

Organizational routines and individual skills: Beyond analogy

Jack Vromen (Vromen@fwb.eur.nl)

EIPE – Erasmus University Rotterdam

First draft, November 2005

(Please do not quote without the author's permission)

Abstract

The discussion of organizational routines has been dominated by a pre-occupation with the roles that organizational routines allegedly play in the production of industry and firm behaviour. This partly explains why analogies of organizational routines with genes and skills have been invoked and relied upon. That such discussions can obfuscate (rather than illuminate) what organizational routines are and how they function is showed by critically scrutinizing the argument that just like the exercise of skills the exercise of organizational routines typically proceeds automatically (as opposed to being the result of deliberate conscious choice). This clearly runs counter to the fact that the smooth operation of routines can involve and in fact sometimes even require deliberate conscious choice at the level both of top management and of the organizational members directly and actively involved in the operation. In an ontological analysis of organizational routines, skills (and genes) also appear, but this time not as analogues but as indispensable basic constituents (or building blocks) of routines. It is argued that next to individuals and their skills, the social organization of individuals within firms and artefacts are also basic constituents of organizational routines. When invoking routines in theories that aim to explain either industry or firm behaviour, it is important to keep in mind that these are the key components of what we are talking about. Yet this should not be taken to imply that whatever it is that we want to explain (industry or firm behaviour, for example) we should always descend and refer to this micro-structure of organizational routines. It is perfectly legitimate to take 'social wholes' as the units of analysis. The work of Edwin Hutchins on distributed cognition is discussed to illustrate how teams can fruitfully be taken as the units of analysis while keeping a keen eye on what precisely is happening at the micro-level within these teams.

JEL Classification: A12, B41, B52, D21, D23

Introduction

When looking at recent attempts to clarify the notion of organizational routine (see, for example, Cohen *et al.* 1996, Dosi, Nelson and Winter 2000, Becker 2003, 2004, 2005), what is striking is that ever more properties and functions are imputed to organizational routines.¹ On a disjunctive reading of all these properties and functions the danger is that ‘organizational routine’ becomes a non-discriminating container notion (or a catch-all) covering everything that can possibly happen in firms.² On a conjunctive reading of the properties and functions the danger is that there virtually cannot be anything in firms that answers to ‘organizational routine’. Anyway, one gets the uneasy feeling that whatever it was that we started talking about when discussing organizational routines gets out of focus.

One of the reasons for ‘organizational routines’ gradually losing its contours, it is argued here, is that the discussion of organizational routines has been dominated by a pre-occupation with the roles that organizational routines allegedly play in the production of industry and firm behaviour. This partly explains why analogies of organizational routines with genes and individual skills have been invoked and relied upon. Unfortunately, the effect of invoking these analogies has been that what routines are in and by themselves, how they typically function and what their basic constituents are has been obfuscated rather than clarified. Clarifying these latter issues requires that we move beyond analogies. We should descend one level of analysis even below that of the firm and look into the interior of routines. This is what I call an ontological investigation of the micro-structure of organizational routines.³

Philosophical preliminaries

Before turning to organizational routines, a few preliminary remarks about philosophical matters are in order. What I want to bring home especially is that in discussing things like micro-foundations and methodological individualism, ontology and methodology should not be confused.⁴ It is not just that ontology and methodology address different issues. Positions taken (or endorsed) in ontology also do not necessarily prejudge positions to be taken in methodology. In particular, endorsing ontological individualism does not imply that methodological individualism should be endorsed too.

On my understanding, ontology is about what is real rather than apparent. Ontology is also about the nature of things that are real and about their basic constituents. Furthermore, identifying the causes, mechanisms and forces underlying processes taking place in the real

¹ Although I’d prefer reserving routines to the things that see to it that certain stimuli automatically invoke certain responses from *individual* organization members, as they were originally conceived of in March and Simon (1958, 140), I yield to the emerging consensus that routines pertain to the level of organizations (in what particular sense this is the case will be discussed below).

² Cf. “When writers try to proffer definitions, they often pack so much into routines and capabilities that they effectively become identical to the organization itself...” (Felin and Foss 2004, 17).

³ Thus this paper joins the chorus of those pleading for a shift in orientation from analogy to ontology (cf. Hodgson 2002). My version of ontology differs from others in that I do not take the issue of the appropriateness of Universal (or Generalized) Darwinism for studying economic evolution to be an ontological issue. See Vromen (2004) for further discussion.

⁴ In what follows I will not be talking much about micro-foundations, for the reason that ‘micro-foundations’ has an established and rather narrow meaning in economics: the program of requiring that macro-economics should be consistent with standard micro-economic theory (i.e., general equilibrium theory: economic agents follow optimal decision rules and the economy is in equilibrium; see Weintraub 1979, Janssen 1990).

world also belongs to the province of ontology. Note that the constitution or composition of really existing things is a *synchronic* issue and that the causation of things is a *diachronic* issue. It is one thing to ask about some entity what it is made of (or what are its basic constituents), it is quite another thing to ask about an entity how it is produced (or what is its causal history). Controversial issues in ontology include, for example, whether or not social wholes (supra-individual entities such as organizations and states) exist; if so, what is the mode of their existence (do they have a derived existence only, or do they have a *sui generis* existence of their own?); what are their basic constituents; do they have emergent properties (and if so, in what sense?); can they be causally efficacious? Extreme ontological individualists, we could say, maintain that unlike individuals, who really exist, social wholes do not really exist. What we call an organization, they believe, is nothing but the sum of the individuals partaking in the organization. More moderate individuals hold that social wholes do really exist, are causally efficacious, but are not agents in the same way as individual persons are. Treating social wholes as if they were agents in their own right, they believe, would be to commit the fallacy of reification or hypostasis (cf. Sober 1981).

Methodology deals with how we should go about explaining phenomena and processes. For example, what sorts of things or entities and properties are we allowed to refer to in our explanations? Methodological individualists hold that in explanations of social (aggregate) phenomena it is only acceptable to refer to individuals, their properties and their (inter)actions (Elster 1986). They hold that in such explanations it is unacceptable to refer to social wholes (or supra-individual entities) and their properties.⁵ Note that it is not specified what properties of individuals should be referred to in explanations. These could be habits, dispositions, but also brain states, for example. Most of the time, however, methodological individualists seem to have mental states such as beliefs and desires in mind. Note also that there is no presumption that different individuals have different properties. It is possible that different individuals have the same relevant properties. Finally, on this definition, methodological individualism does not speak out on the issue of where individuals get their properties from. This may be anything, ranging from inheritance to autonomous own choices.⁶

Individualists, either of the ontological or of the methodological variety, need not be atomists. That is, unlike atomists individualists may grant the existence and significance of social relations between individuals. Methodological individualists may accept implicit or explicit reference to social relations in explanations of social phenomena. Such acceptance might be based on the acknowledgement of the causal importance of social relations in the production of social phenomena. Different social phenomena might obtain if the social relations between the same set of individuals change, for example. Ontological individualists may hold that synchronically social wholes have different properties under different modes of organization of individuals. This is what Wimsatt's (1997, 2000) criterion of *non-aggregativity* is all about. Crudely put, properties of systems (or wholes) are aggregative if they do not depend on the relative arrangements of their parts (but depend solely on the properties of the parts). As Wimsatt observes, real world examples of aggregative system properties are hard to come by. Mass qualifies, for example. Most systems have non-aggregative properties. Non-aggregativity indicates that the mode of organization in a system is important, as Wimsatt

⁵ Sometimes methodological individualists accept reference to supra-individual entities, provided that the emergence of such entities can (in principle, at least) be explained in terms of individuals, their properties and their (inter)actions. Later on in the paper I will return to this.

⁶ See Whitman (2004) for an interesting discussion of other things that methodological individualism does *not* imply.

himself puts it.⁷ It is obvious that (multi-person) organizations such as firms almost always have non-aggregative properties.

Endorsing ontological individualism does not imply that methodological individualism should be endorsed too. There is no inconsistency or incoherence in denying that social wholes are unitary agents like individuals and at the same time referring to social wholes in explanations of social phenomena (without requiring that the social wholes referred to can in principle be reduced to individuals, their properties and their interactions). In fact, there are good reasons for accepting reference to social wholes and their emergent properties in explanations while being an ontological individualist. To put it differently, there are good reasons for an ontological individualist to (at least sometimes) prefer macro-accounts to micro-accounts. First, it might be convenient to refer to social wholes and their emergent properties as shorthand in cases where referring to individuals and their interactions would be clumsy and tiresome. This is a purely pragmatic reason. Another reason, which also can be called pragmatic, is that reference to social wholes and their emergent properties might suit our interest in specific counterfactual scenarios better than reference to individuals and their interactions. For example, we might be interested in the counterfactual scenario: “What would have been firm X’s performance if it had the U form rather than the M form organization (that it actually has)?” Note that a specific type of (explanatory) *pluralism* is implied here (Garfinkel 1981 and Sober 1999). One might legitimately be interested in different counterfactual scenarios and reference to different levels of analysis might suit different interests.

Yet another reason, which often goes together with the latter pragmatic reason but which is by itself less pragmatic and more principled in nature, is that social wholes and their properties might be *multiple realizable* in individuals, their properties and their interactions. The idea here is that social wholes and their properties are not uniquely realized (or instantiated or implemented) in individuals, their properties and their interactions. For each social whole and their properties there might be multiple ‘micro-structures’ (or micro-configurations) in which they are realized. Thus the relation between ‘macro-structures’ and ‘micro-structures’ might be one-to-many rather than one-to-one. Moreover, these micro-structures might be ‘wildly disjunctive’ (Fodor 1974, Sawyer 2002) so that reference to all of those micro-structures might be impracticable if not practically impossible (hence the relation with the former pragmatic reason).

Multiple realizability is often invoked as a convincing argument against reductionism. Reductionism of a macro-structure to the micro-structure in which it is realized, both of the ontological and of the epistemological/methodological kind, is only possible if there is a type – type identity between the two structures. But if there is multiple realizability, there is no such type – type identity. If there is multiple realizability, then for each macro-structure type there are many different (and, moreover, possibly wildly disjunct) micro-structure types. The absence of type – type identity does not imply that there is no token – token identity between the structures at the two adjacent levels either. On the contrary, it is perfectly possible that all possible tokens (specific instances) of social wholes and their properties are (identical to) specific instances of individuals, their properties and their interactions. It is just that another token of social wholes and their properties *of the same kind* might be identical with a token *of a completely different kind* of individuals, their properties and their interactions.

⁷ This is roughly how Wimsatt defines emergent properties: “An emergent property is – roughly – a system property which is dependent upon the mode of organization of the system’s parts.” (Wimsatt 1997, S373)

All this talk of macro-structures being realized in micro-structures might suggest that it is the micro-structure at which things are really occurring and happening and that insofar the macro-structure can be said to have an existence at all it is a derived existence at most. If this suggestion makes sense, then in fact we could argue that micro-structures (as they are understood here) also have a derived existence at most. For with micro-structures (as they are understood here) we have not reached rock-bottom. Instead of focusing on individual organisms and their properties we could descend further in the hierarchy of biological organization to organs (such as the brain), tissues and cells (such as neurons) and molecules (such as genes or DNA). And we could even further by descending to atoms and elementary subatomic particles and argue that only the latter have a non-derived existence. But this type of physicalism (or materialism) is quite controversial. Some prefer an ontologically deflationary reading in which the levels distinguished are not levels of organization, but of description or analysis. On this reading, the issues of real (or basic) versus derived existence and of ultimate constituents are sidestepped: whatever the nature and mode of their existence and their ultimate constituents, the same phenomena and processes can be analysed at different levels of description.

Wrapping up the argument so far, we can say that there do not seem to be good ontological reasons for rejecting references to supra-individual wholes and their properties in explanations of social phenomena. Ontological and methodological issues and positions should not be confused. There do not seem to be strict implications from positions taken in the one to positions taken in the other. Sometimes disallowance of reference to such social wholes in explanations is based on ontological convictions such as that it is ultimately individuals that produce, reproduce and change social wholes. But in fact no such ontological foundation can be given. Even if one is an extreme ontological individualist (denying the existence of social wholes and emergent properties) there still might be good reasons to accept reference to supra-individual entities and their properties.⁸

Nelson and Winter's analogies: "routines as skills" and on "routines as genes"

To get an idea of what I have in mind when pleading to go beyond analogy and to turn to ontology, let us reconsider some of the central elements in Nelson and Winter's discussion of routines.

Nelson and Winter (1982) develop their notion of a routine in analogy both with the skills of individuals ("routines as skills") and the genes of organisms ("routines as genes"). Routines are like skills, Nelson and Winter argue, in that routines embody organizational (tacit) *knowledge* and organizational (working) memory, just as skills embody individual tacit knowledge and individual working memory. They also argue that routines are like skills in that options are typically selected *automatically*, that is, without awareness that a choice is being made (Nelson and Winter 1982, 73). Nelson and Winter take routines to be like genes in that they both designate features (of firms and organisms, respectively) that are *persistent* (or constant) over time. Routines are like genes also in the sense that they both are *key determinants* of behaviour (again, of firms and organisms, respectively). As Nelson and Winter argue, knowledge of a firm's routines is the heart of understanding the firm's behaviour (*ibid.*, 128). In their behaviour firms are programmed by their routines in much the same way as the behaviour of individual organisms is programmed by their genes.

⁸ See also Jackson and Pettit (1992) for an interesting argument for why macro-explanations might be preferred to micro-explanations.

Note that most of the attributes mentioned are (alleged) *functions* rather than properties of routines. With the exception of the attribute of automaticity, they tell us more about what routines do than what routines are and how they function. Saying that routines serve as a firm's memory, for example, is meant to identify one of the main functions of routines in firms. But by itself this does not tell us what mechanisms working in a routine see to it that routines can serve this function. Similarly, saying that routines take care of the kind of behavioural continuity that any evolutionary theory requires does not provide us with information about what mechanisms in routines produce the required behavioural continuity. Note also that the notion of a routine plays several roles in Nelson and Winter's evolutionary theory. Nelson and Winter explicitly argue that their characterization of firm behaviour in terms of routines is instrumental to developing an evolutionary theory about industry (or market) behaviour: "As in orthodoxy, the characterization of individual firms in evolutionary theory is primarily a step toward analyzing the behaviour of industries or other large-scale units of economic organization" (*ibid.*, 18). At the same time, Nelson and Winter go at great length in arguing that their characterization of firm behaviour in terms of automatic routine-following is more realistic than the "orthodox" characterization in terms of flexible, deliberate choice. Given this last claim it is not that strange that some took Nelson and Winter's characterization of firm behaviour as a first step towards a superior new theory about firm (or organizational) rather than industry behaviour.

What this shows is that routines can be (and in fact are) invoked in different theories dealing with phenomena at different levels of organization. Nelson and Winter's evolutionary theory tries to explain industry behaviour and takes firms and their properties (notably routines) as the basic units of analysis. Organization theory, we could say, tries to explain organizational behaviour by referring to what happens inside of organizations.⁹ Thus we can say that (Nelson and Winter-type) evolutionary theory and organization theory deal with adjacent levels of organization. What the one takes as the basic units of analysis, firms and their behaviour, the other seeks to explain by descending one level of organization further below. We can understand Nelson and Winter's claim that routines are key determinants of firm behaviour as a building block of (or as a stepping stone towards) a theory of industry behaviour. But we can also understand it as a claim about the relation between the behaviour of a firm and the routines inside of it. In the latter case the claim put forward is that if one wants to understand (the nature of) firms and their behaviour, it is routines that we have to focus on. Note that this by itself still does not tell us much about routines themselves. If it is routines themselves that we want to understand, what is it that we should focus on? It seems that this necessitates descending one level of organization even further below, asking ourselves what are the crucial constituents of routines. It is this that I have in mind when talking about the ontology of routines.

Summing up, routines can (and *de facto* do) figure in theories at three adjacent levels of organization:

1. An *evolutionary theory of industry behaviour* in which not firms but their routines allegedly are the real players. What is at stake here is the relation between the level of the industry and the level of the firm. The focus is on what happens within an industry (intra-industry) as a consequence of the interactions between firms (and possibly also other parties) in the industry (inter-firm). An alleged feature of routines figuring prominently in such an evolutionary theory is their *constancy* or inertia.

⁹ This is not to deny, of course, that in trying to explain firm behaviour the external environment of firms should also be taken into account.

2. Zooming in further we have an *evolutionary theory of the firm*. What is at stake here is the relation between the level of the firm and the level of routines. The focus is on what happens within a firm (intra-firm) as a consequence of the interactions between routines (and possibly also other factors) in the firm (inter-routine). An alleged feature of routines figuring prominently in such an evolutionary theory of the firm is their *firm-specificity*.
3. Zooming in even further we have a *theory of routines*. What is at stake here is how routines function internally. The focus is on what happens within routines (intra-routine) as a consequence of the interactions between units or individuals in the firm in question (inter-unit, or inter-individual). An alleged feature of routines figuring prominently in such a theory of routines is their *automaticity*. It is often claimed that just as there is typically no conscious deliberate choice involved in the exercise of skills in individual behaviour, there is typically no conscious deliberate choice involved in the exercise of routines in firms. It is to this feature that we turn now.

Analogies leading us astray: Does the exercise of routines and skills really make for automatic option selection?

What can go wrong if analogical reasoning is not backed up by ontological analysis is nicely illustrated in Nelson and Winter's allegation that the existence and exercise of routines make for automatic firm behaviour in much the same way as the existence and exercise of skills make for automatic individual behaviour. Nelson and Winter argue that one of the key characteristics of skilled behaviour is that it typically does not involve conscious deliberation. They also argue that this characteristic carries over directly to routines: firms in routine operation characteristically do not engage in deliberate choice. This depiction more or less presupposes that it is at least an option for firms to engage in deliberate choice. But can an organization, conceived as some sort of *sui generis* behavioural entity, engage in conscious deliberation at all? Organizations as wholes cannot do so, it seems. Deliberate choice seems to be a misnomer for organizations taken as a whole. Organizations are not agents like individual persons in that have the option either to engage in conscious and deliberate choice, or in unconscious (or subconscious) automatic rule following. Assuming the opposite, assuming that is that firms are unitary agents comparable to individual human beings, amounts to what earlier was called reification or hypostasis.

It seems that it is only individual persons, such as organization members or an organization's stakeholders, that can engage in conscious deliberation. Perhaps all that Nelson and Winter mean is that in a firm in routine operation its members typically do not engage in deliberate choice? Do they mean then that attempts at strategic planning, for example, or choices made about expenditures on and investments in R&D activities, which clearly involve deliberate choice, are not part of the routine operation of firms? Textual evidence clearly suggests otherwise (see, for example, Nelson and Winter 1982, 14). There are many passages indicating that Nelson and Winter allow for deliberate choice on the part of firm members in a firm's routine operation. Indeed, some passages even indicate that continued smooth functioning of routines may require deliberate choices made by the top management. Witness their discussion of the smooth functioning of routines as targets for top management, for example, and of attempts of top management to change malfunctioning routines. Maintaining a smooth sequence of coordinated behaviour may require deliberate, flexible choice also on the part of other employees, especially if (potential) disturbances (whether these are caused externally or internally) have to be corrected (Pentland and Rueter 1994).

What Nelson and Winter seem to have in mind when they talk about automaticity of firms in routine operation turns out to be rather something else. In Dosi, Nelson and Winter (2000) the emphasis is once again on the automaticity of exercising organizational capabilities (as contrasted with the development and deployment of capabilities, which normally require conscious and intentional action). Dosi, Nelson and Winter do not assume here that firms are unitary agents in much the same way as individual persons are. What they do have in mind is that not much conscious effort is needed by top management to keep firms in routine operation. Instead there is self-coordination. It may not even be necessary that top management controlled the emergence or development of routines. The top management of a firm is likened to the head or brain in an individual person. What the mind is to an individual person, the locus of conscious control of behaviour, top management is to the firm.¹⁰ The idea is that once capabilities are in place, they can function properly and smoothly without demanding too much effort from those who are in control of the firm (just as skills can be exercised without the individuals at stake being in conscious control of them). Note that this is an altogether different matter than the issue whether or not much conscious effort is needed by those who partake in the routine operation (or exercise or execution). Note also that the emphasis is here on the coordination between units and persons rather than on the mode of operation (automatic *versus* deliberate) of persons. This coordination issue will also take centre stage later on in the ontological analysis of routines.

Thus arguing that the way firms exercise their routines is just as automatic as the way individual persons exercise their skills is misleading to say the least. Strictly speaking those exercising routines are not firms at all. If no specific action of top management is part of firms in routine operation, as Nelson and Winter argue, then it is the individuals partaking in the operation of a routine through which the routine is exercised. And the mode of operation of those individuals partaking in the operation of a routine need not be automatic at all. The smooth operation of routines may require deliberate and thoughtful repair activities. To be sure, these activities in turn require the existence and exercise of certain skills in those individuals. If they do not have the requisite skills, they cannot perform the activities that may be needed to keep a firm in routine operation. As a result, the routine may break down. This shows that individuals having certain skills is necessary and indispensable for a firm being in routine operation. Note that here the routines of a firm and the skills of individuals are linked to each other not in an analogical, but in an ontological way. Skills of individuals are conceived of here as necessary indispensable building blocks of routines.¹¹ What the foregoing discussion also shows is that, contrary to what Nelson and Winter argue, the exercise of skills in individuals is not invariably linked with an automatic mode of behaviour in individuals.¹² Skills can be exercised also if individuals engage in flexible, deliberate choice. Indeed, certain skills must be exercised if individuals engage in flexible, deliberate choice. Flexible deliberate choice simply is impossible without certain skills being exercised.

Beyond analogy: An ontology of routines

¹⁰ Note that this again of course is an analogy.

¹¹ Although Nelson and Winter put the analogical comparison between routines of firms and skills of individuals on the foreground, they start their discussion of parallels between routines and skills with the observation of this ontological link (Nelson and Winter 1982, 124).

¹² To be fair, Nelson and Winter do acknowledge that sophisticated forms of deliberate choice also require certain skills (Nelson and Winter 1982, 86). But the main thrust in their argument is that automatic skilful behaviour and flexible deliberate choice are two mutually excluding modes of operation in individuals.

What the discussion thus far provides us with is the insight that in the normal and smooth functioning of routines

- a) several individuals are involved, and that
- b) these individuals need to have particular skills.

Whether or not organizational routines are analogous to individual skills in significant respects, there is no way of denying that individual skills are indispensable building blocks of organizational routines. Ontologically speaking, that is, skills are one of the constitutive components (or constituents, in short) of routines.

But are the skills of the individuals involved all there is to routines? Clearly not. Consider Nelson and Winter's discussion of an organization in routine operation:

The overall picture of an organization in routine operation can now be drawn. A flow of messages comes into the organization from the external environment and from clocks and calendars. The organization members receiving these messages interpret them as calling for the performance of routines from their repertoires. These repertoires include ones that would be thought of as directly productive – such as unloading the truck that has arrived at the loading dock – and others of a clerical or information-processing nature – such as routing a customer's inquiry or order to the appropriate point in the organization. Either as an incidental consequence of other sorts of action or as deliberate acts of communication, the performance of routines by each organization member generates a stream of messages to others. These messages in turn are interpreted as calling for particular performances by their recipients, which generate other performances, messages, interpretations, and so on. At any given time, organization members are responding to messages originating from other members as well as from the environment; the above description of the process as starting with information input from external sources or timekeeping devices is merely an expositional convenience. There is, indeed, an internal equilibrium "circular flow" of information in an organization, but it is a flow that is continuously primed by external message sources and timekeeping devices.

(Nelson and Winter 1982, 103)

The overall image here is one of streams of messages entering the organization from the external environment and from clocks and calendars (and other timekeeping devices) and running through the organization from the one organization member to the other, where the output of the one member (the message sent) typically is the input (the message received) for the other member.¹³ This once again highlights that individual organization members are key players in the functioning of routines. But it also indicates that there also are two other basic constituents of routines: the *internal organization* of individuals in firms and *artefacts* (under the supposition that the clocks, calendars and other timekeeping devices that are mentioned above can be generalized in terms of artefacts). Let us first focus on the internal organization of individuals in firms.

The above quote from Nelson and Winter (1982) suggests that if a particular message reaches some organization member A, it is crucial for some particular routine that A sends her message to organization member B (and not to organization member C, D, E...). After having

¹³ Note that Nelson and Winter speak of the routines of individuals rather than of the organization here. Note also that deliberate choice on the part of organization members (to signal or communicate something to others, for example) is not excluded in routine operation.

processed the message, it is also crucial for member B to send his message to member M (and not to N, O, P...). And so on. One can easily imagine that if A had sent her message to C, for example, and if he processed the message and sent his output to O, that the 'final' output of the organization would have been different from what it is now. Or the final output could be the same, but it would have made a difference in the efficiency and performance in the production of the output. This seems to hold true especially if C and O belong to a different unit, specialized in conducting other tasks, than B and M, respectively. It is in this sense that division of labour and internal social organization are crucial parts of the functioning of routines. Another way to put the same idea is that routines depend crucially on how the individual organization members partaking in them are integrated or coordinated. It seems that Wimsatt's criterion of non-aggregativity is met here. For the behaviour of the comprising system (the organization, in this case) it makes a difference how the several parts or components (here the organization members) are related to each other.

What this also shows is that it might be misleading to argue that routines are recurrent interaction patterns (Becker 2004), at least if the latter are understood as regularities in behaviour displayed by several individuals (like conventions in the sense of Lewis 1969).¹⁴ For sure, routines can evolve spontaneously in much the same way as conventions do (Sugden 1986, Young 1998). But what is missing in Lewis's analysis is an explicit reference to the specific sequence in which individuals subsequently come into play. If the foregoing depiction of routines is correct, the latter is crucial in the operation of a routine. In a routine it is not just that there is regularity in the behaviour displayed by the individuals involved. What is crucial also is that the individuals are interlocked in a particular pattern. This, of course, is exactly what 'division of labour' and 'internal social organization' are about.

Note that 'division of labour' and 'internal social organization' are *supra-individual* factors in the sense that they cannot be defined by reference solely to individuals. They involve crucial reference to the comprising whole or system. One might be tempted to argue that these supra-individual factors can be explained away by referring to their origin in antecedent individual (inter)actions. After all, isn't it true that the division of labour and the internal social organization are consequences of individual (inter)actions? Maybe so, but once they have come about they make a difference to how the comprising system as a whole functions. Furthermore, once they have emerged they become part of the social environment of individuals and as such they may be causally efficacious in affecting subsequent individual (inter)actions. Thus while it may be true that division of labour and internal social organization originate in individual (inter)actions, once they have come about they can be treated as basic constituents in their own right.

The role of artefacts is not limited to providing reliable and useful informational input about present conditions to individuals (as in the case of the timekeeping devices mentioned above). As Hutchins (1995) points out, many artefacts also serve as an extension and exteriorization of the memories of individuals. What is written down in the bearing record book of a ship can rightly be seen as an important part of the memory of the ship's crew, for example. Hutchins stresses the importance of the environment, both social and natural, in boosting the cognitive powers of individuals and teams: "Humans create their cognitive powers by creating the environments in which they exercise those powers" (Hutchins 1995, 169). For Hutchins the environments essentially comprise the cognitive division of labour within teams (social) and the artefacts used. The dependence of a team's cognitive performance on properly working

¹⁴ The phrase "internal equilibrium "circular flow"" in the above quote from Nelson and Winter (1982) might suggest such a reading.

artefacts is vividly illustrated in Hutchins's discussion of the potentially dramatic consequences of the sudden loss of steam and the subsequent breakdown of electronic power in the USS *Palau*. Because of the loss of electronic power, measuring instruments such as the gyrocompass ceased to work and the navigational team had to find alternative ways to avoid dangerous and even potentially fatal collisions (for a further discussion, see below).

Pulling the threads together now, we can say that organizational routines are recurrent ways in which incoming information (influx of inputs) is handled (or processed) and transformed into certain outputs.¹⁵ This processing and transforming of information within firms characteristically involves (*intra-organizationally*) cascades (or sequences) of units (and within them individuals) as well as artefacts. For different sorts of inputs different routines may be activated, which basically means that different sorts of inputs activate different processes.¹⁶ The differences between the processes pertain to the basic constituents identified here: individuals and their skills, internal social organization (relating to the cascades or sequences) and/or artefacts. This is not to deny that routines can be represented in "If informational inputs of the sort X comes into the organization, then the organization generates output of type Y" type of rules. But doing amounts to a black box treatment of routines and obfuscates exactly what is put centre stage here: the internal working of routines and the basic constituents involved in it.¹⁷

What is the role of automatic, mindless behaviour and of conscious deliberate choice in this? To repeat, routines typically run themselves in that their continued functioning does not require conscious deliberate interventions on the part of top management. It even may be the case that the emergence of routines does not require deliberate efforts from top management. What about the organization members that are actively involved in the functioning of the routine? As argued earlier, the smooth functioning of a routine does not rule out conscious, deliberate choice on the part of organization members. Conscious, deliberate choice may even be required to keep a routine running smoothly (as in correcting mistakes made by others). Here it is important to identify what is (or what could be) the object or goal of the deliberation. Presumably the object or goal rarely if ever is to maintain the smooth operation of the routine. It is more likely that organization members simply want to do a good job in carrying out their part in the routine. To achieve this aim, organization members may consciously deliberate about what message to send to what other members. But, to repeat, the object of their conscious deliberation is not the smooth functioning of the routine as a whole, but only their part in it. Completely mindless or automatic behaviour on the part of organization members is likely to be an extreme limiting case only. On the other hand, too much open unguided search does not seem to square well with routine operation. Routines have a ring of *doing business as usual*. Unexpected changes in the environment, such as the loss of electronic power in the *Palau*, may shake organization members out of doing their share in maintaining routines and may consequently also terminate routine operation.

¹⁵ See also Vromen (2006).

¹⁶ Note that routines are individuated (or carved up) here with respect to (different sorts of) activating inputs. Alternatively, routines can also be individuated with respect to units or teams (we shall later see that Hutchins does this). Here units or teams mark the boundaries of routines.

¹⁷ This internal working is not highlighted in Cohen *et al.*'s definition of routines: "A routine is an executable *capability* for repeated performance in some *context* that has been *learned* by an organization in response to *selective pressures*." (Cohen *et al.*, 1996, 683; emphasis in the original).

Invoking routines in explanations

Now that we have identified individual organization members and their specific skills, internal social organization (including division of labour) and artefacts as the three basic constituents of routines, what follows from this for explanations that invoke routines as one of the *explanantia*? What are the methodological implications of this ontology of routines? In particular, should *micro*-explanations be always preferred over *macro*-explanations? Should we always descend to the level of how individual organization members (inter)act with each other in their social and material environment? Should we strictly follow the *dicta* of methodological individualism?

Felin and Foss provide the following reason why we should do this:

While we have no ontological problems with the existence and potential influence of collective structures on individual behaviour as such (contrary to “strong” forms of methodological individualism, e.g., Watkins, 1952), nevertheless we argue that it is inadmissible to *begin* analysis with structures and routines, because *their origin* should be of interest and the primary focus. (Felin and Foss 2004, 24)

Note first that the *methodological* conclusion that we should begin analysis with individuals and their (inter) actions (rather than with social structures and organizational routines) is based on the *ontological* observation that it is always individuals and their (inter)actions that are at the origin of social structures and organizational routines. I think this clearly is a *non sequitur*. It simply does *not* follow from the observation that social structures and organizational routines originate in individual (inter)actions that we should begin our analysis individuals and their (inter) actions. This would follow only if the origin of social structures and organizational routines were the only legitimate issue to be addressed. But of course there are other issues that can be legitimately addressed such as what happens once social structures and organizational routines have come about. Especially if it true what Felin and Foss themselves suggest, that social structures can influence individual behaviour once they have come about, why would we want to dismiss explanations in which social structures and organizational routines are referred to? The only reason I can think of is that the explanation of the origin of social structures already contains all the information we need to have to explain their subsequent influence on individual behaviour. But this would be tantamount to arguing that the ‘founding acts’ of the ‘founding fathers’ fully determine what happens in subsequent periods. It would rule out the possibility that once founded social structures can acquire ‘a life of their own’ that cannot be traced back to the ‘founding acts’ of the ‘founding fathers’. Thus it would be to deny exactly what Felin and Foss (rightly) do not want to deny, namely that social structures can have a real existence and can have real causal powers.

Furthermore, if it is the origin of social structures and organizational routines that we are interested in, we can ask ourselves whether it is the mental states (such as preferences and expectations) of individuals that we should focus on. Often mental states are the properties of individuals that proponents of methodological individualists have in mind when they argue that in explanations it is acceptable only to refer to properties of individuals. And, of course, preferences and expectations are exactly the properties of individuals that rational choice theorists refer to when explaining individual and social behaviour. But why stop with mental states of individuals? Why not descend further and refer to psychological mechanisms (underlying preferences and expectations), or to neural architecture and pathways, or to genes as different, lower-level properties of individuals? In fact, once one starts thinking about it,

stopping with mental states seems to be somewhat also in another sense. Why stop with referring to the mental states of the ‘founding fathers’ at the time that they produced the social structures and organizational routines? Why not push the explanation back in historical time and ask ourselves why the founding fathers had the mental states they had at that point in time?

Let us address the last question first. It is not uncommon for adherents to methodological individualism to accept explanation in which supra-individual entities (or ‘wholes’) and their properties are referred to, but only if (in principle, at least) an explanation of these supra-individual entities (or ‘wholes’) and their properties can be given in which only individuals, their properties and their (inter)actions are referred to. The latter required explanation supposedly deals with the origin of the supra-individual entities in question. But if an explanation of the origin of supra-individual entities and properties is demanded, why don’t we also demand an explanation of the ‘origin’ of the mental states of the individuals referred to in the explanation? Why do we rest content with treating preferences as given, for example, while we insist that supra-individual entities and properties are to be explained? Why this asymmetry? After all, the mental states also have their own history and also come from somewhere. What is more, supra-individual entities such as culture, the social milieu or environment are likely to have been causally efficacious in the genesis or genealogy of the mental states that individuals have at a certain point in time. Only extreme genetic determinists would deny this.

It may be also be somewhat arbitrary to stop explanations with the mental states of individuals in explanations in the sense that there may be properties also at lower, more basic levels of biological organization that we could refer to. We have seen that for an organizational routine to function particular individual skills are indispensable building blocks. As Nelson and Winter themselves rightly observe, individual skills in turn can be decomposed into subskills, subskills into subsubskills, and so on until we reach the level of physiology and neurology (Nelson and Winter 1982, 88). A further decomposition of the neural (cellular) level into the molecular (and especially genetic) level is also possible. In this sense the relation between routines and genes is similar to that of routines and skills. Whether or not organizational routines in firms are strictly analogical to genes in individual organisms, it cannot be denied in a very real ontological sense that genes are building blocks of routines.¹⁸

Consider the case of individual learning. Often learning is conceived of in terms of the updating or the revision of beliefs. But sometimes learning is analysed also in terms of the psychological mechanisms underlying the updating or revision of beliefs. The focus is then on the underlying competences enabling and constraining particular forms of learning. Knowing the underlying psychological mechanisms (or competences) may help us understand better why we learn some things quickly, (almost) effortlessly and reliably, why we have great trouble in learning other things.¹⁹ In studying learning we can also descend one level further below and look into the neural rewiring that learning brings about. Learning typically tends to strengthen particular synaptic connections in the brain. We can go even one level further below and look into how genes regulate the neural rewiring taking place in learning (Marcus 2004). And we can look into the causal roles that genes play in the consolidation of memory in individuals (Kandel 2001).

¹⁸ Under the supposition that “being building blocks of” is a transitive property. See Vromen (2006) for further discussion.

¹⁹ This is how evolutionary psychologists tend to understand mental modules.

So at what level of biological organization should we study individual learning? As argued in section 2, this ultimately is a matter of pragmatics. It depends in particular on the specific theoretical interests that we have. It is perfectly legitimate to be interested in different sorts of counterfactual scenarios. It is as legitimate to be interested in the question how we would behave and learn if we had other psychological mechanisms (or if other mechanisms were activated) than we actually have, for example, as it is to be interested in the question what neural rewiring would accrue if we had other genes (or if other genes were activated). But if this is true, what prevents us in our explanations from ascending to levels above the individual? What is wrong in addressing the issue, for example, how firms would behave if they had other routines than the ones they actually have? Nothing, I shall now argue, as long as such macro-accounts have a recognisable footing in micro-stories.

Hutchins's distributed cognition: why micro-accounts need not always be preferred to macro-accounts

Edwin Hutchins is the most well-known pioneer of what now is called *distributed cognition* (Hutchins 1995, Hutchins and Klausen 1996). The central premise in distributed cognition is that in studying cognition and cognitive processes it is a crucial mistake to assume that the only thing that should be looked into is what happens within the skull of skin of the individuals involved. It is seriously misleading to follow traditional cognitive science in taking the locus of knowledge to be entirely inside the individual. Where traditional cognitive science also goes wrong, Hutchins holds, is in believing that studying individuals in artificial laboratory situations is a useful starting-point for understanding human cognition. What gets out of sight by doing so, Hutchins argues, is the major contribution of the environment to the cognitive accomplishments of individuals. This contribution can only come to the fore if the cognitive activities of individuals are studied 'in the wild', that is in their natural surroundings. To be more precise, both the social aspects (having to do mostly with the division of labour between individuals and with the mode of coordination) and the cultural and material aspects (pertaining to the use of artefacts) of the environment typically play a large role in the cognitive activities of individuals and contribute greatly to what individuals jointly achieve.

Hutchins aims to show that human cognition is not just influenced by culture and society, but *is* in a very real and fundamental sense a cultural and social process. Artefacts, division of labour and coordination are crucial parts in how individuals go about learning and acquiring information. What the social and cultural (material) environments typically do is that they relieve the cognitive burden put on individuals: "The task of navigation requires internal representation of much less of the environment than traditional cognitive science would have led us to expect" (Hutchins 1995, 132). This is also why it is more appropriate in many situations and for many purposes to take a team as a whole (such as the navigation team on a ship) than the individuals in the team as the basic unit of analysis.²⁰ To this end Hutchins transplants the three types of analysis of cognitive systems introduced by David Marr (1982) from the level of individuals (as the cognitive systems to which they are almost always applied) to that of teams. This implies that first the task of the team is to be specified (task-level analysis), that subsequently cognitive processes within the team are analyzed in terms of information-processing (computational or representation or algorithmic-level analysis), after

²⁰ One of the conditions that have to be met seems to be that the conflict of interests between the team members is relatively small. In the case of a navigation team this condition seems to be met: individuals are in the same boat here in a very literal sense here. This seems to be a restriction of cognitive approaches in general (also applying to the routine and organizational capabilities approach).

which it is studied finally how these information-processing processes are implemented materially (implementational-level analysis). An important part of a navigation team on a ship, for example, is to provide reliable information about the relative (changing) location of the ship *vis-à-vis* other relevant objects and positions. In executing the task, the cognitive activities of the individual members of the navigation team obviously are of paramount importance. But what is crucial also is what special role separate individuals play in the team (division of labour), what measuring instruments are deployed (and how they are deployed) and how all this is coordinated.²¹ Indeed, as Hutchins shows in a painstakingly detailed manner, it is impossible to fully understand the cognitive activities of the separate individuals without taking these other components into account.

How to understand the relation between the team as a computational system and its individual members, who can also be analyzed as computational systems in their own right? Hutchins speaks of a *nested hierarchy* here:

When we consider a system like ship navigation, however, the situation is complicated by a nesting of computational systems. What is the implementational level for the navigation system as a whole is the computational level for the people who operate the tools of the system. The material means in which the computation is actually performed are implementational details for the system, but they set the task constraints on the performance of the navigation staff. (Hutchins 1995, 51; see also 130)

Thus the artefacts (such as gyrocompasses and other measuring instruments) that are a crucial part of the implementational level for the navigation team of the whole provide crucial inputs to the information-processing of individual team members.²² Hutchins is at pains to stress that there is nothing metaphorical in talking about crucial representations in the cognitive operation of the whole team being outside the heads of the individual team members (*ibid.*, 360). The information written down in the bearing record book is a real part of the team's memory, for example, and the erasure of lines drawn in pencil on a chart is a real case of a loss in the team's knowledge (*ibid.*, 364-365).

How crucial artefacts are in the operation of the team, is nicely illustrated in Hutchins's detailed discussion of the crisis that the ship *Palau* ran into when there was a sudden loss of steam pressure leading to an almost total breakdown in the ship's electronic power. The navigation team could not possibly stay in routine operation for the simple and obvious reason that the artefacts that they heavily relied upon stopped working as a consequence of the breakdown. This created a dangerous crisis, for without steam pressure the ship could not be steered effectively and a heavy ship as the *Palau* cannot be brought to a halt instantaneously. What subsequently happened within the navigation team is minutely described by Hutchins and analyzed by him as an instance of *organizational learning*. Without going into detail, what essentially happened is that after some hustling around the navigation team settled in a new organization and coordination of work. All this was a matter of several sequential mutual local adjustments, and not of conscious design of the member in charge, or of conscious concerted reflection on the whole process by team members. Thus no one really figured out the new solution. It just came about by individual members trying something out and

²¹ Note that these three components – individuals, internal social organization and artefacts - correspond neatly and nicely with the three basic constituents of routines identified above.

²² One of the reasons that Hutchins put forward for why individual persons need not always and inevitably be the right unit of analysis is that for some purposes individuals can also be viewed as societies of psychological mechanisms. Thus, one could say, there is in turn also a hierarchy nested within individuals.

subsequently adjusting to each other. This is why Hutchins calls the new work organization the product of adaptation rather than of design.²³

Why is this an instance of organizational rather than of individual learning according to Hutchins? It is clear that this is not because Hutchins takes the whole organization (or the whole team) to be learning as if it were an agent of its own (similar to individual persons). There are no traces of reification or hypostasis in Hutchins's account. As far as I can see, what makes the navigation team's response to the crisis an instance of organizational rather than of individual learning is due to two things. First, the direct consequences of the type of learning involved were not updated or revised beliefs of individuals (which is the case in individual learning), but adjustments in the use that individuals make of their material environment (and possibly also alterations they make in the material environment) and in their social organization. And second, as Hutchins repeatedly emphasizes, the solution that the team settled on was not found by any one particular individual team member. No one fathomed the new work organization, designed it and subsequently took measures to implement it before the new work organization actually materialized. To be sure, there were moments of conscious thinking and design in the adaptation process, but these were of a local rather than a global kind. No one consciously anticipated the new organization of the team *as a whole*.²⁴

Hutchins's discussion shows how we can make sense of organizational learning as a form of learning distinct from individual learning without imputing rather mysterious properties and capacities to organizations (or teams) as a whole. In general, we could say, Hutchins's work shows how it is possible and coherent to have one's feet firmly on the ground, to have a profound understanding of the phenomena at issue and yet to treat supra-individual entities such as teams and organizations as the basic units of analysis. Hutchins's approach might perhaps be called a top-down approach (Ross and Spurrett 2005). But it is based on meticulous micro-descriptions of practices. In fact, Hutchins's detailed reports can be said to be as micro as micro can be. In short, what Hutchins's analyses demonstrate is that there need not be something contradictory in keeping a keen eye on the micro-structure or ontology of routines and in preferring a macro-account to a micro-account.

Conclusions

If we want to find out what roles organizational routines can possibly play in theories of industry and firm behaviour, the first thing to do is to shift from analogy to ontology. We should start out with an analysis of what organizational routines are, how they (typically) function internally and what are their basic constituents (or, to put the latter differently, what is their *micro-structure*). This means that we should descend one level below that of the theory of the organization (or of the firm; in which the interior of the organization is looked into) to look into the '*interior*' of routines. An ontological analysis of organizational routines reveals that, whether or not organizational routines are analogous to individual skills in important respects, individual skills are in a very real sense indispensable basic constituents (or building blocks) of organizational routines. The internal social organization of the individuals involved in the exercise of organizational routines and the artefacts that are used

²³ Thus, whereas Nelson and Winter suggest that the development of new routines (as distinct from their exercise) typically involves deliberate choice (on the part of top management), Hutchins stresses that this need not be the case.

²⁴ Recall the earlier remarks in this paper about the importance to distinguish the local (individual) and global (social or system) level when discussing (the object of) deliberate choice.

and consulted are identified as the other two sorts of basic constituents of organizational routines.

Having thus identified the micro-structure of organizational routines in an ontological analysis of organizational routines, the next (methodological) question is what exactly should be referred to in explanations that invoke organizational routines. A methodological individualist would insist that it always is the micro-structure of routines that we should refer to. Once ontological and methodological issues are distinguished properly, however, it becomes clear that the methodological individualist stance does not follow from the ontological micro-structure of organizational routines identified here. It is perfectly legitimate to refer to supra-individual entities (such as units or teams within organizations) and their properties in explanations in which routines are invoked. What level of analysis is the appropriate one depends more on pragmatic (pertaining in particular to what sort of counterfactual scenarios we are interested in) than on ontological reasons. As especially Hutchins's analysis of the navigation team on the ship *Palau* shows, treating the team as a whole as the unit of analysis in no way need to conflict with a minute description of micro-processes. The requirement that what is ascribed to the team as a whole should be consistent with what we know about these micro-processes should not be confused with the requirement it is always the micro-structure or the micro-processes that we should refer to in our explanations.

REFERENCES

- Becker, Markus C. (2003), The concept of routines twenty years after Nelson and Winter (1982): A review of the literature, DRUID Working Paper No 03-06.
- (2004), Organizational routines: a review of the literature, *Industrial and Corporate Change* 13(4), 643-677.
- (2005), The concept of routines: some clarifications, *Cambridge Journal of Economics* 29, 249-262.
- Cohen, Michael D., Roger Burkhart, Giovanni Dosi, Massimo Egidi, Luigi Marengo, Massimo Warglien and Sidney Winter (1996), Routines and other recurring action patterns of organizations: Contemporary research interests, *Industrial and Corporate Change* 5(3), 653-699.
- Dosi, Giovanni, Richard R. Nelson and Sidney G. Winter (eds) (2000), *The Nature and Dynamics of Organizational Capabilities*, Oxford: Oxford University Press.
- Elster, Jon (1986), *Rational Choice*, Oxford: Basil Blackwell
- Felin, Teppo and Nicolai J. Foss (2004), Organizational routines: a sceptical look, DRUID Working Paper No 04-13.
- Fodor, Jerry A. (1974), Special sciences, or the disunity of science as a working hypothesis, *Synthese*, 9-115.
- Garfinkel, Alan (1981), *Forms of Explanation – rethinking the question of social theory*, New Haven: Yale University Press.
- Hodgson, Geoffrey M. (2002), Darwinism in economics: From analogy to ontology, *Journal of Evolutionary Economics* 12(2), 259-281.
- Hutchins, Edwin L. (1995), *Cognition in the Wild*, Cambridge, Mass.: MIT Press.
- Hutchins, Edwin L. and T. Klausen (1996), Distributed cognition in an airline cockpit, in Y. Engelström and D. Middleton (eds), *Cognition and Communication at Work*, Cambridge, CUP, 1996, 15-34.
- Jackson, Frank and Philip Pettit (1992), Structural explanation and social theory, in David Charles and Kathleen Lennon (eds), *Reductionism and Anti-Reductionism*, Oxford: Oxford University Press.
- Janssen, Maarten C.W. (1990), *Micro and Macro in Economics: An Inquiry into their Relation*, Groningen: Wolters-Noordhoff.
- Kandel, E.R. (2001), The molecular biology of memory storage: A dialogue between genes and synapses, *Science*, 294, 1030-1038.
- Lewis, David (1969), *Convention: A Philosophical Study*, Cambridge, Mass.: Harvard University Press.
- Marcus, G. (2004), *The Birth of the Mind: How a Tiny Number of Genes Creates the Complexities of Human Thought*, New York: Basic Books.
- Marr, David (1982), *Vision: A computational investigation into the human representation and processing of visual information*, San Francisco: Freeman.
- Pentland, Brian T. and Henry Rueter (1994), Organisational routines as grammars of action, *Administrative Sciences Quarterly* 39, 484-510.

Ross, Don and David Spurrett (2004), What to say to a skeptical metaphysician: A defense manual for cognitive and behavioural research, *Behavioral and Brain Sciences* 27, 603-647.

Sawyer, Keith (2002), Nonreductive individualism: Part I – Supervenience and wild disjunction, *Philosophy of the Social Sciences* 32(4), 537-559.

Sober, Elliott (1981), Holism, individualism and the units of selection. (Reprinted in G. Hodgson (ed.), *Economics and Biology*, Cheltenham, UK: Edward Elgar, 1995, 399-427).

---- (1999), The multiple realizability argument against reductionism, *Philosophy of Science* 66, 542-564.

Sugden, Robert (1986), *The Evolution of Rights, Cooperation and Welfare*, Oxford: Basil Blackwell.

Vromen, Jack (2004), Conjectural revisionary economic ontology: Outline of an ambitious research agenda for evolutionary economics, *Journal of Economic Methodology (Symposium: Ontological Issues in Evolutionary Economics)*, 11(2), 213-247.

---- (2006), Routines, genes and program-based behaviour (paper submitted to the *Journal of Evolutionary Economics*).

Weintraub, E. Roy (1979), *Microfoundations*, Cambridge: Cambridge University Press.

Whitman, Douglas Glen (2004), Group selection and methodological individualism: compatible and complementary, *Advances in Austrian Economics (Evolutionary psychology and economic theory)* 7, 221-249.

Wimsatt, William C. (1997), Aggregativity: Reductive heuristics for finding emergence, *Philosophy of Science (Proceedings)*, 64, S372-S384.

---- (2000), Emergence as non-aggregativity and the biases of reductionism, *Foundations of Science* 5, 269-297.

Young, Peyton (1998), *Individual Strategy and Social Structures: An Evolutionary Theory of Institutions*, Princeton: Princeton University Press.