Journal of Institutional Economics (2011), 7: 2, 175–196
© *The JOIE Foundation 2010* doi:10.1017/S1744137410000160
First published online 7 June 2010

Routines as multilevel mechanisms

JACK J. VROMEN*
EIPE, Erasmus University Rotterdam, 3000 DR, The Netherlands

Abstract: It is argued that routines can be fruitfully conceived of as multilevel mechanisms. The merits of viewing routines as multilevel mechanisms are that it helps in putting together a coherent picture of what routines are, what routines do, and how they do it. In particular, it helps in getting a clearer picture of how skills and routines are ontologically (rather than metaphorically) related to each other. It allows us to see that while routines are *generative* mechanisms producing recurrent patterns of firm behavior, as *multilevel* mechanisms they themselves are at the same time recurrent patterns of interaction within firms. Because of its 'behavioral' spirit, viewing routines as multilevel mechanisms (rather than as, for example, unobservable dispositions of firms to energize patterns of behavior in firms) greatly facilitates further empirical research on crucial, as yet unresolved issues, such as how stable and robust routines are and to what extent firm behavior is routine.

1. Introduction

Despite several serious and interesting attempts to clarify the notion of routine (cf. Cohen *et al*., 1996; Becker, 2004), a widely agreed upon definition and understanding are still lacking (cf. Becker *et al*., 2005: 784). Ambiguities that were already present in Nelson and Winter's (1982) seminal discussion are still not fully removed. Part of the reason for this unsatisfactory situation is that the notion of routine was meant to serve multiple purposes. So much was packed into the notion right from the start that it is a small wonder that it has not collapsed under the weight it had to carry. What also did not help is that Nelson and Winter introduced the notion of routine by way of a couple of metaphors. With hindsight, we can say that quite a few scholars working on routines got carried away by these metaphors and were led into unproductive blind alleys. As a consequence, discussions of routines often proceeded in very general and abstract terms. What have been lacking mostly are detailed in-depth and concrete analyses of what routines are and how they (typically) function. Ontological analyses of routines are meant to fill part of this lacuna.

In Vromen (2006), I made a first attempt in this direction. What a shift from a metaphorical to an ontological analysis might imply can perhaps be illustrated

*Email: vromen@fwb.eur.nl

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

most clearly by looking at the relation between routines and skills. As will be discussed in more detail below, Nelson and Winter invoked the 'routines as skills' metaphor mainly to stress that routine firm behavior resembles skilful individual behavior in its 'automaticity'. Some have accepted this metaphor, while others opposed this by arguing that a lot of firm behavior is not automatic. It is not clear who is right. Part of the problem is that it is not even clear what it means to say that firms behave automatically. What is clear and uncontested, though, is that routine firm behavior crucially involves skilful individual behavior. Without skilful individual behavior, routine firm behavior would be impossible. Here routine firm behavior and skilful individual behavior are linked not metaphorically, but ontologically. This seems to be a more promising starting point for further studies of what routines are and how they typically function than Nelson and Winter's original 'routines as skills' metaphor. Yet it is also important to stress that this ontological link provides no more than just a starting point for further study. Much more effort is needed to find out what routines are and how they function than just making the observation that skilful individual behavior is involved in routine firm behavior. Precisely how skilful individual behavior is involved in routine firm behavior is not clarified by just this observation, for example.

This paper seeks to further clarify such ontological issues. It is argued that viewing routines as multilevel mechanisms helps in doing this. It is furthermore argued that viewing routines as multilevel mechanisms also helps in seeing that two seemingly incompatible understandings of routines are in fact compatible. One of the ambiguities still surrounding the notion of routine is that routines are sometimes seen as underlying generative structures, but also sometimes as behavioral patterns. It seems routines cannot be both: something cannot possibly be both a cause and its effect. As we shall see, however, once we view routines as multilevel mechanisms, structures that underlie and generate patterns in firm behavior appear as behavioral patterns within firms in their own right. Some shy away from such a behavioral understanding of routines, and choose to define 'routines' instead as dispositions that energize behavioral patterns, on the ground that routines do not get lost if they are not exercised for a while. It will be argued that this argument does not establish what it purports to establish and, furthermore, that there are no other good reasons to postulate a disposition as an extra unobservable entity on top of the behavioral patterns that they allegedly energize. One of the main merits of the behavioral understanding of 'routines' defended here is that it makes doing more empirical research on routines easier. More empirical research on routines is badly needed and ontological analyses cannot possibly obviate that need. But ontological analyses can help in elucidating not only what are the remaining issues on which more empirical research is needed, but also what should be focused on in attempts to resolve the issues.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

2. 'Routines': from metaphors to ontology

Nelson and Winter (1982) introduced the notion of a routine by way of two metaphors: 'routines as genes' and 'routines as skills'. These two metaphors continue to influence current debates about what routines are, how they function and what theoretical roles they can or should play in evolutionary theories of economic change.

Nelson and Winter (1982) first discuss individual skills before they start discussing organizational routines. Their discussion of skills of individuals and of skilful individual behavior is meant to pave the ground for their discussion of routines of firms and of routine firm behavior. Nelson and Winter argue that routines are for firms what skills are for individual persons: they are programmatic, they embody tacit knowledge, and, above all, option selection in them is automatic in the sense that no deliberate choice is involved (Nelson and Winter, 1982: 73). Routines are non-rational and self-actuating (cf. Becker, 2005). Here we hit on what might be called *the* root 'common sense' meaning of routine: conscious deliberation about how to do something is not needed and *de facto* absent in routine behavior. A novice automobile driver (or billiard player) has no other option than to consciously deliberate about any single move. But once he is experienced, the movements are routinized (or habituated) and conscious deliberation about every single move is no longer needed. Note that routines are taken here to be properties of firms, not of individuals. Nelson and Winter might occasionally have used 'routine' to denote a particular sort of individual behavior, but in the meantime it seems that everyone agrees that routines are collective, not individual phenomena.

It is not immediately clear, however, what it means to say that routine firm behavior is automatic. In individual behavior, 'automatic behavior' is meant to denote the absence of deliberate choice. But firms do not seem to be the sort of 'agents' that are capable of making deliberate choices. Only individual persons acting in or on behalf of firms seem to be capable of doing that. Does 'routine firm behavior' mean then that individual organization members that are involved in the functioning of routines behave on automatic pilot? In some of the best literature on the subject, it is argued that it is the procedural memory of individuals that is a (and perhaps even *the*) crucial component in the functioning of routines (Cohen and Bacdayan, 1994). In one of the best empirical studies to date (Pentland and Rueter, 1994), however, it is emphatically denied that the functioning of routines involve thoughtless or mindless automatic behavior on the part of the individuals involved. It is stressed here that the functioning of routines typically and perhaps even essentially involve thoughtful deliberations on the part of individuals. In their seminal discussion of routines, Nelson and Winter (1982) likewise do not assume that individuals in firms that are involved in the operation in routines behave on automatic pilot.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

We should therefore not assume that individual organization members involved in routine firm behavior do not engage in deliberate choice. What then could it mean to say 'routine firm behavior' is automatic? The most plausible answer, I submit, is that top management of the firm is typically neither installing nor policing routines in firms (Vromen, 2006). Routines typically evolve spontaneously, that is without conscious design by top management, and once they emerged they tend to be self-sustaining, that is external enforcement by top management is typically not needed.[1] Nelson and Winter liken top management of a firm to the prefrontal cortex of an individual person (Nelson and Winter, 1982: 125). Just as the prefrontal cortex enables individual persons to consciously plan their actions and lives, top management is capable of conscious planning and deliberate monitoring in firms. To say that routine firm behavior is automatic is to say that this capability typically is not exercised in routine firm behavior.[2]

Routines tend to be stable, if not inert over time. This is what Nelson and Winter (1982) primarily wanted to convey with their 'routines as genes' metaphor. Nelson and Winter were looking for entities in firms that were stable enough over time to serve as units of selection in an economic evolutionary theory of change. They argue that the routines of firms fill the bill. Like genes in biological evolution, routines are taken to be long-lasting entities (meeting Dawkins's 1976 requirement of longevity). Nelson and Winter stress the rigidity and even inertia of routines. Routines cannot be changed overnight, even if there is a strong desire, on the part of top management for example, to do so.

More recently, Hodgson and Knudsen (2004) argued that routines are genotypes rather than phenotypes of firms and that routines can be seen as replicators, suggesting that routines resemble genes not only with respect to longevity, but also with respect to fidelity in transmission.[3] Although this is contested, Hodgson and Knudsen might have a point here. But it seems Nelson and Winter themselves never thought it worthwhile to engage in a discussion of whether routines are genotypes or phenotypes or in a discussion of whether or not routines are replicators. They rather stress that routines tend to be firm-specific; routines are difficult to imitate by virtue of the fact that they embody inarticulate tacit knowledge. Attempts to imitate routines of other firms are therefore likely to lead to 'mutations'. Nelson (2007) is adamant in arguing that the respect in which routines resemble genes is stability over time (cf. Nelson and Winter, 1982: 135), not faithfulness in copying in inter-firm transmission.

1 This is related to Nelson and Winter's view on routines as truces.

2 This is not to say that installing routines or maintaining them cannot be targets for managerial effort. But Nelson and Winter very much stress that routines tend to emerge and tend to be maintained without managerial effort.

3 Note, though, that Hodgson and Knudsen recognize that the fidelity in copying is considerably lower in the transmission of routines than in the transmission of genes.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

The two features of routines discussed thus far are clearly related. The (alleged) fact that routines tend to be self-sustaining at least partly accounts for the (alleged) fact that routines tend to be stable over time. A third feature that Nelson and Winter attribute to routines, namely that routines are key determinants of firm behavior and firm performance, is less clearly related to these two features. It is perfectly possible that, despite their tenacity, routines have a smaller impact on firm behavior and performance than other factors (such as the policy planned by top management, for example). Nelson and Winter do claim, however, that understanding a firm's routines provides the key to understanding its functioning and performance (Nelson and Winter, 1982: 128), or at least its systematic parts: 'most of what is *regular* and *predictable* about business behavior is plausibly subsumed under the heading of "routine"' (ibid.: 15). Differences in behavior and performance between different firms can be explained for the most part in terms of the firms' specific routines. A firm's specific routines determine how standard inputs (or tasks, or stimuli) to the firm are transformed into regular output patterns. In each firm, there might be different routines for different sorts of inputs. According to Nelson and Winter, there are not only firm-specific first-order routines in a firm determining how things are typically handled in and by the firm, but also firm-specific second-order routines determining how the firm searches for better (first-order) routines if the old first-order routines cease to function satisfactorily. To be sure, there is also an irregular and unpredictable part of firm behavior, but in theorizing about firm behavior that part can only be accounted for in terms of stochastic elements.

Nelson and Winter argue that routines not only function as truces in firms but also as memories. Routines somehow store the information and knowledge that a firm needs to function effectively and reliably in the face of recurrent tasks. Although there are some occasional references to 'organizational learning' and to the (alleged) fact that the memory of an organization cannot fully be reduced to what is stored as memories in the heads of individual participants, the 'routines as memories' idea is not elaborated upon in great detail and depth. The same arguably holds for Nelson and Winter's 'routines as truces' idea. That idea is not worked out in any detail and depth either. This is unfortunate, for it is left unclear what exactly is going on inside firms when they are in routine operation.

Routines figure at several levels of analysis in Nelson and Winter's evolutionary theory. As several commentators have observed, Nelson and Winter were primarily interested in developing an evolutionary economic theory at the aggregate level of industries and whole economies in which their 'theory of the firm' is just one building block (Nelson and Winter, 1982: 18). Their evolutionary economic theory was meant to be a competitor of 'orthodox' economics (cf. Vromen, 1995). As Nelson and Winter point out, one of the crucial assumptions in 'orthodox' economics is that individual agents (firms included) are engaged in flexible conscious choice. Nelson and Winter's assumption that routines provide the key to understanding firm behavior is developed in opposition to

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

this 'orthodox' assumption. But 'orthodox' economics and Nelson and Winter's evolutionary theory situate the phenomena to be explained at the same level and they both share a technological (rather than organizational) perspective on firms and their behavior.[4] Similar to production functions in 'orthodox' economics, routines appear as decision-rules of firms in Nelson and Winter's formal models that transform factor inputs into firm outputs: quantities of goods and services produced.

Nelson and Winter's 'theory of the firm' itself focuses on phenomena one level below. The theory aims at explaining the differences in behavior and performance between different firms by referring to their specific routines. It directs our attention to the specific routines that particular firms have; to when and under what conditions, each of them is activated; and how their activation results in firm behavior. Ironically, although Nelson and Winter meant their discussion of routines to be a building block in an economic theory of industry behavior, it had far greater impact on theorizing in the field of strategic management.

Arguably, working out the 'routines as memories' and 'routines as truces' ideas in some depth and detail would necessitate descending one more level. It would prompt us to look inside routines at interactions between units or teams in firms and between individual firm members in them. Even those who insist that 'organizational memory' and 'organizational learning' cannot be fully reduced to the memories and learning efforts of individual firm members, they do not deny that individual firm members are crucially involved in the storage of knowledge in firms and in the functioning of routines in general. Here the discussion shifts from metaphor to ontology. At stake is the issue of what are the basic constituents of routines.[5] No longer is the issue whether or not organizational routines are like individual skills in relevant respects, for example. The issue rather becomes what more there is to organizational routines than individual skills. No one doubts that individual skills are indispensable elements in the operation of routines. The bone of contention is rather whether individual skills are the only indispensable elements in the operation of routines.

In Vromen (2006), I characterized routines succinctly as 'recurrent *intra*-organizational, *multi*-person *inter*action patterns displayed in specific artificially created physical environments' (ibid.: 551). What I aimed to bring out here is not only that several persons (organization members) with their respective skills are crucially involved in the functioning of routines, but also that at least three further features are also crucial for routines. First, there is the specific way in some particular routine in which the persons are connected and arranged in interaction patterns. In other words, the social organization of the members is crucial. Second, specific artifacts often are crucial too. As especially Hutchins

4 Firms are essentially depicted as production functions in 'orthodox' neoclassical economics.

5 This is not the only possible meaning of 'ontology'. In Vromen (2004), I discuss three different clusters of issues that can all be called 'ontological'.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

(1995) makes vividly clear, units or teams of individuals boost their cognitive powers by creating a suitable physical environment. Third, routines typically are context-dependent in that they are triggered (or activated) only by specific sorts of inputs to the firm. Each routine is typically activated only in a specific environment.

Although this already gives some idea of the basic constituents of routines and of how they interact when routines function smoothly, still a lot is unclear about how exactly the constituents interact in routines. What is still unclear in particular is how different levels of analysis are related when routines function smoothly. In Vromen (2006), I argued that since genes (and their expressions) are somehow involved in the exercise of skills in individual behavior, genes can also be called basic constituents of routines. I also argued that organization, the way in which genes are organized in cascades, and context-dependence are important also at the genetic level. What this suggests is that roughly the same structures – the context-dependent activation of organized entities – appear at several successive levels of analysis. This suggestion was not worked out in the paper, however. I aim to make good on that now. I argue that it is helpful to conceive of routines as multilevel mechanisms.

3. Routines as multilevel mechanisms

The first thing to appreciate is that our understanding of how certain phenomena come about is greatly enhanced if we know the mechanism (or mechanisms) that produced them. We know that we can start the engine of our car by switching a key or pushing a button. But we understand a great deal more if, in addition to this, we know what causal process is set in motion by our switching of the key. Similarly, via all kinds of clever experiments we might come to know that a particular medicine is effective in curing certain diseases. But we understand a great deal more if in addition to this we know how intake of the medicine affects bodily processes, say at the molecular level. This seems to be well appreciated in science. Craver (2007) discusses the celebrated example of Hodgkin and Huxley's model of the action potential (the so-called HH model) in neuroscience. It is generally acknowledged that with their equations, Hodgkin and Huxley made a significant contribution to neuroscience. Numerous experiments confirmed that the HH-equations accurately describe how specific ion conductances (which are relevant for the electronic potential differences across the membranes of neurons) change with voltage and time. Despite their enormous success, however, Hodgkin and Huxley themselves never regarded their equations as genuinely explanatory. Instead, they looked upon their model as a merely phenomenological model that failed to identify the molecular mechanisms that causally link changes in voltage to changes in ion conductances. Without an account of the underlying mechanism, the HH-model only tells us how the changes correlate. The model itself does not explain why ion conductances change, when they change.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

Scientists often believe they understand a phenomenon thoroughly only once they are able to give a mechanistic explanation of it. But what exactly are mechanisms? From the foregoing introductory discussion, we can already conclude that mechanisms pertain to the processes that causally connect certain triggers, like the switching of the key or changes in voltage, to their eventual effects, the starting of the car's engine and changes in ion conductances respectively. A natural way of putting this is to say that the triggers are the causes that set into motion a particular mechanism that in turn produces the phenomenon to be explained. But this leaves a lot wanting. It does not tell us what is this 'thing' that does the causal connecting and how it does it.

One of the most outspoken proponents of mechanistic explanation in the social sciences is Jon Elster (1989, 1998, 2007). Elster has consistently argued that a thorough causal explanation of a social phenomenon does not only specify the phenomenon's cause (or causes), but also the mechanism that connects the cause to the phenomenon. Elster (2007) argues that there are many possible candidates for such social mechanisms, rational choice being only one of them. One looks in vain for a general analysis of what mechanisms are and how they connect causes to their effects in Elster (2007), however. What is lacking in particular is a clear general view on what can be called the internal constitution of mechanisms. What are the parts, or components, that mechanisms are composed of, how are they typically related to each other and how do they work together in transforming particular sorts of inputs into particular sorts of outputs? Can light be shed on how mechanisms are constituted and how they function *in general*?

Fortunately, in recent years a lot of progress has been made in the philosophical literature in clarifying what mechanisms are and how they function.[6] Although no unanimous agreement has been reached about a precise definition of 'mechanism', there are several things that all scholars in the field seem to agree on. The points on which agreement has been reached are substantive enough to get a clear picture of what is involved in the functioning of mechanisms. There seems to be consensus that mechanisms consist of parts (or components) that are organized in such a way that their activities and interactions 'are productive of regular changes from start or set-up to finish or terminal conditions' (Machamer *et al.*, 2000: 3). Thus, what are important in mechanisms are not just the parts they are composed of, but also the specific ways in which the parts are organized or arranged. In this sense, mechanisms as 'wholes' are typically more than the sum of their parts.

Giving a mechanistic explanation implies that what happens in between a system's start or set-up conditions and its finish or terminal conditions is not treated as a black box, but is highlighted. The way in which the mechanism's

6 The works of Bechtel (see, e.g., Bechtel and Abrahamsen, 2005), Craver (see, e.g., Craver, 2007), Darden (see, e.g., Darden, 2001), Glennan (see, e.g., Glennan, 1996) and Machamer (see, e.g., Machamer *et al.*, 2000) can be mentioned here.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

parts are organized and behave and interact is peered into. As Craver puts it, 'mechanistic explanation is inherently componential' (Craver, 2007: 131). In this sense, mechanistic explanations have an inbuilt reductionist bias. In mechanistic explanation, one is looking not just at what the system is doing, but also 'into the system', at how it is doing, what it is doing. Thus, giving a mechanistic explanation implies that a system is decomposed into its constituent parts (see also Simon, 1969 and Cummins, 1975 and 2000 for similar views). Often the parts of a mechanism can themselves be seen as mechanisms with parts. Thus, systems are often nested and mechanisms are often multilevel. If a mechanism is multilevel, it can be analyzed at different levels.[7]

Mechanistic explanation always involves the decomposition of a system into its constituent parts. But the level at which a system is to be decomposed into is not thereby pre-specified. In multilevel mechanisms, there are several options available. Multilevel social mechanisms in which humans are involved need not necessarily be analyzed in terms of individual persons and their (inter)actions. They might be analyzed instead in terms of supra-individual units (such as teams and groups of individuals) or in terms of intra-individual entities (such as the individual's neurons). Pragmatic considerations, pertaining to the specific things we might be interested in, legitimately play a crucial role in determining the level to which we want to decompose the system. A relevant consideration is also how stable or invariant over time the components at various levels are. Components at lower levels need not be more stable than those at higher levels. Thus, in a multilevel social mechanism, teams of individuals might be more stable than the individuals in the teams, in the sense that the functioning of the team remains more or less the same, even though some members of the team are replaced by others. In such a case, an analysis of the system in terms of teams rather than in terms of the individuals comprising the teams might be preferred.[8] It can be added that mechanisms themselves also must display some stability or invariance over time. It does not make sense to speak of a mechanism, it seems, unless it has some durability. We will not be inclined to call the action of a configuration of parts a mechanism if the configuration acts only once.

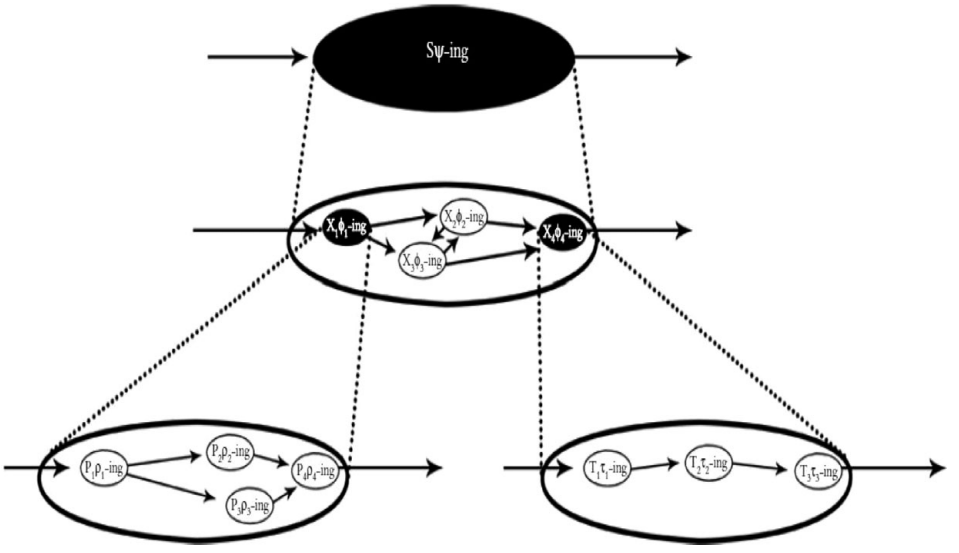
In mechanisms, the relation between the levels is always of a particular kind: it is always that of a whole and its constitutive parts. As Craver makes clear, there are in principle many different ways in which levels can be related to each other. One such relation is realization. Some argue that mental properties are realized in physical properties, for example. The relation between levels in

7 This is not to be confused with a notion such as multilevel selection. The key idea in 'multilevel selection' is that there *several* selection processes simultaneously going on at different levels (say the individual and group level), whereas the key idea in 'multilevel mechanism' is that *one and the same* mechanism (or causal process) spans several levels (and can thus be analyzed at different levels).

8 Thus, contra Abell *et al*. (2008) (see also Felin and Foss, 2005, 2006; Felin and Hesterly, 2007), Vromen (2009) argues that the notion of mechanistic explanation does not support an unqualified preference for individualist ('micro') explanations over collectivist ('macro') explanations.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

Figure 1. Multilevel mechanism (Craver, 2007: 189, Figure 5.7)



mechanisms is not one of realization, however, but a part-whole one. Indeed, in mechanistic explanation 'levels of mechanisms are defined componentially within a hierarchically organized mechanism' (Craver, 2007: 191). The constitutive relation between levels as envisioned in mechanistic explanation is not a causal one, at least if we understand a causal relation in a common, standard way. On a common, standard interpretation of a causal relation, causes and effects are logically independent. Wholes and their component parts are clearly logically dependent. Moreover, causation is standardly taken to be diachronical (i.e., causal processes take time), whereas 'constitution' or 'composition' are synchronical. At any given point in time, as in a snapshot, a mechanism has interacting component parts, which become visible if we zoom in on the mechanism. As Craver and Bechtel (2007) point out, all this implies that there cannot be interlevel causation in mechanisms. Thus, in mechanisms we have neither upward nor downward causation. As the starting of the engine is an integral crucial part of the starting of the car, it would be misleading (if not downright false) to say that the starting of the car 'downwardly' causes the car's engine to start functioning, for example. The starting of the car's engine *is* to a large part the starting of the car. The starting of the car's engine is what we see if we zoom in on a starting car.

We can visualize the working of a multilevel mechanism as shown in Figure 1.

At the top level, the mechanism (or system S) 'as a whole' is depicted. The mechanism reliably and predictably links inputs (denoted by the arrow on the left) to outputs (denoted by the arrow on the right). The middle level is what we

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

see if we zoom in on the functioning of the mechanism. What becomes visible then are particular organization and interaction patterns in which entities at a lower level Xs matter. If the mechanism is truly a multilevel mechanism, we could zoom in further on the activities of the entities at the middle level. Then patterns of organization and interaction of entities (the Ps and Ts) at an even lower level become visible. What the figure nicely brings out is that the levels involved in a multilevel mechanism are not causally, but constitutively or componentially related to each other. The same causal process can be looked into at different levels. The figure also shows that by zooming in on mechanisms we can choose to have a closer look at parts of the mechanism. We can study parts or chunks of the causal chain in a more detailed way. In this sense, lower-level analysis is finer-grained than higher-level analysis.

In sum, multilevel mechanisms crucially display the following features:

- They display some stability over time.
- At the level at which the mechanism is identified, particular sorts of inputs are transformed regularly, reliably, and predictably into outputs of system.[9]
- They have parts (or components) that are organized in particular ways.
- The actions (or activities) of components themselves can be seen as mechanisms in their own right.

Routines, as I discussed them in Section 2, do seem to display these features.

- What Nelson and Winter primarily wanted to convey with their 'routines as genes' analogy, we saw, is that routines tend to be stable and persistent over time. Indeed, they even argued that routines tend to be robust, in the sense that they tend to survive personnel turnover.
- We saw that in standard situations, in which 'business (in and for a firm) is as usual', routines regularly, reliably and predictably transform particular sorts of incoming messages to the firm into the firm's behavior and performance.
- As multi-person phenomena, routines have parts that are organized in particular ways. These parts can be either units, teams, or groups of individual firm members. Since units, teams, and groups in firms consist of individual firm members, individual firm members and their skills are component parts of routines at any rate.
- Regardless of whether individual firm members or units are seen as the proper parts of routines, the parts of routines can be further decomposed into (sub-) parts. The activities of individual firm members can be seen as the functionings of mechanisms that crucially involve the exercise of the members' skills as component parts, for example. The member's skills can themselves in turn also be understood as mechanisms with 'subskills' (Nelson and Winter, 1982: 88) as their component parts and so on.

9 From a computational point of view, mechanisms can be seen as algorithms in that they produce more or less guaranteed results when presented with particular sorts of inputs.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

4. Routines are generative structures *and* behavioral patterns

One of the merits of viewing routines as multilevel mechanisms is that it helps us see better how different and perhaps even seemingly conflicting conceptions of routines can be combined in a coherent way. In what follows, let us keep in mind that routines are collective phenomena. Only multi-agent collectives such as groups and organizations can have routines. Individual agents (or persons) can have (and actually do have) skills and habits, but they cannot have routines. 'Routines' are reserved exclusively to denote possible properties of collectives.

In the literature, we can find the following two characterizations of routines:

1. Routines are structures (or capabilities, or capacities) underlying and generating behavioral patterns.
2. Routines are specific recurrent behavioral patterns.

For each of these characterizations, it is easy to find opponents and proponents in the literature. It is easy to see why proponents of the first characterization tend to be opponents of the second characterization and *vice versa*: either routines are structures generating behavioral patterns, or routines are themselves behavioral patterns. It seems one cannot have it both ways. If routines are behavioral patterns, it seems they cannot be at the same time structures generating behavioral patterns. Yet viewing routines as multilevel mechanisms help us see how these seemingly contradictory characterizations can be reconciled in a coherent understanding of routines. It allows us to see how routines can be generative structures and behavioral patterns at one at the same time.

The first characterization, that routines are structures underlying and generating behavioral patterns of firms, can be related to Nelson and Winter's argument that routines are like the memories of firms. A firm's routines are repositories of standard ways of handling things in and by the firm, which are stored somewhere in the firm. Each of the routines is activated (or triggered, or energized) in a particular context by specific sorts of stimuli or inputs. In this conception, routines are what leads to firm behavior rather than firm behavior itself. Firm behavior is what results when (some of the) firm's routines are expressed (cf. Cohen *et al*., 1996: 673). In this vein, 'routines' are understood as executable capabilities for repeated performances in particular contexts (Cohen *et al*., 1996: 683).

There are also passages in Nelson and Winter (1982), however, suggesting that 'routines' refer to patterns of behavior rather than capabilities underlying such patterns. In Nelson and Winter's depiction of a firm in routine operation (ibid.: 103), for example, 'routine' seems to stand for specific sequences of interlocked activities and practices within firms. What is emphasized here is that messages that flow into the organization are processed by particular organization members, which pass the results from these processings on to particular other organization

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

members, who in turn process these results … and so on, until in the end something productive is accomplished. Different patterns of behavior within a firm might be triggered by different stimuli (or inputs). Thus, an order that reaches a firm will be processed routinely in a different way in the firm than a customer's complaint about the quality of goods delivered. In this vein, 'routines' are understood as 'multi-actor, interlocking, reciprocally triggered sequences of actions' (Cohen and Bacdayan, 1994: 554), or as 'repetitive, recognizable patterns of interdependent actions, carried out by multiple agents' (Feldman and Pentland, 2003: 95).

Now it seems that 'routine' cannot consistently mean both: the repeated occurrence of particular patterns of behavior and the structures underlying and generating such patterns. As especially Hodgson points out, routines cannot possibly be both generative structures and the outcomes or products of such structures (Hodgson, 2008: 19–20). This impossibility for 'routines' to mean both is taken to imply by Hodgson and many others that if routines are patterns of behavior, they cannot be generative mechanisms or causes. But once we conceive of routines as multilevel mechanisms, we can see that no such thing follows from the unassailable truth that something cannot possibly be both a cause and an effect of that very cause.

Once we conceive of routines as multilevel mechanisms, we see that routines are themselves patterns of behavior that are capable of generating patterns of firm behavior. We have two behavioral patterns here and it is important to see that they are different ones. Routines are patterns of behavior within firms. They involve particular sequences of individual interactions inside firms. As was argued earlier, viewing routines as multilevel mechanisms is a natural and plausible way of cashing out this insight: routines *are* particular patterns of individual interaction in particular artificial environments. The particular patterns of individual interaction in particular artificial environments are what we see if we zoom in on the functioning of particular routines. It is these behavioral patterns that we see if we look inside a firm in routine operation at what is happening at the level of individual interactions. What these 'intra-firm' behavioral patterns eventually generate as their outputs are particular patterns in the firms' production of goods and services. If we zoom out, we see that routines reliably and predictably link particular sets of inputs to firms to particular outputs. As such they can be conceived of as decision rules of firms.

'Firm behavior' is often taken to refer only to the outcomes of such decision rules: the goods or services actually produced. But 'firm behavior' might be understood in a wider or broader sense than just the goods or services ultimately produced. If a firm is in routine operation, the functioning of routines (or, if you want, the expression of routines) arguably is also part of firm behavior. As patterns of behavior occurring inside the firm, the functioning of a routine can itself be called a species of firm behavior. Thus, in a more integral understanding

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

of 'firm behavior' the working of routines are themselves part of firm behavior.[10] In this sense, highlighting the routines in action is highlighting how a firm in routine operation behaves in transforming its inputs into its outputs.

Another way of putting this is to say that routine firm behavior can be observed at different levels of analysis. We can look at routine firm behavior at a relatively high level of analysis. What we then see is how particular sorts of inputs to the firm as a whole are repeatedly transformed into determinate outputs. Routines of firms then appear as undifferentiated wholes, as decision rules of firms. What goes on within firms, which is what is going on in routines if the firms are in routine operation, is then out of sight. This becomes visible only (the 'black box' is opened) if we decide to look at routine firm behavior at lower levels of analysis. We then see particular sequences of interactions in firms in which units or teams, and individual employees (or firm members) in them, figure as entities. Thus, it is the same firm behavior that we are looking at. It is just that we either look at the whole causal process from above (or from the outside) or look at particular chunks of it from below (or from the inside).[11]

In sum, viewing routines as multilevel mechanisms helps us see how two different and even seemingly contradictory characterizations of routines can be reconciled in a coherent way. It enables us to see that routines are generative mechanisms that are themselves behavioral patterns of (or rather within) firms. As behavioral patterns they reliably and predictably produce firm outputs when activated by standard inputs.

5. What about routines as organizational dispositions?

Hodgson (2008) (and Hodgson and Knudsen, 2006) also view routines as generative mechanisms. Contrary to what is argued in this paper, however, they argue that generative mechanisms cannot possibly be patterns of behavior. Arguing that generative mechanisms are patterns of behavior would amount to confusing potentiality and actuality.

Hodgson chooses to define routines as dispositions to activate patterns of behavior rather than the patterns of behavior themselves: 'Routines are organizational dispositions to energize conditional patterns of behavior within an organized group of individuals, involving sequential responses to cues' (Hodgson, 2008: 21, see also 25). I take it that this definition implies that 'routines' do not refer to patterns of behavior at all, neither to patterns of individual interactions inside firms, nor to their outputs. What the definition suggests is that the total causal process starting from certain inputs (or 'cues') to a firm and finally

10 Nelson and Winter speak of 'a collection of individual member performances that actually constitute a productive performance for the organization as a whole' (Nelson and Winter, 1982: 104).

11 Many have noticed that there is a gap between Nelson and Winter's subtle and rich informal discussion of routines and their 'crude' formal treatment of them in their models. Viewing routines as multilevel mechanisms helps in bridging this gap.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

resulting in some output of a firm looks as follows: particular inputs to a firm activate a particular disposition, which energizes a pattern of behavior within the firm, which finally issues in a particular output of the firm. Note that this is quite a roundabout way of conceiving routine firm behavior compared with the conception of routine firm behavior envisioned in this paper. In the conception envisioned here, particular inputs directly activate a pattern of behavior within the firm (i.e., a particular routine), which issues in a particular output of the firm. Hodgson seems to build in an extra link in the causal chain: (conditional) patterns of behavior are 'energized' by inputs, not directly but indirectly, via the activation of organizational dispositions.

What is the use of postulating an extra 'entity', a disposition to energize patterns of behavior, on top of this? Hodgson's definition seems to be at least partly motivated by what he takes to be requirements of a Darwinian economic evolutionary theory. In their Generalized Darwinism project, Hodgson and Knudsen are looking for replicators in the economic domain.[12] And they believe that routines (and habits at the individual level) fill the bill. But Hodgson's main reason for rejecting an interpretation of 'routines' in terms of behavioral patterns seems to be that routines do not cease to exist when they are not exercised for a while. A firm that is inoperative during weekends, for example, does not all of a sudden loose its routines. Yet if routines were behavioral patterns, we would be forced to say that firms loose their routines over weekends. If we had to conclude on the behavioral understanding of 'routines' defended here that routines disappear during weekends, how then would it be possible that routines tend to reappear after weekends? This would be miraculous. Hodgson and Knudsen conclude that routines are not behavior, but capacities or dispositions underlying behavior (Hodgson and Knudsen, 2004: 290).

This conclusion does not follow from their argument, however. Hodgson is right that some capacities must be stored in a firm for a firm to display the same patterns of behavior on Mondays that the firm also displays on Fridays. But it does not follow that we should call such dispositions routines rather than the behavioral patterns that they 'energize'. We are still free to choose to call patterns of behavior routines. In the understanding of 'routine' argued for in this paper, routines are (part of) what is going on in a firm when a firm is in routine operation. On this understanding, there are no routines functioning if a firm is not in routine operation. Thus if a firm is inoperative during weekends, there are no routines functioning during weekends. It is not that we would be forced to say that routines mysteriously disap pear in weekends. There simply are no routines functioning during weekends. Contrary to what Hodgson suggests, understanding 'routines' as behavioral patterns does not imply that it is inexplicable that the same routines that are triggered by some inputs on Fridays

12 For critical discussions of Generalized Darwinism, as especially of the alleged necessity of specifying replicators in an economic Darwinian theory, see Vromen (2007, 2008).

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

are also triggered by similar inputs on Mondays. Understanding 'routines' as behavioral patterns does not imply that firms cannot retain their capacities to display the same routines when the routines are not activated for a while.

Hodgson might rightly insist that what we still need to explain then is how firms can retain their capacities to display the same patterns of behavior when the capacities are not exercised for a while. It is true that the behavioral understanding of 'routines' in terms of behavioral patterns within firms by itself does not give us such an explanation. But it is not true that a behavioral understanding makes it impossible to give such an explanation. The behavioral understanding in fact gives a hunch as to where we should be looking for such an explanation. It prompts us to look at the conditions that must be met in a firm for the same behavioral pattern to be activated by similar inputs (even if the time intervals between the occurrence of the inputs are considerable). Looking into the firm at what must happen at the level of individual interaction sheds some light on this. The similar input must first reach and be processed in the same way by the same participant (or by a participant who has been assigned the same job or role), the output of this must be transferred to a specific other participant (or by a new employee who has taken on this specific job or role), who must process the output of the first participant in the same way as before, and so on. For all this to work as it did before, the participants must have the same or similar procedural memories (Cohen and Bacdayan, 1994) and expectations as they had before. These psychological and mental states of participants are presumably not all that matters to the capacities of firms to display the same patterns of behavior when faced with similar inputs. 'Organizational memory' is greatly enhanced by artifacts and by stable arrangements of participants (especially the division of labor in different roles and positions). But it seems undeniable that the procedural memories and expectations of participants are crucial for this. Thus, the behavioral understanding of routines defended in this paper suggests the broad contours of an explanation of why firms tend to retain their capacity to display the same behavioral patterns when presented with the same stimuli even if the capacity is not exercised for a while.[13]

It has to be noted also that just stipulating that routines are dispositions rather than patterns of behavior, as Hodgson and Knudsen do, or just mentioning something like 'organizational memory' as an answer, by themselves do not offer much of an explanation either. Merely understanding routines as dispositions or capacities does not bring us closer to a satisfactory explanation. Their understanding allows us to say that firms can retain their routines also during times in which the routines are not exercised. But saying so by itself does not address the issue of how the routines are retained. This issue is as unresolved as before. Calling dispositions routines does not tell us either what exactly is

13 It is generally acknowledged that there are limits to this: if capacities are not exercised for too long, they might vanish.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

organizational memory, where in the organization is it stored, how does it relate to individual memory and the like. These are the questions we want answers for. This remains as puzzling or mysterious as before. I conclude from this that the concern that Hodgson and Knudsen raise makes sense, but that, contrary to what they argue, this concern does not provide an argument in favor of their own preferred understanding of routines as dispositions.

The advantage of viewing routines-as-mechanism-as-patterns of interaction within firms (as I propose) over viewing routines-as-mechanisms-as-dispositions that are themselves not patterns of behavior (as Hodgson and Knudsen propose) is that it provides a more direct and less mysterious and roundabout way of looking at firm routine behavior. Instead of arguing that certain cues in the environment trigger (or activate) particular dispositions in firms that subsequently energize conditional behavioral patterns (as Hodgson and Knudsen do), we can simply say (as I do) that particular members have been assigned the job to process incoming messages and to pass through the results of their processing to other members. Postulating the existence of yet another entity (Hodgson and Knudsen's disposition) raises more questions than it answers. What is this disposition (over and above the patterns of individual interaction), where does it reside, how and when is it activated, how does it energize patterns of behavior, and so on. My behavioral understanding of mechanisms (and, hence, also of routines), which does not postulate such a rather mysterious entity, greatly facilitates empirical research.

6. Some unresolved issues

Viewing routines as multilevel mechanisms is not a panacea that solves all our problems and answers all our questions. Far from that. It rather presents a way of looking at routines that is helpful for resolving some issues. Much of the contents and substance that are needed to resolve many remaining issues must come from elsewhere. Yet viewing routines as multilevel mechanisms does bring some unresolved issues in sharper focus. In discussing routines as multilevel mechanisms and fleshing out their implications, we stumbled across issues that are still awaiting resolution. Let us recapitulate a few of them.

One such issue is where and how the knowledge is stored in a firm that is needed for a smooth functioning of its routines, even if the routines have not been exercised for a while. Above I conceded that that viewing routines as multilevel mechanisms by itself does not provide an explanation for this. Viewing routines as multilevel mechanisms does suggest, however, where to look for such an explanation: not just at individual firm members and their procedural memories, but also at the ways in which the members are organized and interact in the firm and at the artifacts that are used.

Another unresolved issue pertains to the boundaries of routines. If routines are viewed as patterns of individual interaction, as it is argued here, then where

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

are we to draw the boundary of routines? If a firm is in routine operation, there is a long chain of individual interactions leading to the firm's inputs to its outputs. We could choose to call the whole chain a routine. Alternatively, we could choose to carve up the chain into several routines. How fine-grained or coarse-grained should our partitioning of these long chains into separate routines be? Pentland and Feldman put it as follows: 'Identifying a particular routine is a bit like trying to isolate the Gulf Stream from the Atlantic Ocean' (Pentland and Feldman, 2005, 798). It is not clear what would be reasonable, non-arbitrary criteria to individuate routines. It might be argued that not much hinges on this. But it is clear that a failure to resolve the issue does not facilitate empirical research on routines.

The issue of how to individuate routines is a conceptual issue. Other issues that we have stumbled across are rather empirical in kind. The issue of how stable and robust routines are is a case in point. We saw that the very notion of a routine arguably implies some stability or durability of the patterns of individual interaction that are denoted by the notion. It seems it does not make sense to call a sequence of individual interactions that occurs only once a routine. For the same reason, the very notion of a mechanism also seems to imply some stability. Viewing routines as mechanisms thus draws extra attention to issues of stability and robustness. But precisely how stable or enduring a pattern of individual interaction is, is an empirical issue that cannot be prejudged on the basis of theoretical or conceptual grounds. Just how long a firm can retain its capacity to display some patterns of interaction if the capacity is not exercised for a while is an issue that can be settled only on empirical grounds, for example. The same holds for the issue to what extent routines can survive personnel turnover.

Another empirical issue is how much of firm behavior is routine. Again, viewing routines as multilevel mechanisms does not prejudge this issue. It leaves open the possibility that most of firm behavior is non-routine; and it also leaves open the possibility that even though firms are in routine operation most of the time, the decisive operations of firms, which determine their relative performance for the greatest part, are non-routine. If so, explanations of differential firm performance in terms of the firms' routines would get things wrong. In general, viewing routines as multilevel mechanisms need not necessarily support Nelson and Winter's claim that knowledge of the routines of firms is the heart of understanding firm behavior (Nelson and Winter, 1982: 128). This, I submit, is a merit rather than a weakness of the view on routines argued for in this paper. It is a merit of the behavioral understanding of routines in terms of observable patterns of interactions defended here that it facilitates empirical (a posteriori) research on these issues. Any attempt to prejudge these issues on a priori grounds should be regarded with suspicion.

Yet another empirical issue, finally, is what specific sorts of inputs (or stimuli), if any, activate specific routines. Many scholars and commentators have argued that routines are context-dependent. A particular routine is activated only by

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

specific sorts of inputs. This raises the issue by what specific sorts of inputs a particular routine is activated. Just observing by what specific inputs a particular routine has been activated in the past does not suffice to resolve this issue. The possibility cannot be ruled out that quite unexpectedly, the same routine might be activated by unprecedented novel stimuli. Routines can 'misfire'. A reasonable working hypothesis for finding out more about this seems to be that routines evolved to solve recurrent past problems in (or for) firms. The challenge then would be to specify the problems that routines evolved to solve. If we know these problems, we might try to infer the sorts of stimuli that activate particular routines. But although this strategy might work fine in general, there are rather obvious reasons for believing that the strategy does not work perfectly. We might misspecify the recurrent evolutionary problem that a routine supposedly evolved to solve, for example. This only further underwrites the need for empirical research.

7. Conclusions

The merits of viewing routines as multilevel mechanisms are that it helps us resolve some conceptual issues and that it helps in getting other, as yet unresolved, empirical issues in sharper focus. It helps in putting together a coherent picture of what routines are, what routines do, and how they do it. In particular, it helps in getting a clearer picture of how genes, skills, and routines are ontologically (rather than metaphorically) related to each other. We saw that it can weave together two seemingly contradicting conceptions of routines in a clear and coherent way. It allows us to see that routines are both generative structures and behavioral patterns. As *generative* structures, routines produce recurrent patterns of firm behavior (in the narrow sense of firm output), but as *multilevel* mechanisms, routines themselves are also recurrent patterns of interaction in firms.

Because of its 'behavioral' spirit, viewing routines as multilevel mechanisms (rather than as, for example, unobservable dispositions of firms to energize patterns of behavior in firms) greatly facilitates empirical research. It also suggests what future empirical research on routines should focus on. The specific recurrent sequences of individual practices that familiar, standard tasks activate in specific firms are what we should be looking for. The properties of these sequences should also be investigated empirically. Just how stable over time are they? How often en frequently are the sequences to be activated for them to keep recurring? Do the sequences tend to survive personnel turnover? A further merit of the 'behavioral' view defended here is that it does not prejudge the issue of whether routines are indeed the decisive factors in determining the relative performances of firms that many take them to be. It does not rule out the possibility that the decisive activities in firms are non-routine (or are not part of routines). This again is an empirical issue that cannot be settled on theoretical or conceptual grounds.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

References

Abell, P., T. Felin, and N. J. Foss (2008), 'Building Microfoundations for the Routines, Capabilities, and Performance Links', *Managerial and Decision Economics*, **29**: 489–502.

Becker, M. C. (2004), 'Organizational Routines: A Review of the Literature', *Industrial and Corporate Change*, **13**(4): 643–678.

Becker, M. C. (2005), 'A Framework for Applying Organizational Routines in Empirical Research: Linking Antecedents, Characteristics, and Performance Outcomes of Recurrent Interaction Patterns', *Industrial and Corporate Change*, **14**(5): 817–846.

Becker, M. C., N. Lazaric, R. R. Nelson, and S. G. Winter (2005), 'Applying Organizational Routines in Understanding Organizational Change', *Industrial and Corporate Change*, **14**(5): 775–791.

Bechtel, W. and A. Abrahamsen (2005), 'Explanation: A Mechanist Alternative', in C. F. Craver and L. Darden (eds.), 'Special Issue: Mechanisms in Biology', *Studies in History and Philosophy of Biological and Biomedical Sciences*, **36**: 421–441.

Cohen, M. and P. Bacdayan (1994), 'Organizational Routines are Stored as Procedural Memory: Evidence from a Laboratory Study', *Organization Science*, **5**: 554–568.

Cohen, M. D., R. Burkhart, G. Dosi, M. Egidi, L. Marengo, M. Warglien, and S. Winter (1996), 'Routines and Other Recurring Action Patterns of Organizations: Contemporary Research Issues', *Industrial and Corporate Change*, **5**: 653–698.

Craver, C. F. (2007), *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*, Oxford: Oxford University Press.

Craver, C. and W. Bechtel (2007), 'Top–Down Causation without Top–Down Causes', *Biology and Philosophy*, **22**(4): 547–563.

Cummins, R. (1975), 'Functional Analysis', *Journal of Philosophy*, **72**: 741–764.

Cummins, R. (2000), 'How Does It Work? vs. What Are the Laws? Two Conceptions of Psychological Explanation', in F. Keil and R. Wilson (eds.), *Explanation and Cognition*, Cambridge, MA: MIT Press, pp. 117–145.

Darden, L. (2001), 'Discovering Mechanisms: A Computational Philosophy of Science Perspective', in K. P. Jantke and A. Shinohara (eds.), *Discovery Science*, *Proceedings of the 4th International Conference*, New York: Springer-Verlag, pp. 3–15.

Dawkins, R. (1976), *The Selfish Gene*, Oxford: Oxford University Press.

Elster, J. (1989), *Nuts and Bolts for the Social Sciences*, Cambridge: Cambridge University Press.

Elster, J. (1998), 'A Plea for Mechanisms', in P. Hedstrom and R. Swedberg (eds.), *Social Mechanisms: An Analytical Approach to Social Theory*, Cambridge: Cambridge University Press, pp. 45–73.

Elster, J. (2007), *Explaining Social Behavior: More Nuts and Bolts for the Social Sciences*, Cambridge: Cambridge University Press.

Feldman, M. S. and B. T. Pentland (2003), 'Reconceptualizing Organizational Routines as a Source of Flexibility and Change', *Administrative Science Quarterly*, **48**: 94–118.

Felin, T. and N. J. Foss (2005), 'Strategic Organization: A Field in Search of Micro-Foundations', *Strategic Organization*, **3**: 441–455.

Felin, T. and N. J. Foss (2006), 'Individuals and Organizations: Further Thoughts on a Microfoundations Project', in D. Ketchen and D. Bergh (eds.), *Research Methodology in Strategy and Management* 3, Amsterdam: Elsevier, pp. 253–288.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

Felin, T. and W. S. Hesterly (2007), 'The Knowledge-Based View, Nested Heterogeneity, and New Value Creation: Philosophical Considerations on the Locus of Knowledge', *Academy of Management Review*, **32**: 195–218.

Glennan, S. S. (1996), 'Mechanisms and the Nature of Causation', *Erkenntnis*, **44**: 49–71.

Glennan, S. S. (2002), 'Rethinking Mechanistic Explanation', *Philosophy of Science (Supplement)*, **69**: S342–S353.

Hodgson, G. M. (2008), 'The Concept of a Routine', in M. C. Becker (ed.), *Handbook of Organizational Routines*, Cheltenham and Northampton: Edward Elgar, pp. 15–28.

Hodgson, G. M. and T. Knudsen (2004), 'The Firm as an Interactor: Firms as Vehicles for Habits and Routines', *Journal of Evolutionary Economics*, **14**(3), 281–307.

Hodgson, G. M. and T. Knudsen (2006), 'Why We Need a Generalized Darwinism, and Why Generalized Darwinism is not Enough', *Journal of Economic Behavior & Organization*, **61**(1), 1–19.

Hutchins, E. L. (1995), *Cognition in the Wild*, Cambridge, MA: MIT Press.

Kogut, B. and U. Zander (1992), 'Knowledge of the Firm, Combinative Capabilities, and the Replication of Technology', *Organization Science*, **3**: 383–397.

Latour, B. (1986), 'The Powers of Association', in J. Law (ed.), *Power, Action and Belief*, London: Routledge & Kegan Paul, pp. 264–280.

Levitt, B. and J. March (1988), 'Organizational Learning', *Annual Review of Sociology*, **14**: 319–340.

Machamer, P., L. Darden, and C. F. Craver (2000), 'Thinking about Mechanisms', *Philosophy of Science*, **67**: 1–25.

Nelson, R. R. (2007), 'Universal Darwinism and Evolutionary Social Science', *Biology and Philosophy*, **22**(1): 73–94.

Nelson, R. R. and S. G. Winter (1982), *An Evolutionary Theory of Economic Change*, Cambridge, MA: Harvard University Press.

Pentland, B. T. and M. S. Feldman (2005), 'Organizational Routines as a Unit of Analysis', *Industrial and Corporate Change*, **14**(*5*): 793–815.

Pentland, B. T. and H. Rueter (1994), 'Organisational Routines as Grammars of Action', *Administrative Sciences Quarterly*, **39**: 484–510.

Simon, H. A. (1969), *The Sciences of the Artificial*, Cambridge, MA: MIT Press.

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

Vromen, J. J. (1995), *Economic Evolution: An Enquiry into the Foundations of New Institutional Economics*, London: Routledge.

Vromen, J. J. (2004), 'Conjectural Revisionary Economic Ontology: Outline of an Ambitious Research Agenda for Evolutionary Economics', *Journal of Economic Methodology*, **11**(2): 213–247.

Vromen, J. J. (2006), 'Routines, Genes and Program-based Behaviour', *Journal of Evolutionary Economics*, **16**(*5*): 543–560.

Vromen, J. J. (2007), 'Neuroeconomics as a Natural Extension of Bioeconomics: The Shifting Scope of Standard Economic Theory', *Journal of Bioeconomics*, **9**: 145–167.

Vromen, J. J. (2008), 'Advancing Evolutionary Explanations in Economics', in H. Kincaid, and D. Ross (eds.), *Oxford Handbook on Philosophy of Science on Economics*, Oxford: Oxford University Press, pp. 337–368.

Vromen, J. J. (2009), 'Microfoundations in Strategic Management? Squaring Coleman's Diagram' (forthcoming in *Erkenntnis*).

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press

Wimsatt, W. C. (1994), 'The Ontology of Complex Systems: Levels, Perspectives, and Causal Thickets', in M. Matthen and R. Ware (eds.), 'Biology and Society: Reflections on Methodology', *Canadian Journal of Philosophy*, **20**: 207–274.

Wimsatt, W. C. (2007), *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*, Cambridge, MA: Harvard University Press.

Young, H. P. (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*, Princeton, NJ: Princeton University Press.

https://doi.org/10.1017/S1744137410000160 Published online by Cambridge University Press