the “typological thinking” mentioned in the following section, just like Darwinian accounts are associated with “population thinking”.

between variation and selection processes is difficult to maintain – and even when they can be distinguished, they frequently fail to be fundamental, since another kind of process, negotiation, underlies both of them. In addition, it is often impossible to identify a relevant “population”, which is the critical starting point for population thinking.

In the fourth section, we describe a shift in perspective, from population thinking to *organization thinking*. After analyzing some critical differences between biological organization and sociocultural organization, we propose an innovation modality alternative to the Darwinian account. We call this modality *organizational self-transformation*, and we argue that it bears a relation to sociocultural evolution similar to the Darwinian account’s relation to biological evolution. We conclude by describing some features of organizational self-transformation, in particular the positive feedback dynamic to which it gives rise. This dynamic, which we call *exaptive bootstrapping*, has generated the proliferation of artifacts and organizations that construct our current sociocultural world – and us.

2. Darwin’s theory of biological evolution: what it is, why it works

For over a hundred years before Darwin published the *Origin of Species*, students of natural history had collected a huge body of evidence illustrating what they regarded as a nearly perfect match between the morphological characteristics and behaviors of biological individuals and the environmental opportunities that these individuals exploited to earn their living. For example, the woodpecker’s beak seemed *designed* to drill into bark to extract the insects the woodpeckers ate, while the hummingbird’s beak and capacity to hover could hardly be improved upon as a way to sip nectar from flowers. Such matches between structure and functionality were generally interpreted as evidence for the existence of a benevolent Designer, who had constructed a world able to sustain in harmonious equilibrium all the products of His creation, including the many species of plants and animals that the natural historians were busy describing and classifying.

For these natural historians, a species was an immutable, ideal organization, exquisitely tuned to the exigencies of the environment in which members of the species sustained and reproduced themselves. This organization was more or less perfectly instantiated in all the individuals that belong to the species. Ernst Mayr (1959, 1991) labeled this way of construing biological organization *typological thinking*.

Darwin’s first great accomplishment was to displace typological thinking by what Mayr called *population thinking*. Population thinking *un-reifies* the species. Instead of a timeless *kind*, the species becomes just an *aggregation of individuals*. The species has a beginning, in an act of speciation, and sooner or later it will have an end, when it becomes extinct. Its boundaries – who is in, and who is not – are determined by pedigree: once the species has come into being, no individual may enter unless its parent or parents are already members, and no member can defect, except through death. For Darwin, what counts for a species as *population* is not the commonalities that its members share, but the *variation* among conspecifics. That is, the species as population is characterized not

by a shared organization, but by the *statistical distribution of those features that differ* among members of the population. Thus, not only is the *species* no longer identified with an *ideal* organization, but the very *real* organization of an *individual* is disregarded in Darwin's story, and the individual figures in it merely as a set of distinct features. From the point of view of population thinking, biological evolution is the story of the changes in the distribution in the population of these features over time, including the introduction of new features.³

The species as population is not composed of a fixed set of individuals: the members of the population change over time, through birth and death events. It is critically important to the success of Darwin's story that the features of new individuals are statistically related to the current distribution of features in the population. Fortunately, for features that are *heritable* this turns out to be the case. That is, absent changes due to the mechanisms of variation and selection described in the next paragraph, reproduction – whether it be sexual or asexual – does not change⁴ the distribution of differing heritable features, which, as we have seen, Darwin's theory takes as the variable of interest in the species transformation story.

Darwin analyzed the change in feature distribution over time by distinguishing between two fundamental classes of mechanisms through which change happens: mechanisms of *variation* and mechanisms of *selection*. The former introduce *new* variant features that individuals in the population may possess, while the latter determine which features will *increase in frequency* in the population over time. New heritable features persist in the population, unless they are eliminated by selection mechanisms.⁵ In *Origin*, he called attention to one particular selection mechanism, which he regarded as the primary determinant of the direction of change in feature distribution: *natural selection*. According to Darwin, individuals in the same species were always in competition among themselves to obtain the resources necessary to survive and reproduce. Thus, a heritable feature that helps individuals possessing it to get a reproductive edge over their conspecifics would tend to increase in frequency generation by generation. Such features came to be described as increasing the *fitness* of individuals who possessed them, and natural selection came to be conceived in terms of a stochastic process guided by a *fitness function*, whose value for a set of features represented the relative competitive reproductive advantage of an individual possessing this feature set.

³ Of course, this begs some key questions, which Darwin did not address: just what constitutes a new *feature*, and how can it be integrated functionally with previously existing ones; or which new features – or distributional changes in existing features – constitute the origin of a new species? Biologists since Darwin have addressed these questions, sometimes supplementing but never successfully challenging the basic suppositions underlying population thinking.

⁴ In expectation, with random mating in sexually reproducing species (that is, absent sexual selection, a selection mechanism). Of course, in sexually reproducing species, offspring will not have the same features as both their parents, and so the *realized* distribution after a new individual is born is not identical to what it was before the birth – it is conventional to refer to “random drift” as a selection mechanism when such random oscillations contribute to the directionality of change in the feature vector distribution.

⁵ Here, as in the previous footnote, we follow the convention that classifies random drift as a selection mechanism.

According to the currently canonical version of Darwinian theory, the so-called neo-Darwinian synthesis that welded together ideas from Mendelian genetics to Darwin's evolutionary theory in the first three decades of the 20th century, the key variation mechanisms for heritable features in biological evolution derive from genetic operations during reproduction. Thus, a particular innovation is initiated as a new variant *genotype*, which, in interaction with pre-existing structural or behavioral features and particular environmental contexts, happens to result in a new structure or a new behavior at the *phenotypic* level. This new structure or behavior may get incorporated in some kind of process with already existing behaviors, structures and environmental features, and this new process may provide a survival or reproductive advantage to the individual who possesses it. If so, by natural selection, the frequency of individuals with the innovation will increase over time in the population, and the innovation will count as a success.⁶

From the moment of *Origin's* publication, Darwin's ideas were discussed everywhere, not just in the restricted circles of the natural historians.⁷ This initial high level of interest ensured that many other scientists would continue to probe these ideas – to support them, oppose them, extend them. It took several generations for the initial *succès de scandale* to develop into scientific orthodoxy, but by the 1930's, this too had been achieved. How can we account for the great success of Darwin's theory of biological evolution? We think four reasons, described below, were primarily responsible. The first two of these account for the unusual degree of interest that Darwin's ideas encountered in the first decades after *Origin* and kept them under the spotlight of intense scientific exploration, and the second two were responsible for the construction of the scientific consensus around the Darwinian account as a foundation for a theory of biological innovation:

R1. A non-teleological explanation for structure-function fit

Darwin's theory provided the first really plausible alternative to the hypothesis of intelligent design as an explanation for the extraordinary match in the biological world between structure and function. It differed in two fundamental ways from intelligent design. According to intelligent design, biological organization is *immutable* and *teleological*: each structural feature was designed *in order to* carry out a particular function, and once designed, it need not – and did not – change. In contrast, in Darwin's story, structural change was the *key* to the match between structure and function. New variants might appear in any possible direction in feature space, but only those that provided an advantage with respect to the overarching functional imperatives – to survive and to reproduce – were retained through the generations.

The reaction to a non-teleological alternative to intelligent design was explosive. It empowered scientific radicals like Thomas Huxley and atheistic lay(wo)men like Harriet Martineau to take the initiative in dissolving the claims for divine causality that had

⁶ Note that *how* the genotype contributes to the construction of the phenotype is not part of Darwin's account. This important problem, under the rubric of "evo-devo", is currently a very hot topic in biological research. Evo-devo extends rather than revises the Darwinian account, and the separation between variation and selection mechanisms serves as a strong constraint on the kinds of structures that evolution can produce – for example, new variations, which occur at the genotypic level, cannot be directed toward the provision of specific functions at the phenotypic level.

⁷ See, for example, Desmond and Moore (1991), chapter 33.

subordinated natural history – and science more generally – to religion. Correspondingly, it deeply upset, even scandalized, proponents of traditional values, from scientists like Richard Owen and Louis Agassiz to clergymen like Samuel Wilberforce, who in an 1860 Oxford free-for-all discussion on Darwin's ideas asked Huxley from which side of his family was he descended from apes. Though debates of this sort generated substantially more heat than light, the idea of directional change without a directing intelligence became a central theme in the ongoing cultural battle between an emerging materialistic scientism and traditional deference to established authority and revealed religion. For scientists on both sides of this battle, finding evidence and arguments that bore on the plausibility of Darwin's ideas, whether to support or demolish theory, had a very high priority – and the guarantee of a large, very interested public should they accomplish a breakthrough in this quest.

R2. Demonstrating the plausibility of organizational innovation from gradual feature change: Two analogies

However internally consistent Darwin's ideas were, they were certainly not directly demonstrable from empirical evidence. Indeed, to make his argument plausible, Darwin had to come to grips with a very difficult empirical problem: many natural historians were prepared to argue that the immutability of species was no hypothesis, but observable fact. From Aristotle to Linnaeus to Darwin, no observer of the natural world had seen a new species come into being. True, paleontologists had found what seemed to be fossilized remains of plants and animals that did not seem to belong to any known existing species, and enough progress had been made in the relative dating of geologic strata to convince most scientists that not all presumed life forms – living or fossil – had come into being in the same epoch. But the possibility of extinction and separate epochs of creation did not contradict the assumption that, once created, a species-as-ideal-organization didn't change. Where was the evidence for change?

Darwin answered this question with a spectacular rhetorical move, accomplished through the juxtaposition of two analogies. The first analogy was intended to show that biological form is indeed malleable.⁸ Darwin reminded his readers of the wonders achieved by plant and animal breeders – in particular, dog and pigeon fanciers – in creating new breeds or varieties.⁹ Of course, the breeder plays an essential role here, since he selects individuals to reproduce on the basis of the features he wants to favor. In the wild, Darwin asserted, competition to survive and reproduce takes the place of the breeder. Obviously, this *natural* selection is much less intense than the breeders' *artificial* selection, and consequently will require much more time to accrue observable differences in bodily structure or behavior. Here, Darwin introduces his second analogy: we know, he claims, from the revolutionary work of Charles Lyell in geology what can happen when presently observable microprocesses act over huge spans of time. As he puts it in *Origin*, “a man should examine for himself the great piles of superimposed strata, and watch the rivulets

⁸ This analogy, as Darwin himself points out in the introduction to the sixth edition of *Origin*, had been introduced for this purpose by other authors before him. No other author, as far as we know, had coupled it as did Darwin with the analogy with Lyellian geology.

⁹ Darwin admits that, like all analogies, this one is incomplete, since varieties aren't species: when interbred, their offspring are not sterile.

bringing down mud, and the waves wearing away the sea-cliffs, in order to comprehend something about the duration of past time, the monuments of which we see all around us.” Given enough time, small quantitative changes can produce large qualitative change. So: the same process of differential reproduction, favoring particular features, which we can observe when breeders produce new varieties, can operate over Lyellian “geological time” to produce new species by natural selection.

With this argument, Darwin changed the status of the immutability of species from an empirical fact to a mere hypothesis, and an increasingly implausible one at that – as the torturous evasions of Owen around this issue in the years after 1859, ridiculed in later editions of *Origin*, testify.¹⁰

R3. Carving nature at the joints: the variation–selection dichotomy

When Darwin decomposed the evolutionary process into variation and selection mechanisms, he knew almost nothing about how the former actually worked. He was convinced, though, that they satisfied two properties, which were really all his theory at its most abstract level required of them: they generated variant features independently of their potential functionality, and these features were heritable. These properties sufficed to justify his analysis of the *directionality* of evolution – and hence, in his theory, the origin of species – solely in terms of selection mechanisms, in particular natural selection. The research initiated by the rediscovery of Mendel’s work around the beginning of the 20th century resulted in the identification and detailed explication of the genetics underlying evolutionary variation mechanisms such as mutation and cross-over.

As a result, it became clear that Darwin’s decomposition was not merely an analytic or conceptual move, but really carved nature at the joints. As Darwin had foreseen, they were functionally orthogonal: the directionality of evolution was supplied just by selection, not variation, mechanisms. But much more was true. The principal variation and selection mechanisms incorporated into the neo-Darwinian synthesis differed from each other with respect to their *ontological level* and their characteristic *spatiotemporal scales*. The variation mechanisms involved random changes in the genome, which took place essentially instantaneously, while selection operated in individuals interacting with their biological and physical environments and occurred on a time scale of many generations. Moreover, variants that produced viable individuals with new phenotypic features upon which selection could operate were exceedingly rare, so rare that the time between such events was sufficiently long that selection could in general process them one at a time, attaining equilibrium frequencies with respect to a new innovation without interference from another.

The significance of these level and scale differences between variation and selection mechanisms was two-fold. First, they made it possible to *parallelize* evolutionary research: laboratory scientists explored the genetic basis of variation mechanisms, while field researchers investigated the past history and present operation of directional change through selection. Secondly, the absence of strong interaction effects between the two classes of mechanisms very much simplified the work of theoreticians who sought to put

¹⁰ See Desmond and Moore, 1991, chapters 33 and 34.

together their effects to deepen and extend evolutionary theory – in particular, towards a quantitative theory.

R4. Mathematization

Since Galileo, the epistemological gold standard of science has been the construction of mathematical theories that provide succinct description and quantitative prediction for empirical phenomena. By the end of the 19th century, physics was enshrined as the king of sciences – in large part, because it had married the queen, mathematics. Other emerging sciences, like chemistry and economics, did their best to emulate the king in this respect. Biology seemed hopelessly behind. Despite the efforts of a few fringe players like D’Arcy Thompson, it still resembled its ancestor, natural history, much more than its successful rival science, physics. Though the theory presented in *Origin* was far more general and precise than any other yet introduced in biology, it was anything but mathematical. With the incorporation of Mendelian genetics into evolutionary theory, the situation changed. Ronald Fisher, Sewall Wright and JBS Haldane pioneered the quantitative theory of population genetics, which provided the beginning of what has become a flourishing mathematical foundation for Darwinian evolutionary biology, with a consequent upgrading of its scientific status. In this theory, natural selection is represented by a fitness function, with a natural interpretation in terms of expected offspring from individuals with alternative genomic configurations; the stochastic components of the models derive from genetic theories that can be calibrated with frequencies from genetic experiments.

Darwinian accounts

As with other successful scientific theories, Darwin’s ideas have inspired scientists working on different problems than his. For example, biologists and cognitive scientists have developed interesting and fruitful theories based on Darwinian reasoning to explain phenomena ranging from the construction of immunological and neural organizations during individual ontogeny to cognitive processes like perceptual categorization and even induction.¹¹ Such theories begin with the identification of the essential elements of what we will call a *Darwinian account*: a relevant *population*; and *variation* and *selection mechanisms* for features that vary among individuals in the population. If the members of the population change over time, there must be some mechanism, corresponding to reproduction in evolutionary theory, that guarantees the stability of the frequency distribution of features in the population over time, absent the operation of the identified variation and selection mechanisms.

One notable example of a Darwinian account is Edelman’s theory of neuronal group selection, which describes the construction during an individual’s ontogeny of a neural organization that can support perceptual categorization, the basis for many innovative context-specific behaviors that the individual may generate during its lifetime. In

¹¹ E.g. clonal selection in immunology (Jerne, 1966), neural Darwinism (Edelman, 1987), classifier systems for induction (Holland, Holyoak, Nesbitt and Thagard, 1987). In addition, there are some interesting and successful Darwinian accounts for particular social phenomena, for example Croft’s (2001) theory of language change.

Edelman's theory, the population consists of repertoires of neuronal groups, "collections of hundreds to thousands of strongly interconnected neurons."¹² These groups arise according to mechanico-chemical processes of cell division, movement, death and differentiation, which guarantee that "no two individual animals are likely to have identical connectivity in corresponding brain regions." These processes thus constitute the theory's variation mechanisms. Selection operates on the neural groups, as synaptic strengths increase or decrease, within and between the groups, in response to patterns of activation correlated via re-entrant signaling with sensory and motor activity. In this theory, the repertoires of neural groups are constructed via the variation processes and remain stable (unless they disappear) during the subsequent operation of selection, so there is no need for any mechanism corresponding to reproduction. Edelman's theory shows that biological evolution not only follows the Darwinian account, but that it "engineers" systems which themselves operate coherently with that account.

In general, the relation between successful theories based on a Darwinian account and the phenomena they purport to explain shares the characteristics we claim account for success of Darwin's theory: macrolevel function-carrying structure emerges from non-teleological interaction microprocesses; proposed variation and selection mechanisms are ontologically distinct and causally independent; fundamental aspects of the phenomena of interest are expressible in tractable mathematical or computational models. It is of course *conceivable* that a Darwinian account for some class of innovation phenomena might succeed even if none of the conditions of success for Darwin's theory hold, but it would seem prudent to assign a very low *a priori* probability to such an outcome – and to look elsewhere for the foundations of a theory of innovation in such a context.

Now, neither R1 nor R2 seem particularly relevant for human sociocultural innovation. In contrast to the biological world, the fit between structure and function in the social world is not so evident that it cries out for explanation. Rather, it is *function* itself that seems to need to be explained: what is the functionality associated with such social constructions as cathedrals, horror movies and dog shows? This is not a question that a Darwinian account is equipped to address.¹³ Moreover, as far as R2 is concerned, it is quite unnecessary to demonstrate the existence of large-scale organizational innovation in the sociocultural context: we all *know* that social systems, and the kinds of agents and artifacts that inhabit them change, sometimes drastically and suddenly. Nor will a Lyellian analogy work to relate observable microprocesses to large-scale sociocultural innovation, for two reasons. First, it is unlikely that all or even most large-scale sociocultural innovation proceeds by the gradual accumulation of changes induced by microprocesses. Second, it is even more doubtful whether these microprocesses are themselves sufficiently stationary over long time scales that they could generate large-scale changes, without undergoing such significant transformations that their "observability" becomes irrelevant to predicting long-term effects.

¹² Edelman (1987), p. 5.

¹³ Unless it imposes the Procrustean bed of subordinated functionality and defines the "master functionality" to which cathedral building, horror movies and dog shows are subordinated.

We also can – at least provisionally – claim that R4, mathematization, has yet to tell in favor of a Darwinian account for sociocultural phenomena. As far as we know, the attempts to provide such an account have yet to introduce any new mathematics, beyond variations on population genetics and evolutionary game theory; and the mathematics that have been applied have yet to produce the kind of verifiable predictions and unifying conceptions that marked the work of Fisher, Wright, Haldane and their successors.

Thus, the main issue for a Darwinian account of sociocultural innovation is that raised by R3: do the foundations of a Darwinian account carve nature (or in this case, society) at the joints? We take this issue up in the next section.

3. A Darwinian account for sociocultural innovation?

To give meaning to the question that provides the title to this section, we need to say something about what we mean by “sociocultural innovation”. As it happens, the previous section refers to an interesting, if somewhat surprising, example of what we have in mind.¹⁴ The example comes from Darwin himself, through the analogy he introduces between natural and artificial selection, and in particular the “fancies” that were coming into prominence as a popular pastime and commercial enterprise in Darwin’s epoch – and which so intrigued Darwin for the sheer variety of animal forms the fanciers were able to generate.¹⁵ In our discussion of this example, we highlight four features that are signatures of sociocultural innovation processes. After we describe these features and explore some of their implications, we confront our findings with the foundational requirements for a Darwinian account.

Artificial selection and the dog fancy

As we saw, Darwin discussed artificial selection in *Origin* to highlight the wide range of heritable features that nature provides as grist for selection’s mill. Depending on the selection criterion used, artificial selection can generate varieties from the same ancestral stock that are eventually as dissimilar as, say, tiny Maltese dogs and huge Saint Bernards. For the purposes of Darwin’s argument, the criterion for selection is just a free variable: it could be any feature, as long as it is observable, heritable and stable over time. But if we are interested in analyzing artificial selection as a process of *sociocultural* innovation, then of course we need to understand what kinds of features are employed as selection criteria, and how they become established.

According to Darwin, there is a single primary functionality that underlies all selection criteria in natural selection: reproductive potential -- the capacity to produce the

¹⁴ Of course, Darwin’s theory and the processes through which it became scientific orthodoxy offer another good example. Though at first sight it might seem very far removed from the dog fancy – after all, one seems to be just about concepts and the evidentiary standards of scientific research, while the other seems to be about the generation of new breeds of dog – in fact it shares all the key features that we identify with that example.

¹⁵ Though Darwin was particularly interested in the pigeon fancy, we will mainly concentrate on the dog fancy, which has continued to grow in social and economic importance to the present day, although somewhere in the process the label “fancy” has largely disappeared from common usage – even though “Dog Fancy” is the name of a popular magazine for fanciers.

maximum possible number of surviving and reproducing offspring. All other evolutionary functionality is *subordinated* to this primary functionality. For example, individuals must survive to reproduce; hence survival takes on (subordinated) evolutionary functionality in Darwinian terms – so long as the features that ensure it do not compromise reproductive potential. Similarly, the woodpecker’s sharp strong beak makes food gathering more efficient, and so enhances the survival of a (proto-) woodpecker that has this feature; thus, a sharp strong beak has (subordinated) evolutionary functionality for the woodpecker.

For much of human history, such subordinated functionality probably provides an adequate first-order explanation for the criteria employed in artificial selection as well. People used the plants and animals they domesticated and bred as *tools* to help themselves (or the social organizations of which they were a part) to carry out functions related to reproduction or survival (of the relevant individual or social organization): ensuring an adequate food supply, for example, or defeating enemies in combat. Any feature that rendered domesticated plants and animals particularly effective with respect to such functionality could become a selection criterion for artificial selection – that is, the basis for differential treatment by their human masters that enhanced the fecundity of individuals possessing the feature, either directly (e.g. by determining which seeds were planted, which animals were allowed to mate, or which offspring were not intentionally eliminated) or indirectly (e.g. by food allocation practices).

At some point, though, in the history of human interaction with domesticated plants and animals, a different kind of selection criterion emerged. These criteria were no longer subordinated to reproduction or survival. For example, some plants were bred to enhance the beauty of their flowers. Similarly, some breeds of dog probably arose from selection criteria related to their capacities to provide pleasant companionship for their human masters: already in ancient Rome, tiny Maltese dogs were simply household pets, noted – and almost surely selected¹⁶ – for their affectionate temperaments and useless but luxurious silky coats.¹⁷

In England, in the mid-19th century, a new kind of non-subordinated selection criterion began to emerge, associated with a new kind of social activity, the dog fancy. In 1859, the same year that the first edition of *Origin* was published, the first official dog show was held in Newcastle, followed a few months later by another in Birmingham in which 80 dogs (and their human handlers) competed in 14 different classes. These competitions proved very popular, and they rapidly grew in number, as well as in the number of competitors and the size of the public who attended them. In the earliest competitions, the dogs were judged with respect to their competence in performing class-specific activities related to such subordinated functionalities as pointing, retrieving or herding. However,

¹⁶ Probably from ancestors selected for their subordinated functionality of eliminating vermin that could attack the master’s food supply.

¹⁷ A status they continued to enjoy in the Renaissance (and beyond). Carpaccio painted a beautiful diptych, in one frame of which a group of men hunt ducks in the Venice lagoon with the help of water spaniels – and in the other two women dreamily await their men’s return on the terrace of a Venetian palazzo, accompanied by a fluffy, cuddly Maltese.

this quickly began to change, and in the final decades of the 19th century, the most prestigious competitions had a completely different kind of criterion: the winners were those animals that were judged to best exemplify the “standard” conformation (physical and temperamental) of their class! Thus, the selection criteria, both for the judges in conformation competitions and for breeders seeking to produce winning animals, was not only totally unrelated to Darwinian primary functionality, but depended on *attributions* about the *attributions of others*: what judges, and fanciers in general, believed were the ideal features of a particular class, and what determined how close they believed a given animal might be to this ideal.

Such selection criteria of course required considerable alignment among the attributions of the people who participated in the competitions – owners, breeders, judges and the public who paid to view the conformance competitions – about just what counted as a class and what constituted the ideal conformation for each class. And this raises an exquisitely social question: where did these attributions come from, and how did they come to be sufficiently aligned among fanciers to make conformation competitions possible, attractive and profitable?

It is important to understand that, with just a few exceptions, what we now know as dog “breeds”¹⁸ did not pre-exist the rise of the dog fancy and the conformance competition, even if dog *breeding* did, as we have seen. As we saw, the classes in competitions were initially defined by function, and sometimes also by size, not genealogy. For example, in the earliest competitions, spaniels were divided into “springer” and “field” classes, depending on whether the animals were trained to flush game or simply locate and retrieve it. Because of the requirements of their task, springers were generally larger and more agile than field spaniels, but it was perfectly possible for a dog to compete in the springer spaniel class in the same competition in which its sibling was entered in the field spaniel class. Moreover, because of the variety in form of dogs entered in the field spaniel class, some competitions began to distinguish between larger dogs, called “field spaniels” and smaller ones, called “cocker spaniels” (supposedly because they were used to hunt smaller prey, like woodcocks). As *conformance* competitions increased in popularity, the issue of judging criteria for classes like these became particularly rancorous, since winning these competitions did not depend on what the dogs did, but how they appeared.

The solution to these questions of rules and definitions lay in *organization* – and *negotiation* channeled by organizations. Fanciers, especially breeders, established societies that debated and established procedures for determining rules for entering, classifying and judging competitions and conventions for sponsoring or recognizing competitions based upon these rules. The first such societies were based upon interest in particular classes of dogs, but soon the desire for overall coordination led to the creation of a new national society, the Kennel Club, founded in 1873, which appropriated the responsibility to oversee all “official” dog fancy competitions. The following year, the Club published its first Stud Book, which included results of past competitions and rules

¹⁸ Defined by a particular set of characteristics that “bred true” and backed up by certified breed genealogies.

and calendars for future ones. Moreover, the Kennel Club, working with the class-based societies, began to establish conformation *standards* for the classes it recognized. To handle entries in a standardized way, each individual dog was restricted to membership in one particular class. The basis of this enrolment quickly became genealogy. That is, the *classes* were transmuted into *breeds*. In 1880, the Kennel Club began to register dogs as “purebred” members of the newly standardized breeds.

All this, of course, goes exactly in the opposite direction to Darwin’s move to *unreify* species. In effect, the dog fancy societies, coordinated by the Kennel Club, reified the *breed* – and endowed each breed with an ideal organization, expressed in its published conformation standard. This reification of the breed was not a recognition of some existing underlying *natural* reality, but rather, through the activities of the clubs, the competitions, and in particular the Kennel Club’s breed registry, it *created* a new *social* reality.

This reality, although it is based upon typological thinking, is far from static. The process of determining what constitutes a breed and its associated standard was – and still is – ongoing, and it can be highly contested. The history of what is now known in the US as the *English cocker spaniel* breed (and elsewhere in the world, as simply the *cocker spaniel*) illustrates this point. Judges in early competitions in the field spaniel class favored larger dogs, and as a result most breeders selected for size. An article in a dog fancy magazine in the early 1880’s described the result: small spaniels exhibited in the last few years were just “weeds and wastrels of larger breeds,”¹⁹ and this kind of dog was well on the way to extinction. Fortunately for the many 20th century admirers of cocker spaniels, one small spaniel, Obo, enjoyed considerable success in competitions in the mid-1880’s, and his stud services became in high demand. A group of Obo-philic breeders broke away from the Spaniel Club to form their own Cocker Spaniel Club, which in 1902 drafted the first cocker standard and pushed it through the Kennel Club’s ratification procedure. With its reality acknowledged and its ideal organization described, the cocker spaniel was positioned to attain instantiation in competition-winning champions, a public attracted by these dogs’ appearance and “active, merry” disposition (as the standard proclaimed), an organized group of breeders willing to produce animals for this public to purchase, and – as a result – a future.

While standards may seem to their drafters to provide a clear vision of a breed ideal, they are in fact subject to interpretation. Breeders interpret standards through the dogs they bring to competitions. Judges award victory to the competitors who come closest to their own interpretations of what breed standards mean – which, in fact, can even change as particular individual dogs reveal previously unimagined potential to display such standard features as these, from the Kennel Club’s current cocker spaniel standard:

- “General Appearance: Merry, sturdy, sporting; well-balanced; compact ...”
- “Temperament: Gentle and affectionate, yet full of life and exuberance”

¹⁹ Quoted in Caddy, 1995.

- “Head and Skull: Square muzzle, with distinct stop set midway between tip of nose and occiput. Skull well developed, cleanly chiseled”
- “Eyes: Full, but not prominent ... with expression of intelligence and gentleness but wide awake, bright and merry; rims tight.”

And the dog-fancying and -buying public tends to seek out puppies with champion pedigrees, or at least which (the sellers assure them) will resemble the images of champions they have seen and admired. So new instantiations of the standard can change attributions about what the standard means, which can change what kind of features instantiations tend to display, and so on and on. As a result of this recurring feedback between changes in attributions and changes in instantiations, current English cockers don't look at all like Obo: for example, according to Caddy (1995), Obo was 10 inches tall and weighed 22 pounds, while the current Kennel Club cocker standards (drafted in 1992) decree a height between 15.5 and 16 inches and a weight between 28 and 32 pounds for males of the breed.²⁰

Sometimes, the gap between conflicting attributions about a breed can be so great that the only recourse is through a re-negotiation of the text of the standard itself. In the case of cocker spaniels, the Kennel Club's initial standards lasted for almost fifty years. However, during this period, some American breeders began to introduce significant changes in the configuration of their cockers. They were less interested in producing dogs well adapted for hunting than in satisfying the growing demand for household pets, and so they selected for features that appealed to the non-sporting dog-loving public: “cute”, human-like facial features and a glamorous full coat. Other breeders, committed to their attributions of a cocker as a small but powerful sporting dog, tried to resist this trend, as usual organizing societies and functional competitions devoted to pushing their attribution of what an ideal cocker should be, but the judges in the big, prestigious conformance competitions awarded “Best of breed” and even “Best of show” prizes to the newcomers, and the “American” cocker quickly became one of the most popular dogs in the US. Moreover, the American Kennel Club cocker standard was revised to favor the new type of rounder-faced, smaller, full-coated dog. The dispute among cocker fanciers and their clubs in the US about what was the “real” cocker was finally resolved through the creation of a “new” breed with its own standard, the “English cocker spaniel” – which of course was essentially the same as the Kennel Club's cocker spaniel. The Kennel Club in the meantime revised its own cocker standard to eliminate the “undesirable” American innovations, but as the great popularity of the newer version in the US began to spill over across the ocean, it too finally admitted the “American cocker spaniel” as a new authorized breed, joining its “older” cocker spaniel cousin.

To non-fanciers, all these activities – forming societies, sponsoring competitions, drafting breed standards, maintaining registries – can seem like an eccentric and socially marginal exaggeration of man's long-time relationship with the dog. But this is illusory: the “pet” phenomenon – the reinvention of the domestic animal as a companion rather than a servant – and its emergence in the past century and a half as a mass movement in Europe

²⁰ The American Kennel Club standard height for male English cockers is 16-17 inches.

and North America is a sociocultural fact with increasing political²¹ and economic consequences. Conformance competitions and breed standards lie at the heart of the modern dog industry. In the US, American Kennel Club registration can make the difference between a puppy selling for several thousand dollars – or being given away for free. The American Kennel Club itself is a large and powerful organization, which sits at the apex of a hierarchy of many societies and clubs dedicated to the purebred dog. In fact, since its inception in 1884, its members are not individual dog fanciers, but organizations. According to its 2006 Annual Report, the Club now has 594 member organizations, sanctions and regulates over 20,000 events annually (including over 1500 conformation competitions), sponsored by nearly 5,000 affiliated organizations, and it registers nearly a million puppies each year as purebred members of the 157 breeds it currently recognizes.

Some features of sociocultural innovation

We will now describe four characteristic features of sociocultural innovation that the dog fancy story illustrates. But first, we introduce several concepts that play a key role in formulating these features: *artifact*, *attribution*, and *agent*.

- By *artifact*, we mean something that human beings produce for the use of (generally other) human beings. We thus use the term in a very broad sense: a purebred puppy, for example, is for us an artifact. Artifacts may be physical, informational or performative. A breed standard, for example, is an informational artifact – as is Darwin's theory of evolutionary biology or the text of *Origin*. To be useful to others, informational artifacts generally require some form of physical or performative instantiation: a printed copy of *Origin*, for example, is a physical instantiation of Darwin's text. A conformance competition is an example of a performative artifact.
- We will use the term *attribution* to describe how people or social organizations represent to themselves the entities that inhabit their world. In particular, attributions specify the identity of social agents and artifacts: for agents, what they do and how they do it; for artifacts, how they are made and used, by whom. People, and social organizations, interact with things in the world on the basis of the attributions of identity they assign to them.
- By *agent*, we mean an organization of human beings and artifacts, in the name of which social action is initiated and executed. We defer our discussion of the concept of *organization* to the final section of this chapter.

1. The emergence of functionality: from interaction to new needs

According to a functionalist perspective, pre-existing *needs* lead social actors to participate in *interactions*, which *satisfy* the needs that induced them. In processes of sociocultural innovation, the causal arrow connecting social interaction and needs can point in the opposite direction as well. In this case, the operative causal chain has a few more essential links: from *interaction*, to new *attributions*, to new *values*, from

²¹ For example, the animal rights movement.

which new *needs* emerge. Once the new needs arise, they can induce the formation of new patterns of social interaction, through which social actors seek to satisfy them. In this way, new *functionality*, which is not subordinated to any pre-existing functionality, can come into being in the sociocultural world.

In the example, the dynamic we just described was initiated by the new activity of canine competitions. Initially, as we saw, dogs competed on the basis of their functional competences, but soon new attributions came into being, through which dogs were seen as embodying an ideal type, which came to be identified with its “breed.” These attributions were then incorporated in breed standards, which represented new values: the purebred dog was no longer merely a tool or even a valued companion for its master; rather, he was an instantiation of his breed, more valued and more valuable the better he conformed to the particular fancier’s interpretation of the dog’s breed standard. These values led to certain needs – viewing, owning, breeding particular types of purebred dogs – appropriate to a new social role, the dog fancier, needs that found expression in such activities as attending competitions, purchasing, training and showing a dog, operating a kennel, engaging with other fanciers in breed-based organizations in activities designed to enhance the breed’s status and value and in discussions about how best to interpret (and if necessary to revise) the breed standard.

The preceding paragraph thus links new activities with the emergence of new attributions, which in turn give rise to new values, and, for those individuals recruited into the new social roles opened up by activities around these values, new needs. Like most stories of sociocultural innovation, our dog fancy story begins *in medias res*, since it does not explain why canine competitions were initiated in the first place;²² and it concludes inconclusively, since it finishes with some new activities that we might expect (rightly) will lead in their turn to new attributions, new values, new needs – and new activities. In sociocultural innovation, it is often the case that *one new thing leads to another*. We will discuss the modality through which this positive feedback dynamic plays out in the next section. Suffice it here to say that the “reverse functionalism” we describe above plays a critical role in this.

2. Agents and artifacts: the reciprocity principle

Our dog fancy story is about artifact innovation: the emergence of conformance competitions, breed standards, and purebred dogs. As we saw, these artifact innovations were accomplished by new kinds of agents – in particular, local and regional breed societies and national kennel clubs. These agents structured the negotiations and generated the rules through which participants to the fancy, human and canine, were recruited, their respective roles defined, and their activities coordinated. On the other

²² Though we could go farther back, and talk about an increasingly economically and politically challenged aristocracy developing cult activities that reinforce its social status – in particular, fox hunting, the new attributions about dogs and their breeding that masters of the hunt and their acolytes developed, the adoption of these cult activities and attributions by socially ambitious members of the rising bourgeoisie and their extension beyond the circles of those who actually participated in the hunt itself; and other streams of change in attributions, values and needs.

hand, these agents' activities were also made possible, indeed constructed around, an array of artifacts: from breed registries and registered dogs to the blue ribbons and trophies that provided incentives and acknowledged success for show competitors and to the built sites, furnishings and rule books that structured agent interactions, from committee meetings to association conventions to dog shows.

We claim that this intertwining of innovations in artifacts and in the organization of the agents that make and use them is very general. We can express this claim in the form of the following ***reciprocity principle***: *the generation of new artifact types is mediated by the transformation of relationships among agents; and new artifact types mediate the transformation of relationships among agents*. In particular, the reciprocity principle implies that any causally convincing narrative about artifact innovation will constantly jump back and forth between transformations in the space of agents and transformations in the space of artifacts. The proper domain for such a narrative, and for a theory of artifact innovation, is thus neither of these: rather, it is agent-artifact space (Lane et al., 1996).²³

3. Tangled hierarchies

Around half a century ago, Herbert Simon developed a theory of organization for complex systems, based on the idea of nearly decomposable hierarchy. According to Simon, complex systems are composed of entities and processes arranged in a sequence of nested hierarchical levels. Entities are recursively structured, in that level n entities are composed of components, which were entities of lower levels. Inclusion is strict: each level $n-1$ entity can be a component of only one level n entity. Processes involve series of interactions among entities. Near decomposability implies that processes too can be localized hierarchically: they consist mainly of interactions among entities *at the same hierarchical level*. Moreover, each level is characterized by a particular spatial and temporal scale for its entities and processes respectively. In particular, this permits scientists to study (Simonian) complex systems level by level: to follow processes at level n , properties and configurations of entities at level $n+1$ can be regarded as *constants*, since the processes through which they change are too slow to matter at the characteristic time scale for level n processes;²⁴ and level $n-1$ processes are so rapid relative to this time scale that the scientist can assume they are at their equilibrium values as far as level n processes in which he is interested are concerned.²⁵

²³ It is perhaps worth reminding the reader that many current theories of artifact innovation do not respect this seemingly obvious consequence of the reciprocity principle. For example, the vast literature on S-shaped adoption curves from “innovation diffusion theory”, as well as most neoclassical economic research on “technological innovation”, regard only processes on agent space; while work in evolutionary economics around the idea of “technological trajectories”, as well as many attempts to explain such supposed empirical regularities as Moore’s and Wright’s Law, have to do just with transformations of artifact space.

²⁴ This of course does not preclude downward causation: different values of these “constants” can lead to very different outcomes for the processes under consideration, and the “constants” differ over spatial and temporal changes relative to *their* level’s characteristic spatiotemporal scales.

²⁵ Which again does not imply the absence of “upward causation”: indeed, the relevant “equilibrium” might include the very *emergence* of the level- n entities under study from interactions among lower level entities!

We have already seen how the two fundamental classes of processes in the Darwinian account of biological evolution occupy two different levels in a Simonian hierarchy: variation processes happen at the genotypic (molecular) level, while selection processes take place at the phenotypic (individual) level. As we saw, the difference between the entities and time scales for these two classes of processes is critical for the Darwinian account's success in cutting nature at the joints.

The social world is also characterized by different ontological levels. In the dog fancy story, breeders belong to breed societies and other organizations that sponsor competitions, while these societies and organizations may be members of higher-level agents like the American Kennel Club. But social organization is not in general Simonian, for three principal reasons. First, strict inclusion of entities rarely holds. For example, the same breeder may be a member of many different dog fancy organizations. Second, social processes often cannot be localized in single ontological hierarchical levels,²⁶ and as a result near decomposability fails. For example, to follow the story of the emergence of the cocker spaniel, we have to move *back and forth* between events that involve Obo and his descendants, individual breeders and judges, competitions, breed clubs, and the Kennel Club. Third, even when processes can be assigned mainly to a single hierarchical level, the correlation between hierarchical level and intrinsic temporal scale (i.e. the “larger” the entities, the slower the processes) does not necessarily hold in the social world. That is, processes involving higher-level entities need not be slow relative to processes restricted to lower-level entities.²⁷

These non-Simonian properties of social organization have important implications for social science: they undermine the strategy of studying social systems “level by level”.²⁸ In particular, they make it very unlikely that anything like the argument for the fundamental distinction between variation and selection processes that worked in the case of biological evolution – namely, that they involved interactions amongst entities at a different hierarchical level and hence with different characteristic time scales – could apply to sociocultural innovation.

4. Negotiation structured by rules structured by negotiation

A conformance competition is a set of social interactions, involving breeders, presenters, judges, dogs and the public. The interactions that constitute the competition follow a set of rules, which determine what, when and how each of the participants in the competition

²⁶ Despite the best efforts of social scientists to define them so: for example, the attempt to build a (sub)disciplinary divide between microeconomics and macroeconomics.

²⁷ Indeed, with the technological advances in communication and information processing over the last several centuries, and the capacity of higher-level organizations to exploit these technologies, social processes are now taking place at ever larger spatial scales, involving interactions among new higher-level entities, with increasingly *rapid* time scales.

²⁸ In fact, it is interesting to observe that the social sciences are organized very differently from the physical and biological sciences: the latter tend to divide their material by hierarchical level (elementary particle physics, atomic physics, condensed matter physics, chemistry; molecular biological, cellular biology, physiology, various whole organism specialties, ecology), while the first-order divisions in social sciences are functional rather than hierarchic-structural: anthropology, political science, sociology, economics.

is allowed or required to act. Some of the rules are explicit (like the movements and poses through which the presenter exhibits her dog to the judges), some not (for example, when the members of the public should applaud and when they must be silent). Without rules to channel the interactions among the participants, a conformance competition – indeed any structured social interaction event – cannot happen.

Where do the rules come from? Many of them, and all the explicit ones, are determined through processes of negotiation, which take place within the organizations that sponsor and sanction the competition. In these negotiations, members of the organization, drawing on their own prior experience from their participation in the dog fancy as well as other domains of their lives, and directed by their attributions about what a conformance competition ought to be, propose alternative interpretations and argue about their relative merits, until the relevant group reaches closure. How these negotiations proceed – who is allowed to say what, to whom, for how long, in which illocutionary mode, and how closure can be attained – of course depend on some other set of rules, which are determined by the same organizations in which they happen, based perhaps on rules followed by other organizations and encoded in such manuals as Robert's Rules of Order. And the process through which these rules are determined proceeds through negotiation, structured by rules...

We can even think of the conformance competition itself as a kind of negotiation. Each presenter is offering her interpretation of what a breed standard means – the interpretation in this case is in the form of a suitably groomed dog going through its prescribed paces. The judges carry out their evaluations, in the light of their own interpretations of the standard – which sometimes might change, usually slightly, as the competitors enact the interpretations they embody. The result of this mute negotiation is a judgment, which will have its effect on the choices the presenters and others make for future competitions, about which dogs to show and how to show them – that is, on their interpretation of what the breed standard means.

From this point of view, our dog fancy story – and sociocultural change in general – is nothing but a story of negotiations structured by rules structured by negotiations, if we are willing to consider as a negotiation process any structured confrontation among social agents over alternative attributions about the structure of the agent-artifact space they jointly inhabit. As the lady said to William James about the succession of supporting turtles that support the turtle that holds up the world on its back, it's negotiations all the way down.

We now argue that a Darwinian account is foundationally inappropriate for phenomena of sociocultural innovation with features like those we have just described:

Are variation and selection *separable and fundamental*?

Think of the establishment of the cocker spaniel breed as a sociocultural innovation. What processes can we identify that produced this innovation and accounted for the multiplication of tokens of the 'cocker spaniel' type in the UK around the beginning of the 20th century? At least these: the particular genetic combination that produced an

exemplary small spaniel, Obo, who could be perceived by some judges and breeders as representing the ideal properties of the spaniel breed as currently conceived, and who begat progeny sufficiently similar to him in appearance; the formation of an organization of breeders dedicated to breeding and showing small spaniels; the drafting of a breed standard restricted to small spaniels; the approval of this draft standard by the Kennel Club; the admission of registered dogs conforming to this standard in prestigious conformance competitions; the initiation and diffusion of the attribution that a small hunting dog makes a good household pet.

In this list, there are some elements that we could identify as variation processes: the genetic processes that give rise to an Obo, the introduction of a new standard that breeders and competition organizers can adopt as a basis for action; a new attribution of hunting dog as household pet. And some that seem like selection: judging in dog shows, choosing which kind of breed (if any) to raise, show or buy as a pet. But in both cases, the situation is very different from the case of biological evolution. The variation processes listed above happen at very different levels of organization and involve completely different mechanisms: where is the analytic bite in labeling them with the same term? Moreover, while a lot of selection is going on in this story, the criteria through which the selecting happens are changing on nearly the same time scale as the selections themselves, at least in the first and crucial stages of the establishment of the breed. The change in criteria, which depend upon new attributions of identity and functionality, themselves involve variation (the generation of new attributions) and selection (or more precisely aligning attributions among heterogeneous agents), so variation and selection are inextricably intermingled within what we have initially labeled as selection processes. The more intermingled they are, the less analytic value there is in considering them as distinct processes, since they cannot be analyzed separately – as was possible in the case of biological evolution, due to the ontological and temporal separation of the relevant entities and processes.

Worse, if we return to the list of processes contributing to the establishment of the cocker spaniel breed, several of the most important of them do not seem to be decomposable into variation and selection components, intermingled or not. For example, the formation of a new organization dedicated to promoting the (proto)breed was certainly a critical step towards the establishment of the breed. But it represents a variation only in the trivial sense that anything new is a variation, and it leads to no subsequent selection events, except again in the trivial sense that it itself didn't disappear. Rather, it is a construction, achieved through negotiations among a group of heterogeneous agents with aligned directedness (Lane and Maxfield, 1997), who were able to project the effect that such an organization might have in inducing changes in the attributions of other relevant agents about small spaniels and in carrying out activities consistent with their aligned directedness, such as drafting a breed standard and lobbying for its adoption by the Kennel Club. Moreover, the approval of this standard by the Kennel Club and the consequent establishment of a certified breed registry are organizational transformations, achieved through negotiations structured by the Club's rules – and these transformations, essential as they are to the innovation process under consideration, cannot be classified as

either variation or selection processes, again except in the most trivial and analytically useless sense of these terms.

To summarize, variation and selection processes do not seem to carve this sociocultural phenomenon at its joints. While both kinds of processes do occur, they are not distinct with respect to ontological level and time scale, and, in particular because of the endogenous generation of new attributions of functionality and hence criteria for selection, they intermingle in a way that make them analytically inextricable. Moreover, other kinds of processes, in particular organizational transformation achieved through structured negotiations, seem even more fundamental in achieving the kind of sociocultural innovation in which we are interested. Indeed, if we look carefully at our list of possible candidates for variation and selection processes in the establishment of the cocker spaniel breed, we see that almost all of them are actually brought about through underlying processes of organization transformation and structured negotiation. These conclusions seriously undermine the possibility of a Darwinian account for this kind of phenomenon.

Where is the population?

Darwin succeeded because he un-reified the species. Un-reifying higher-level entities is undoubtedly a good thing: only real historical entities should play causal roles in accounts of historical processes. But not all higher-level entities are as causally inefficacious as species end up after their un-reification by population thinking. In particular, the dog fancy is supported by a multilevel set of organizations – from breeders, to breed societies, to national Kennel Clubs – which together form a system that is itself a kind of organization, as we explain in the next section. As we have seen, processes like establishing new breeds are enacted at the *system level* and rely on negotiations and other forms of structured interactions within and among the component organizations that comprise the system. These processes may transform the structure and functionality of the organizations of which the system is composed. Thus, to regard these higher-level systems as mere aggregations that passively monitor changes in frequency distributions of their components' properties is to ignore the most salient features of the dynamics of multilevel organizational change. Indeed, in our dog fancy example – and in the examples of innovation in urban and market systems discussed later in this book – it is rarely the case that the processes we wish to understand can be localized to a single level of organization, never mind to a population of entities all inhabiting the same hierarchical level, as population thinking requires. As we argue in the next section, in these cases what we wish to study are examples of *organizational self-transformation*. When an organization transforms itself, where is the population?

4. Organization thinking

We agree with Herbert Simon that complexity science stands in need of a theory of organization. As we have already observed, Simon's proposal for such a theory, based on his idea of nearly decomposable hierarchies, turns out to be unsatisfactory to account for important aspects of human social organization. Several strands of recent complexity research offer great promise in developing a deeper theory of organization: complex

networks, modularity, degeneracy, scaling laws, and renewed approaches to hierarchy. While we cannot yet offer such a theory, it is our hope that some of the ideas presented in this book might contribute to its construction. In the meantime, in this final section of the chapter, we present some concepts and proposals that are intended to illustrate how organization thinking can provide a foundation for a theory of sociocultural innovation.

4.1 What is organization thinking? Some concepts

The worlds that scientists study consist of a flux of energy, matter and information. The flux is generated by transformations, through which the patterns constructed from energy, matter and information change. These transformations result from interactions among these patterns. We call these patterns *organizations*, and the relations among energy, matter and information that organizations construct through their interactions we call *organization*. Organization thinking seeks to understand how these patterns form and transform through interaction.

We can describe organization in terms of the relationship among three fundamental aspects: *structure*, *process* and *function*. The *structure* of an organization describes its parts (energetic, material and informational), the interaction modalities among its parts, and the modalities through which the organization interacts with other organizations. The *processes* associated with an organization describe the transformations (in organization) in which the organization may participate. The *function* of an organization provides *directedness* to its actions, through its role in determining *which* processes the organization enacts, when it is in a context in which it is possible to enact more than one process.²⁹ None of these elements are necessarily static; indeed, organizations may have processes through which they themselves transform some or all of them.

Processes are supported by structure. To participate in a process, parts of the organization must engage in a sequence of interaction events, each of which requires some particular interaction modality. Instantiating the structural support for a given process may require the activation of *management processes*, of which there are three principal types: *recruitment*, which induces (perhaps even forms) the parts that will participate in the process; *differentiation* or *specialization*, which provides these parts with the requisite interaction modalities; and *coordination*, which controls that requisite interactions happen in the right spatiotemporal order to achieve the appropriate transformation.

It is helpful to describe structure in terms of three subcategories. We will use the terms *representations*, *rules* and *relationships* to describe these subcategories, even though in some contexts some common meanings of these terms may carry inappropriate connotations. *Representations* comprise what we may call the organization's cognitive or classification system: they provide an organization with its view of the world it inhabits. *Rules* determine the organization's behavioral – that is, interactional – possibilities. *Relationships* arise from the history of the organization's interactions with other organizations; they describe how an organization is linked to these other organizations.

²⁹ Obviously, the concept of function is irrelevant for organizations that are never in such contexts – or if, when they are, chance rather than the organization determines which process is enacted. This is the case for physical systems. Function begins with biology.

4.2 What's special about human sociocultural organization?

To begin to answer this question, we use the concepts described in the previous paragraph to describe three idealized types of organizations and the worlds they inhabit. These types are meant to represent in a very simplified, even caricatured way, physical, biological and human sociocultural organizations respectively, so we call them P-, B- and S-organizations.

In a world composed of P-organizations, the rules for each organization have the form of associated fields of forces, which taken together determine, perhaps with some randomness, the kinds and outcomes of the interactions in which the organizations engage with each other. The organization of such a world emerges from the interactions so determined. P-organization processes depend on structure (positions and forces) and chance. There is no need to introduce representations, function or management processes in a description of such a world.

B-organizations actively monitor the contexts in which they find themselves. They do this through representations in the forms of *categories*, whose levels or values the organization can register. These categories are then employed in condition-action rules of the form “If(cat) then(act),” where cat describes a context in terms of categories observable by the organization, and act describes a particular interaction modality that the organization can enact.

B-organization processes consist of chains of interactions arising from rules of this form; such chains can form if each act reliably generates the cat-condition for the next rule in the chain. When various parts of the organization are responsible for different actions in such a chain, we can say that these parts communicate via “signaling” – since the transformations in context that result from one part’s act “signal” the condition that recruits the next part to make its contribution to the process enactment. There is no need for semantic interpretation here: the signal is not intended by the sender to refer to “something else”, which the receiver must infer correctly if it is to respond “correctly.” Indeed, the signaler need not have any representation that would indicate to it the existence of the receiver or the nature of the response to the signal, and vice versa.

Since certain contexts may trigger more than one rule, some of which might indicate mutually impossible interactions, B-organizations must have a way of deciding which triggered rule they will implement. This is the role of function: we can imagine function as some component part of the organization that assigns values to rules, which guide the management process that selects which eligible rule to enact. In this way, function provides directedness to the organization’s interactions.

S-organizations represent their contexts by modeling them: that is, they “populate their worlds” with entities, to which they assign attributions of identity (what kind of organization they are, what they do); moreover, they can operate in the “putative mode”, in which they use their representations (attributions of identity, plus a stock of narrative forms that express what happens when different kinds of entities interact through

particular kinds of modalities) to “simulate” real interactions, projecting changes in context that result when particular sequences of interactions among entities take place. As modelers, S-organizations employ organization thinking: their attributions of identity to S-organizations include attributions of functionality. That is, in their models, S-organizations *do* things because they *want* something. In particular, they make attributions about *their own* functionality, and they monitor whether what they do tends to produce what they want. When it doesn’t, they seek to generate new interactions that, according to what they experience in putative mode, may do better.

So far, what we have described could be read as a simplified description of what human beings do – indeed, if we except the attributions of functionality to others, what some other animals do as well. But our description is not meant to be restricted to the individual level, and our claim is that human sociocultural organizations are the only *supraindividual* organizations that are capable of sustaining the kind of representation we have described. We will justify this claim soon, but first we want to emphasize how important it is. Almost all of what we human beings have achieved, in terms of the incredible expansion of artifact space and all its attendant phenomena as we sketched in the opening paragraphs of this chapter, have been accomplished by *organizations*, not by individual human beings – whatever we might conceivably mean by an individual human being operating independently of the sociocultural organizations of which he is a part. We are claiming that the capability of human sociocultural organizations to innovate depends on the representations, rules, relationships, management processes and function associated with these organizations, which are different from, and have vastly more transformative and generative capability than, those at the individual level. Just as a single neuron may contribute to the expression and interpretation of a concept, but is not equipped to itself express or interpret one, so individual human beings contribute to the formation and enactment of organizational representations and processes, but cannot form or enact them themselves.

If representations of S-organizations do not reside in individual brains, where are they – and how are they related to interaction rules? Clearly, they are distributed among many brains – and also memory artifacts, including books, manuals and memos, and more recently computerized databases. The processes through which they are generated, modified on the basis of experience and exercised in the putative mode are many, but the most important and oldest of these is negotiation. Indeed, negotiation is the process that underlies many management processes in S-organizations. We discussed negotiation – and its recursive relationship with rules – in the previous section. Here, it will suffice to contrast it with signaling, which plays a homologous role in B-organizations. In negotiation, as opposed to signaling, semantics counts. The message that passes between sender and receiver may be completely novel in its form, and yet the sender anticipates that the receiver will be able to interpret it – indeed, that it will have the same meaning for the receiver as it had for the sender. On what is this expectation based, and how can it be right? The answer to this question is of course very complex; we describe a possible beginning point in the next chapter, based upon cognitive capabilities at the individual level and coordination requirements at the social level. Here, we note only that there is a bootstrapping involved, based upon the generative structures of language (and not only:

certain kinds of joint action and communication by other symbolic representation systems based for example on pictures or number have similar negotiation efficacy) together with past experience in which negotiation led to mutually satisfying joint action. Most important, though, is the fact that many individuals who are part of the same organization sufficiently share their attributions and narrative forms on which their messages are based that sender and receiver can be sufficiently aligned around the meaning of the messages they exchange to generate agreement on appropriate joint action.

So negotiation can generate novel possibilities for joint action, at least when representations of the negotiating parties are sufficiently aligned to develop mutually comprehensible projects – and their directedness sufficiently aligned to make the projects attractive to all of them. But negotiations can also lead to the formation of new representations, when the negotiators’ existing representations are not so closely aligned. This is due to the fact that for any human being or S-organization in general, it is very difficult to conceive of the possibility that its representation of the world is not the world! Representations change when this fact (since it must always be a fact, the world being what it is) becomes evident. Of course, occasionally the world has a way of making it evident, by producing big (and in general unpleasant) surprises. Another such occasion is when meaning breaks down during negotiations – that is, when one party to a negotiation uses an attribution or frames a narrative in a way that makes no sense in (or if it makes sense, is contradictory to) the framework of the representations of others. This sort of semantic uncertainty can lead the negotiators to “open up” attributional space that previously was closed, and explore different possibilities that transform existing representations – not necessarily, of course, leading to agreement among them.³⁰

In a supra-individual S-organization, the putative mode operates in general through negotiations (not only between people, but between people and mathematical and computational models, and among people, models, data analytic algorithms and data). Thus, as organizations explore their action possibilities in the putative mode, the generative capabilities of negotiation may result in the discovery of new interaction modalities and new entities with which to interact. As a result, rules for S-organizations necessarily have a more fluid, open-ended character than the condition-action rules that we posited for B-organizations. Indeed, for S-organizations, rules are better described in terms of *permissions*: who may (or may not: unusually, here we use the word permission in a negative as well as positive sense) interact with whom, about what, in which interaction modality. Many permissions are expressed explicitly in S-organizations, through the hierarchical command structure that is part of the structure supporting various coordination processes in many of these organizations. Many more are shared attributions within the S-organization, not explicitly stated but accepted nonetheless. Particular S-organizations (including of course individuals) may of course arrogate to themselves the permission to engage in an interaction; when they do, this permission may be contested by other S-organizations and then negotiated, with the negotiations structured by the rules of the organization responsible for coordinating action among the relevant disputants.

³⁰ See Agar (2001) and Lane and Maxfield (2005) for extended discussions of this idea.

There is of course a great deal more to be said in answer to the question raised by the title of this subsection. Some of these issues will be addressed in subsequent chapters of the book; a general answer awaits future research. Even in the stylized form of this comparison between three caricatured organizational types, it should be evident that S-organizations are constantly transforming themselves and their relationship with other S-organizations, through the negotiations in which they engage and the permissions they arrogate to themselves or grant to others. We have already seen examples in Section 3 of the fact that S-organizations also generate new attributions of functionality, for themselves and for other S-organizations. They may then attempt to realize this new functionality by means of the development of new artifacts. We now return to this theme, concluding the chapter by describing a positive feedback dynamic for innovation in agent-artifact space.

4.3 Positive feedbacks in agent-artifact space: exaptive bootstrapping

We already noted in our dog fancy story that one new thing leads to another: innovations occur in cascades, and involve transformations not only in artifact types, but in organizational forms and attributions as well. In this section, we sketch the theory of exaptive bootstrapping, which explains how such cascades happen. The theory, based on organization thinking, provides a qualitative description of a positive feedback dynamic in agent-artifact space, which we claim accounts for the explosive growth in that space that characterizes human sociocultural change, particularly over the past several centuries. The recognition of the importance of the positive feedback dynamic for artifact innovation and its implications for organizational innovation, including the growth of cities, underlies almost all the research discussed in this book and represents its principal organizing theme.

We begin by distinguishing between two different kinds of invention activities: those that are intended to deliver an existing functionality “better-faster-cheaper” than the artifacts that currently do so, and those that are designed to deliver *new* kinds of functionality. An innovation cascade can be initiated by either type of invention, and in any cascade, both types are present.

In our dog fancy story, the cascade was initiated by the first dog show in 1859, which was intended to deliver a functionality not previously provided by other performative artifacts. For an example of a cascade that began with a better-faster-cheaper invention, we recall one of the most important innovation cascades in human history, which began with the invention of printing by movable type. This was a “better-faster-cheaper” innovation: Gutenberg’s workshop figured out how to produce multiple copies of a manuscript more quickly and cheaply than was possible with the previous method (hand-copying). But almost immediately, the first printing enterprise, headed by Gutenberg’s ex-partner Fust and ex-assistant Schoeffer had to solve a series of organizational and business problems that required new attributions of functionality: for agents, who had to pay up front for the paper for over a hundred copies (soon hundreds to a thousand) of a text, before selling any of them, and needed to work out new techniques for financing, selecting, marketing and selling their products; and for artifacts – what kinds of texts to

print, and how to present them, in order to attract new customers who could not afford hand-copied manuscripts, but could pay enough for the right kind of printed book. And the solutions that the early book producers developed to these problems established new kinds of texts (and hence “reading functionalities”) that in turn induced the development of better-faster-cheaper improvements and novelties, in both the physical and informational forms of books.

Though typically innovation cascades contain both types of innovation, we claim that the positive feedback dynamic depends on the existence of the second kind – in particular, on the role of new attributions of functionality in bringing these about. The theory of exaptive bootstrapping posits the following stages for the positive feedback dynamic:

1. New artifact types are designed to achieve some particular attribution of functionality.
2. Organizational transformations are constructed to proliferate the use of tokens of the new type.
3. Novel patterns of human interaction emerge around these artifacts in use.
4. New attributions of functionality are generated – by participants or observers – to describe what the participants in these interactions are obtaining or might obtain from them.
5. New artifacts are conceived and designed to instantiate the new attributed functionality.

Since the fifth stage concludes where the first begins, we have a *bootstrapping* dynamic that can produce cascades of changes in agent-artifact space. These cascades inextricably link innovations in artifacts, in organizational structure, and in attributions about artifact and organizational functionality.

Exaptation happens between the third and the fourth stage in this process, whereby new attributions of functionality arise from observing patterns of interaction among agents and already existing artifacts. The idea here is that artifacts gain their meaning through use, and not all the possible meanings that can arise when agents begin to incorporate new artifacts in patterns of use could have been anticipated by the designers and producers of those artifacts: the combinatorial possibilities are simply too vast when a variety of different agents intent on carrying out a variety of different tasks have available a variety of different artifacts to use together with the new ones – not to mention that the designers and producers do not share the experiential base and the attribution space of all the agents that will use the artifact they produce, in ways that depend on their experience and attributions, not those of the artifact’s designers and producers! Meaning in use is one thing – the *recognition* that that meaning might represent a functional novelty is another. For this to happen, some participants in (or observers of) these patterns of interaction must come to understand that something more is being delivered – or could be delivered, with suitable modifications – to some class of agents (perhaps, but not necessarily, including themselves) other than what the participants were thinking to obtain through the interactions in which they were engaging – and which these agents might come to value. Thus, the generation of new attributions of functionality is

grounded in an *exaptation*: from the interactions between existing structures (agents and artifacts), new functionality emerges. It may then become recognized by appropriately situated and motivated agents, and (re)cognized as a new attribution of artifact functionality.

To illustrate the stages described, consider the following example from the early days of printing. In this example, stage 1 corresponds to the printed book, and stage 5 to the printed advertisement. The linking stages can be summarized as follows. Before printing, almost all manuscripts were produced in response to orders from a commissioning agent. Not surprisingly, this was initially the case also for the first printing firm, established in Mainz using the printing technology developed by Gutenberg and his co-workers, which was headed by the financier Johann Fust and the printer Peter Schoeffer (Gutenberg himself was an early example of an inventor who failed to make the transition to innovating entrepreneur). Fust and Schoeffer had one important client, the archdiocese of Mainz, which commissioned many works from them, including religious books, references in canon law, and texts for the new humanistic school curriculum in which their clerical workers were trained. Fust and Schoeffer realized early on that they could probably find purchasers for additional copies of these books. They faced the problem of how to reach these potential purchasers and convince them to buy the printed books. One organizational solution to this problem that the firm explored was to hire traveling representatives, which constituted stage 2 of the exaptive bootstrapping cycle. These representatives of course visited fairs and festivals, but they also stopped at towns along their route. When they did so, they would have to make known to potential purchasers their whereabouts and their wares – cycle stage 3. One approach that the firm took to this problem was exapted from their primary ongoing activity, in cycle stage 4: they conceived the idea of using printing, the same technology they employed to produce their wares, to enhance distribution. The new artifact type they developed (stage 5) was the printed advertisement. Their earliest surviving printed advertisement dates from 1469. It is a one page broadside, which begins as follows: “Those who wish to purchase for themselves the books listed hereafter, which have been edited with the greatest care and which are set in the same Mainz printing type as this announcement ... are invited to come to the dwelling place written in below” (Lehmann-Haupt, 1950). Thus, the advertisement attests not only to the nature of the wares (the list of books that it provided), but also to their quality (the “same Mainz printing type as this announcement”). Note that the name of the inn where the representative could be found had to be hand-written, as it changed with time and town. The printed advertisement instantiates the new attribution of functionality: the possibility of mass-circulating information about a product to recruit potential purchasers. Other instantiations of this attribution, for other classes of products, followed, and the circulation of printed catalogues soon became an important means of disseminating product information and organizing exchange activities.

Innovation cascades involve many cycles of the exaptive bootstrapping process. In addition, these cascades also include processes that are purely adaptive: given an attribution of functionality and an artifact that realizes it, apply a known technology to improve the artifact or its method of production to render it better (according to the

values associated with the given attribution of functionality), faster or cheaper. Such processes do not require the generation of new attributions of functionality. Note, though, that better-faster-cheaper invention is not necessarily purely adaptive. Many require new attributions of functionality as well: for example, Gutenberg had to exapt a variety of techniques he had learned as a jeweler in quite different contexts, even with different materials, for the new functionality of type-casting. In such cases, not only the exaptation of new attributions of functionality, but also organizational transformations like those in stage 2 are required, for example in assembling a team of agents that collectively embodies the different competences necessary to achieve a complex better-faster-cheaper invention – and in developing the procedures whereby this team can sufficiently align their directedness and then attributions about each other and the artifacts with which and towards which they work to accomplish what they have come to intend to do together.

References

- M. Agar (1994), *Language Shock: Understanding the Culture of Conversation* (Harper Collins)
- E. Beinhocker (2006), *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics* (Harvard Business School Press)
- J. Caddy (2000), *Cocker Spaniels Today: Book of the Breed* (Ringpress Books)
- W. Croft (2001), *Explaining Language Change: An Evolutionary Approach* (Longman Linguistics Library)
- A. Desmond and J. Moore (1991), *Darwin: The Life of a Tormented Evolutionist* (Warner Books)
- G. Edelman (1987), *Neural Darwinism: The Theory of Neuronal Group Selection* (Basic Books)
- J. Holland, K. Holyoak, R. Nisbet, P. Thagard (1989), *Induction: Processes of Learning, Inference and Discovery* (MIT Press)
- D. Lane, F. Malerba, R. Maxfield, L. Orsenigo (1996), Choice and action, *Journal of Evolutionary Economics*, 6, 43-76
- D. Lane and R. Maxfield (1997), Foresight, complexity and strategy, in W.B. Arthur, S. Durlauf and D. Lane, eds., *Economy as a Complex, Evolving System II* (Addison-Wesley), pp 169-198
- D. Lane and R. Maxfield (2005), Ontological uncertainty and innovation, *Journal of Evolutionary Economics*, 15, 3-50.
- H. Lehmann-Haupt (1950), Peter Schoeffer of Gernsheim and Mainz (Leo Hart)
- E. Mayr (1959), Darwin and the evolutionary theory in biology, in B. Meggers, *Evolution and Anthropology: A Centennial Appraisal* (Anthropological Society of Washington), 1-10.
- E. Mayr (1991), *One Long Argument: Charles Darwin and the Genesis of Modern Evolutionary Thought* (Harvard University Press)

