

Routines and Other Recurring Action Patterns of Organizations: Contemporary Research Issues

Michael D. Cohen, University of Michigan,
and SFI, (session chair)

Roger Burkhart, John Deere and Company, and SFI
Giovanni Dosi, University "La Sapienza",
Rome, and IIASA

Massimo Egidi, University of Trento

Luigi Marengo, University of Trento

Massimo Warglien, University of Venice

Sidney Winter, University of Pennsylvania

and with comments by Benjamin Coriat,
University of Paris XIII¹

A report of a working group meeting held under the joint
sponsorship of the Santa Fe Institute (SFI) and the International
Institute of Applied Systems Analysis (IIASA) in August 1995.
Released as a Santa Fe Institute Working Paper, November 1995

¹ Complete addresses of the authors are in Appendix B

This working paper reports the main lines of discussion that developed during a four-day workshop held August 10-13, 1995, at the Santa Fe Institute, Santa Fe New Mexico. The workshop was devoted to exploring the difficulties and promise of current research on organizational routines and related concepts.

We decided during the concluding day of our sessions that the productive conversations we had held could be quite valuable to others if we could craft an effective format for sharing them. Many of us were surprised, both by our points of agreement and disagreement, and all of us felt that arguments made during the discussions contributed novel and valuable linkages to results and methods in related fields. It was our sense that other researchers working on these issues could well benefit from an account of what was said.

This working paper is therefore designed with the audience of fellow researchers — and especially graduate students — chiefly in mind. It reorganizes (and compresses) the chronological flow of our discussions, in an effort to bring out major themes. In each of its sections it also contains commentary, pointers to literature, and items of dispute, all added by the participants as they reviewed the working paper draft. Some comments of one participant evolved into a much larger argument not easily sliced into separate chunks. So those appear as Appendix A. Thus the document distills not only the four days of discussion, but some reflections on those discussions during the few weeks following the meeting while the draft was created and circulated for commentary.²

The seven gathered at the workshop were drawn from only a subspace of the many researchers working with the concept of routine. All shared a commitment to viewing actors through the lens of research in cognitive psychology on short-term memory limits, reasoning powers, and differentiated forms of long term memory and learning — an approach we will label "cognitive realism". All shared an approach to change processes in organizational systems as broadly evolutionary in their character ("an evolutionary approach to change"). All regarded as highly significant "political" and social forces within organizations, the many channels through which individual and local interests assert themselves at the expense of more global organizational concerns ("the diversity of "fitness" forces").

These common commitments probably contributed to the many points of agreement reached over the four days. Differences in emphasis between the three commitments probably account for many of the remaining points of contention.

² It might be thought of as a sort of "flattened hypertext." We are exploring the possibility of making it available for further commentary via the World Wide Web.

What follows is a synthetic reconstruction of the conversations, indicating both significant agreements and remaining open issues (marked with a *). Each section closes with remarks contributed by workshop participants during revision of this paper, and attributed directly to them via their initials. Those participants (with their identifying initials) were: Roger Burkhart [RB], Michael Cohen [MC], Giovanni Dosi [DD], Massimo Egidi [ME], Luigi Marengo [LM], Massimo Warglien [MW], and Sidney Winter [SW]. Some of Egidi's comments form a longer statement and are collected in Appendix A. (An eighth participant, Benjamin Coriat, was unable to attend due to an injury. Giovanni Dosi labored heroically to indicate at numerous points in the workshop positions he felt sure our colleague Coriat would have taken if he could have been with us. And Coriat has supplied some comments, generated during a review of our draft, and labelled [BC].)

Here are the principal themes of the workshop, each of which is discussed in a section below:

1. Why is it important to carefully examine research on routine?
2. What "action patterns" are under discussion?
3. How can routines and other recurring patterns be usefully categorized?
4. What are the research implications of recent cognitive results?
5. How can evolution be seen in relation to action patterns?
6. How can simulation contribute to better theory in this area?
7. What are examples of various approaches to empirical research that reveal key problems or results?
8. How might 'routine' be defined?³

³ We use single quotes are used when referring to the label for the concept indicated by the quoted term. Double quotes are used in their normal fashions.

1. Why is it important to carefully examine research on routine?

Our meeting grew out an earlier session held during the winter of 1995 in Laxenburg, Austria, with the sponsorship of the International Institute of Applied Systems Analysis (IIASA). That meeting gathered numerous researchers working with evolutionary models of economic phenomena. During the course of the meeting it became clear that the concept label ‘routine’ was being used very widely, and — unfortunately — in widely varying senses. Routine is a fundamental concept in the Nelson and Winter (1982) approach to economic evolution, which was serving many of the researchers present as an intellectual point of departure. So it was especially of concern that many were using the concept with meanings that had drifted well off from the Nelson-Winter definition. And it soon became clear that, with the benefit of a dozen years of additional research by several hundred readers of their book, Nelson and Winter, who were present, also saw the definition of ‘routine’ as ripe for some reappraisal.

By the end of the meeting there was agreement on the potential of another, smaller gathering that would bring key participants together again for an extended session focused strictly on routine: trying to clarify the concept in order to make empirical studies and theoretical models more coherent. The Santa Fe Institute and IIASA were eventually persuaded to jointly sponsor the meeting, since both have serious interests in research bearing on the concept.

Several of us attended other, more general conferences on organization held in the intervening months and felt those meetings confirmed the timeliness of our decision. ‘Routine’ seems destined to be a keyword in a very large number of papers appearing in the mid-nineties. But examining the papers shows little progress so far in reaching agreement on what routines are — and therefore on how or why social scientists should study them.

Commentary

on "needed LM To some of us (to me at least!) routines, (and, more generally, rule-guided behavior) appeal as a basic "foundational" concept for an alternative theory of decision making to the neo-classical one. This endeavour goes far beyond the observation of organizational behavior, and requires a special effort in carefully defining the concept and formalizing it (maybe by boldly simplifying and losing a great part of the richness of different aspects which emerge from empirical observation). For this purpose the definition and formalization of routines given by **examination"**

computer scientists seems particularly appealing, but only partly overlapping with the one emerging from the organization science/evolutionary economic tradition. (See the detailed comments along this line by ME in Appendix A). Clarification on these definitional issues is, in my view, most needed.

MW The concept of routine is probably paying the price for its success: as it diffuses, its meaning gets increasingly vague and subject to arbitrary extensions. This is to some extent the unavoidable side-effect of popularity (after all, the "butterfly effect" has gone through much more cruel vicissitudes). However, some reasons for our discomfort are rooted in one of the original metaphors characterizing routines in the Nelson and Winter book. Much of the temptation of naming 'routine' anything that is not new under the sun of organizational life may belong to their suggestion that routines are the "organizational DNA": a metaphor that conveys the sense that routines are *the* organizational replicator, the fundamental structure undergoing reproduction. It thus becomes a natural temptation to claim that, whenever you observe some feature which is stable and reproducible over time or space, it has to be a routine. I think that the wise caveats that in Nelson and Winter's book came along with the introduction of the genetic metaphor have been largely overridden by this last in readers' minds. In contrast, cultural phenomena (to which economic evolution clearly belongs) present a wide variety of replicators, whose nature and level of aggregation is very diverse; some of the richness of cultural evolution resides in the way multiple replication processes interweave in quite complicated ecologies. So, I think that recognizing the diversity of replicators and replication processes in action within organizations and starting to define a more articulate language for speaking about them is a useful way to improve evolutionary thinking about firms, and is a good service to the notion of routine itself.

2. What action patterns are under discussion?

Our clear principal focus is on action patterns that can be called 'routine'. However, "other recurring action patterns" became part of the title of the workshop (and this paper) because we wanted to acknowledge that there are other recurring actions in firms that also merit close study, but may not exactly be routines. For example, there might be widely shared "heuristics", rules of thumb used in common situations. They could give important regularity to firm actions and be important topics for research, but might lie outside a definition of 'routine'.

Winter explored the topic of heuristics and other non-routine patterns in a short memo prepared for the group. A revised version of that memo is in the section on categorization below.

This meeting occurred because all the participants believe that useful theory about firms (and many non-profits and government agencies) must be solidly based in a realistic, empirically informed, account of multi-person action. As Herbert Simon observed (1981), agents capable of perfectly tailoring their actions to their world would require no theory. Since they would optimize, one would need only a theory of that environment to which they flawlessly responded (and a statement of their goals). None of the participants believes that a good theory of firms (etc.) is likely to be found this way. All believe that firms are not frictionless reflections of their momentary environments, but rather highly inertial action repertoires, responding to — indeed perceiving — today's environment largely in terms of lessons learned from actions in days gone by.

In the consensus view, it is essential to understand how such action repertoires are assembled, maintained, and modified. The perspective points a researcher's interest toward action patterns with these three properties:

1. Recurring (actions taken by the same actors at different times — or by different actors — that we want to call "the same actions", e.g., approving a purchase order)
2. Selectable (actions — or patterns of action — that could be subject to forces that would make them more or less likely to recur, e.g., setting the R&D budget at six percent of total expenditure)
3. Set in organizational context (the actions are not those of isolated individuals, e.g., trying alternative tools on an assembly line, not in one's basement).

* In addition, many felt the actions being discussed were typically not deliberative choices. They wished to preserve the prominence that tacitness and automaticity have generally had in distinguishing routine action from choice and problem solving. Some felt the "other recurring action patterns" label in the workshop title was useful in leaving a place for deliberative choices while keeping "routines" for the non-choice cases, and yet another group (the "imperialists") were happy to include deliberative choice under routines leaving little in the other category. They note that even the standard of "acting rationally" is evolving as decades roll by and fields like finance and operations research become more sophisticated. Thus young managers work hard to learn how to "choose rationally"

in school and early jobs, and tend to stick to the patterns they have learned through the remainder of their careers, making rational actions a recurring pattern and potentially a species of routine.

* Finally, some embraced a distinction of ‘action’ and ‘behavior’ in order to resolve ambiguities that otherwise arise. When the same observable behavior occurs while very different intents are being accomplished (e.g., faxing an affected colleague the press release just before or after the public announcement) the cases might be coded as the same raw behaviors (fax the release), but interpreted by the actors as meaningfully different actions (pre—notified, or not). And different observable behaviors may occur in order to accomplish the same underlying intent (e.g., "notify her before it's public", once by phone, another time by fax). Others felt it was acceptable to treat the two terms as interchangeable. Some felt that the distinction was meaningful but that behavior rather than action should be given priority in view of its greater observability. In the remainder of this account ‘action’ is generally used in a sense including intention. ‘Behavior’ is used to indicate directly observable events.

**Commentary
on "actions
under
discussion"**

MW The problem whether to name ‘routine’ all "recurring action patterns" in organizations or just those that imply non-deliberate, automatic action may be a question of terminological taste. In any case, I think the substantive question is to start analyzing the structure of such action patterns and to make useful distinctions. Whether we want or not to call ‘routine’ the working of a kanban system in an assembly line, what strikes me is that we don’t have yet a language for going too far in analyzing its working. To be sure, there is a lot of tacit knowledge and of non-deliberate behavior in the way workers can achieve coordination in a kanban system: there’s more in it than workers or production engineers can say. But stability and replicability is assured to a large extent also by the following of explicitly stated rules and instructions (like "stop working if you don’t have any production kanbans authorizing it"), and by a carefully designed artifactual environment within which workers can learn coordinated behaviors. All these ingredients are "recurring, selectable, and set in organizational context" - but they differ both in their cognitive nature and in the way they are reproduced. Moreover, if one wants to characterize a kanban system as a recurring action pattern, one has to look not only at those diverse ingredients, but also (and maybe chiefly) at the architecture of relations among those ingredients - how they fit reciprocally and

with other processes such as set-up time reduction and working capital reduction. This relational architecture is itself recurring, selectable, and set in organizational context. Thus, whatever may be the terminological standard, the problem is finding a language for defining the ingredients of recurring action patterns and the architecture tying those ingredients together and giving them coherence.

LM 1) I have some doubts on "Selectable" as a key ("definitional") property of routines. That all kinds of actions and action patterns are selectable (under a broad definition of "selection") is rather obvious and I don't find it at all peculiar of routines. Certainly the role of selection pressure on the shaping and evolution of routines is a key issue, but I wouldn't see selectability as a key discriminating characteristic of routines.

2) The discussion on automaticity vs. choice (imperialists vs. ?) is a basic one and should be perhaps more emphasized. To simplify the positions I see it in this way: those who stress the automaticity have in mind a very clear distinction between what is and what is not a routine. A standard psychological test (performing some other demanding activity while one is executing a routine) would perfectly discriminate. But this would probably narrow down (certainly too much for the economists, cf. my point under "needed examination") the domain and the interest of the concept. On the other end the "imperialists" (who basically have in mind the computer science view of routines) tend easily to an all-embracing definition (routine=executable procedure).

GD In (Dosi, Marengo, Bassanini, and Valente, 1995) we add two further characteristics of routines (which are there considered in terms of behavior patterns): 1) their context-dependence and 2) their invariance vis-a-vis fine informational change, once given the context. I think the former attribute is important in that it defines a set of behaviors which meets the logic of appropriateness (March 1992). Hence, Mr. Jones elicits specific behavioral repertoires which are appropriate to his being on an assembly line or having to answer a phone call, without having to select them by some backward inference from the ultimate consequences that they entail and often also without much intentionality. Regarding the second attribute, I have been quickly convinced by the Santa Fe meeting that it might induce some misunderstanding in that routines and other patterns of action might sometimes be extremely sensitive to signals from the environment (e.g. the behavioral pattern of a secretary might involve fine tuning on the mood of the

boss), although other action patterns display much higher invariance (e.g. spend 6% of sales on R&D; never hire a communist). However I would maintain the in both cases the behavioral patterns themselves are rather inertial and relatively invariant to signals and feedbacks which could notionally call for their revision. Putting it another way, I'm suggesting that the class of behavioral patterns we are discussing involve a major activity of information-framing and -censoring which sustains their reproduction over time, often well beyond the circumstances which spurred their introduction.

3. How can routines and other recurring patterns be usefully categorized?

Two approaches to this classification problem were developed for the workshop. One was based on the differing cognitive demands of various expressed action patterns. It appears immediately after this section as the first comment.

This first approach gave priority to cognitive processing loads as its forcing principle, and has the advantage of connecting to a rich and growing body of research on cognition. The second approach emphasized the demands of evolutionary change by focusing on classifying types of representations, each presumed to give rise to action patterns and thence to be subject to variation-yielding modification. It appears as the last comment after this section.

Thus if we are interested in an action pattern carried out by a group of actors in an organizational setting, we can ask how the representation of the action is maintained in the organization. Some answers, roughly clustered, might be:

1) in the memories of the individual actors for their respective roles in the overall pattern. 2) by means of locally shared language, with special significance given to terms used to trigger or carry out the actions. (Weick (1979) is especially informative on creation of shared meaning as a form of organizational memory). 3) via physical artifacts, such as tools, spatial arrangements, written codes of standard operating procedures, or computer systems. 4) via organizational practices such as archives, rotations of personnel, maintenance of working examples, or by building key assumptions into organizational structure (e.g., "architectural knowledge" of Henderson (1992)). 5) by means of globally shared language forms, such as formalized oral codes or pledges, or widely retold "war stories".

An interesting advantage of this type of classification arises from the importance in evolutionary analysis of sources of variation. By making explicit how an action pattern is stored between expressions, it suggests the processes that

might be creating useful, or deleterious, variation. For example, forgetting is a far more active process for declarative memories of theories than for procedural memories of skills. This suggests greater stability over time for action patterns consisting largely of skilled actions by group members evoked in a performance context (e.g., an assembly sequence), and relatively greater variability for action patterns maintained via jointly held theories of how the environment works (a market projection for a new product). The formulation also helps makes sense of frequent observations that changes in routines can be triggered by substitution of what were thought to be equivalent tools or raw materials, since these may have been serving a representational function. In turn, this observation provides a potential rationale for otherwise "irrational" resistance to the turbulence of supply chains that flows from emphasizing price competition in purchasing.

Commentary

on categorization

SW (“Routines resorted and glossed”) In their discussion of routines, Nelson and Winter apparently created an impression of themselves as "lumpers" with respect to this conceptual category. Although a lot of different aspects and types of routines were discussed, there was also language in the book that created the impression in many readers that some unitary concept of "a routine" was the key to our analysis. This note attempts to correct this misapprehension by pursuing a "splitter" approach to routines (actually, to a category broader than "routine" as I would now use that term). The distinctions involved in the "splitting" here are at the cognitive level; they categorize different ways that a routine may relate to the cognitive functioning of the individuals involved in its performance.

This cognitive approach neglects, and hence risks obfuscating, the contextual aspects of routines. Context dependence is fundamental; the effectiveness of a routine is not measured by what is achieved in principle but by what is achieved in practice; this generally means that the routine might be declared effective in some specific contexts, but perhaps not in others. One important aspect of total context is the physical, which includes both the local/artifactual complements to the routine (e.g. the requisite plant and equipment) and the broader physical environment that was not produced for the benefit of the routine (e.g., climate, air pollution, radiation, etc.) A second major aspect of context is motivational/relational: what is the explanation for the fact that the human beings involved in the performance are *willing* to do what they do? The "routine as truce" story is helpful here: once upon a time there was overt conflict, but in most cases

it is largely over when the observer comes on the scene. What the observer sees is therefore the product of cognitive functioning *constrained by sensitivity to the sources of conflict*. However, the broader the temporal and geographical scale on which one seeks to address the motivational issues, the less of the problem is unlocked by the "routine as truce" formulation.

The absence of these contextual factors from the following discussion implies that there are additional dimensions (perhaps several) on which some resorting and glossing of the routines concept might be called for. The case for making the distinctions noted below is not that they provide an exhaustive taxonomy, but that they provide a taxonomy that is useful, though partial.

The starting point here is the question "what does an evolutionary theory really require about firm behavior if its basic evolutionary logic is to track?" The answer is the presence of some sort of . . .

QGTs, "Quasi-genetic traits": Any trait that remains approximately constant in the organization over a period long enough for significant selective feedback to accumulate at a level where "outcomes" are tested by an "environment". (An analytical distinction between an outcome that is "tested" and an environment that does the testing is indispensable to an evolutionary view — but a change of unit/level of analysis may reclassify actual behaviors between these two categories.)

Some Types of QGTs, Distinguished by their Cognitive Aspects:

I Routines (in the broad sense, falling into 2 subcategories):

I.A Routines (narrow sense): complex, highly automatic (and at least in that sense "unconscious"), behaviors that "function as a unit" and typically involve high levels of information processing that is largely repetitive over separate invocations of the routine.

I.B Rules of thumb: quantitative, relatively simple decision rules that are consciously invoked and require low levels of information processing. Often, a rule

of thumb yields a "first approximation" answer to some quantitative question and shapes actual behavior only through this role as a target or norm.

I claim that the distinction is important, especially in understanding earlier literature; the idea that boundedly rational actors respond to a complex and uncertain or unknowable environment with *simple* rules captures what is going on with rules of thumb, but is far off the mark with routines (narrow sense), which are skill-like and may be awesomely complex and involve awesomely large amounts of unconscious information processing. Thus, the boundedly rational response has a big component of making the best of the situation by narrowing focus and suppressing deliberation, something quite different from the simplification response reflected in rules of thumb.

II Heuristics and strategies: Concepts and dispositions that provide orientation and a common structure for a range of similar problem-solving efforts, but supply few if any of the *details* of individual solutions. (This is a very broad category; examples are identified by the combination of broad scope and prescriptive force: guidance toward a solution is provided over a substantial range of perceived problems. By contrast, a routine provides a detailed response to a narrow problem that may not even be perceived as a problem, since its "solution" is at hand.)

III Paradigms, cognitive frameworks: Mental models that are so fundamental to the cognitive activity of the actor that they affect perception as well as problem-solving and other cognitive functions; presumptions that underlie the actor's ability to see the world as meaningful and understandable. Although these models may have large articulable elements (as in scientific paradigms), they can also be inculcated as the result of long consistent experience, and in that case are often tacit. The specific influence on *perception* is generally tacit regardless of the amount of articulation possible in other respects: people cannot explain with any completeness why they perceive things as they do.

IV Examples:

IV.A Routines (narrow sense): The actions of a retail clerk in accepting payment, recording the sale, perhaps bagging the purchase, etc., or of a bank teller in executing the bank's side of a transaction. The sequence of steps executed in clearing a check. Accounting systems. Airline reservation systems. The actions of a team of workers at a particular station on an assembly line.

IV.B Rules of thumb: Rules for determining R&D expenditure or advertising expenditure in relation to sales. The markup number in a markup pricing formula. (By contrast, the cost number that is marked up is often the product of a cost accounting system that is an elaborate routine (narrow sense).)

IV.C Heuristics and strategies: In the realm of decision making: "Do what we did the last time a similar problem came up." In technology: "Make process equipment larger (in physical volume) or faster (in cycle time)." "Make the product smaller" (e.g., in semiconductor devices.) In business strategy: "Always seek to be the number one or number two firm in every market you participate in — divest units that, over a period of time, fail to meet that standard." "Invest heavily in capacity to serve growing markets, so as to obtain learning curve advantages." (BCG doctrine). "Stick to the knitting." "Quality is job No. 1." In bargaining: "Always ask for more than you would be willing to settle for." "Weigh the effect that a proposed settlement may have as precedent or signal." In management: "Centralize to achieve greater responsiveness to top management, and other advantages." "Decentralize to achieve greater responsiveness to customers, and other advantages." As these examples suggest, a heuristic or strategy may permit concise articulation, but its consequences in behavior often depend on a much richer background of understanding that is not fully articulable. Heuristics and strategies

IV.D Paradigms and cognitive frameworks: At a fundamental level, commitments to concepts like causality, with causes temporally preceding effects, or to two-valued logic, or to category systems (as in Lakoff, 1987), or to "seeing is believing" (the naive version of the idea that there is a "neutral language of observation").

At a not-quite-so-deep level, scientific paradigms in the Kuhnian sense (Newtonian physics, phlogistic chemistry, neoclassical economics), and their counterparts in technology. (See especially Kuhn's Chapter 6, p. 63-4.) "A scientific paradigm could be approximately defined as an 'outlook' which states the relevant problems, a 'model' [i.e., an "exemplar"? — SW] and a 'pattern' of inquiry." (Dosi, 1984, p 14). In broad analogy with the Kuhnian definition of a "scientific paradigm", we can define a "technological paradigm" as a "model" and a "pattern" of solution of *selected* technological problems based on *selected* principles derived from natural sciences and on *selected* material technologies.

Or, the uni-causal principle — the widely held intuitive belief that every identifiable and value-laden "effect" must have a conceptually unique and commensurately value-laden "cause" — a principle that provides ground for conspiracy theories, often picks out "pilot error" as the "probable cause of an airline disaster, attributes cancers and birth defects to handy environmental problems created by big, money-grubbing corporations, etc.

At a more superficial level, the innumerable patterns of unconscious inference that derive from experience, typified by the tendency to perceive anomalous playing cards as normal — the black four of hearts is seen as a normal four of hearts or of spades (Bruner and Postman). Related mechanisms probably account for more consequential perceptual phenomena such as stereotyping on the basis of race, gender or ethnicity (note that the "experience" base of the inference scheme need not be first-hand or statistically valid). In the worlds of management and policy, the powerful role of perception of a proposal as "something that this organization is not likely to do" or "something that will not attract majority support".

Other QGTs

Although not having distinguishing cognitive features analogous to those of the types listed above, there are other types of QGTs that can provide the basic continuity needed for an evolutionary theory. In particular, there are *stable systemic traits*: traits that are not themselves "structural" but that are stabilized over long periods by some combination of features of the organization. An operational definition of "stable" is that the intra-organizational variation over time is small relative to the cross-sectional variation at a given time, or alternatively, that the intra-organizational variation over time for the trait in question is small compared to other traits of the same organization. In general, then, identification of a trait as "stable" depends on the availability of some standard of comparison.

Examples of stable systemic traits: Budget allocations in a government bureaucracy or a multi-divisional corporation with a weak center (Williamsonian "H-form" rather than "M-form") — often describable by a simple formula involving historical allocation patterns and perhaps sub-unit contribution, even though no such formula is explicitly acknowledged in the decision process. Financial ratios reflecting, on the average, the joint effect of a number of facets of business operations rather than any process focused on that ratio as such — sales/assets,

working capital/assets, times interest earned, etc. (If these are explicitly attended to, perhaps because outsiders are believed to be attending to them, then they would be in the category of rules of thumb.)

Of course, mere constancy is not enough to identify something as a "trait" of an organization or organism — just wait until that chameleon moves to a background of a different color and you will find what sort of "trait" his color is. As the examples above suggest, it is *something about the organization* that is producing the observed stability and will tend to preserve it in the face of environmental change, even if the changes are ones that might suggest the need for some adaptation.

GD 1) While I like the Winter taxonomy, I have some trouble on the current definition of the stable systemic traits (SST). The way I would interpret them is as relatively invariant ensembles of routines (narrow sense), rules of thumb and possibly heuristics. SST's, in this definition, would then be the central level of observation in order to identify corporate capabilities as currently discussed in the business literature. That is the proposition "company X is good at doing Y" represents a statement on the collective properties of its routines, etc. Likewise it is at the level of SST's that one might be able to discriminate among discrete types of organizational forms and model behaviors (e.g. Fordist or Toyotist).

2) It is true that rules of thumb are relatively simple, information poor rule. However, behind the simple rule there are likely to be several routines (narrow sense) that make the rule operational.

MC If we are interested in an action pattern carried out by a group of actors in an organizational setting, we can — in addition to looking at cognitive loads — ask how the representation of the action is maintained in the organization. Some answers, roughly clustered, might be:

1. in the memories of the individual actors for their respective roles in the overall pattern
2. by means of locally shared language, with special significance given to terms used to trigger or carry out the actions. (Weick (1979) is especially informative on creation of shared meaning as a form of organizational memory)
3. via physical artifacts, such as tools, spatial arrangements, written codes of standard operating procedures, or computer systems

4. via organizational practices such as archives, rotations of personnel, maintenance of working examples, or by building key assumptions into organizational structure (e.g., "architectural knowledge" of Henderson, 1992), or
5. by means of globally shared language forms, such as formalized oral codes or pledges, or widely retold "war stories".

An interesting advantage of this type of classification arises from the importance in evolutionary analysis of sources of variation. By making explicit how an action pattern is stored between expressions, it suggests the processes that might be creating useful, or deleterious, variation. For example, forgetting is a far more active process for declarative memories of theories than for procedural memories of skills. This suggests greater stability over time for action patterns consisting largely of skilled actions by group members evoked in a performance context (e.g., an assembly sequence), and relatively greater variability for action patterns maintained via jointly held theories of how the environment works (a market projection for a new product). The formulation also helps make sense of frequent observations that changes in routines can be triggered by substitution of what were thought to be equivalent tools or raw materials, since these may have been serving a representational function. In turn, this observation provides a potential rationale for otherwise "irrational" resistance to the turbulence of supply chains that flows from emphasizing price competition in purchasing.

4. What are the research implications of recent cognitive results?

One line of recent work in psychology has developed in a way that nicely reinforces traditional organization theory views of routine. Work on procedural memory in human individuals has shown that it has distinctive properties. It is centered on skills, or know-how, rather than on facts, theories, or episodes (know-that) which seem to be more the province of an alternate, "declarative", memory system. Procedural memory differs from declarative in its long decay times, and greater difficulty of transfer and of verbalization. This fits nicely with properties of routines observed in the field and in the laboratory (Egidi, 1994; Cohen and Bacdayan, 1994). And it appears to provide a firmer foundation in individual psychology for the characterization found in Nelson and Winter of routines as "tacit" and highly stable analogs of individual skills.

While it is useful that procedural memory aligns with many characteristics of highly routinized action patterns, it seems likely that much effective organizational

performance involves a mixture of such "automatic" or "tacit" elements together with a certain amount of "decision making" or "problem-solving" that is much more deliberative and self-aware in its character.

The group engaged in detailed analysis of hypothetical examples in which choice points were embedded within relatively automatic sequences. The aim was to clarify questions about whether the choices were part of one routine or were better viewed as sutures joining together several smaller routines. This issue became known as the problem of "grain size", whether a routine should generally be taken to be a rather large block of action, or a small block typically interleaved with other small routines and other forms of non-routine action.

The issues were illustrated as well by observations from the Egidi-Cohen-Bacdayan experiments, where coherent chunks are small at early stages of play and become larger with experience. But even at late stages of experience, it is problematic whether players of the card game could be said to have one routine or rather a repertoire of several which are called upon as contexts vary. Different choices about the grain size or boundaries of routines lead to different predictions about how the system will maintain coherence as it encounters unexpected situations or problems.

Commentary on cognition The extended commentary by ME, though relevant to most sections of the paper, is especially germane here. The reader who has not yet looked at Appendix A is encouraged to look there now.

GD What generalizations can one draw from the existing literature in cognitive and social psychology with regards to "learning modes"? It seems plausible to conjecture that the way people learn depends to some extent on the content of what is learned and the context in which it happens. Is it true? Is there some clear evidence?

MW On the one hand, research in cognition has supported the idea that part of our knowledge is of the tacit (or, as is more commonly said, "implicit") form. This provides a remarkable support to the routine hypothesis, although we need more empirical and experimental evidence, along the lines traced by the Cohen-Bacdayan-Egidi work, that characteristics of individual cognition can be mapped onto collective behavior. On the other hand, added conceptual dimensions have emerged that provide a richer picture and might help us in differentiating more

sharply several distinct forms of organizational cognition. Four such dimensions look particularly useful to me.

Implicit/explicit. The notion of implicit knowledge matches quite closely Polanyi's "tacitness" issue. The cognitive literature has widely debated such issues in recent years, providing a precise characterization of implicitness and some empirical criteria for assessing it. I'd like to report a concise but useful definition by Keith Holyoak and Barbara Spellman: implicit knowledge is knowledge about covariations in the environment, learned by exposure to stimuli exhibiting the covariations, obtainable without attention or awareness, . . . demonstrated by improved performance, . . . but not fully verbalizable, and not fully manipulable, in the sense that it cannot be re-represented explicitly to serve as input to other procedures (Holyoak and Spellman 1993, p. 278)⁴

Distributed/local. The local/distributed distinction has been tightly related to the emergence of the debate on connectionism. A "distributed" representation is one in which there is no single place where a concept or a pattern is stored and there is no explicit symbol for representing it, but instead representations arise from the joint activation of multiple elements (typically, neurons). In a social context, the notion of distributed knowledge can be extended to the idea that a given routine has no place where it is represented (for example, there is no explicit rule, stored in some record, fully describing it) and no one participant could entirely describe it, since the routine arises from the interactions among only partly overlapping, incomplete and inarticulate (possibly non-symbolic) individual knowledges (coordination in a team can easily have this character). Karl Weick's discussion of the concept of "mutual equivalence structures" suggests the distributed nature of knowledge embedded in most basic forms of stable organizational interaction (Weick 1979, pp. 98-103).

Situated/context independent. The idea that social knowledge can be not only distributed among individuals but also distributed among individuals and their environment brings us to the notion of situatedness: knowledge which is tightly coupled to a specific context, because such context acts as external memory and information processor. Research in cognitive anthropology has suggested that many "routinized" competences can be of the situated kind (Lave 1988; Suchman 1987; Hutchins 1995). This point was also present in Nelson and Winter (1983) when they argued that

⁴ Holyoak and Spellman also review some assessment criteria derivable from such a definition.

knowledge stored in human memories is effective only in some . . . organizational context. It typically includes a variety of forms of external memory — files, message boards, manuals, computer memories, magnetic tapes — . . . second, the context includes the physical state of equipment and of the work environment generally . . . finally, the context of information possessed by an individual member is established by the information possessed by all other members (Nelson and Winter 1983, p. 105)

However, I think the implications of such an assumption still have to be fully drawn in evolutionary theories of the firm (see the discussions of the role of artifacts in the following section).

Knowledge intensive/search intensive. Any cognitive system experiences a trade-off between using search or stored knowledge as a basis for its actions (Newell 1990). Newell also labels this as the preparation/deliberation trade-off. He points out that "each situation leading to a response calls for some mix of deliberation and preparation" (Newell 1990, p. 102), showing how a particular system can be located as points in the immediate knowledge (preparation)/ search (deliberation) space. A long tradition in the bounded rationality literature discusses these issues in the organizational context (March and Simon, 1958, is still the basic reference). It still seems to me to provide useful distinctions. In particular, March and Simon suggest that search can also be routinized to a greater or lesser degree, implying that the preparation/deliberation distinction applies also within the boundary of routines rather than defining routines themselves.

Most of these dimensions are already present in the Nelson and Winter chapter on organizational routines; however, they seem to implicitly assume that such dimensions are highly correlated. But this appears to be a quite strong assumption: there's no reason to imagine that a routine lies on the left side of each axis. For example, there are routines that imply highly implicit forms of knowledge, but rely on shared rather than distributed knowledge. Search intensive activities can rely on implicit rules of search while knowledge intensive behaviors may rely on explicit standard operating procedures. Thus we can gain a richer description and useful distinctions by trying to plot our observations on those multiple axes.

Also the "grain size" issue is strongly affected by how we plot actual patterns of action along those dimensions. For example, the definition of "meta-routines" as routines that use routines as inputs to be processed would conflict with the assumption that implicitness (or tacitness) is a fundamental feature of routines; it

is hard to imagine how implicit or distributed rules could be used as inputs of other rules —and in fact cognitive psychology suggests that only explicit, locally represented rules can be processed by higher level rules (see the above cited definition by Holyoak and Spellman).

LM Just a small point (strictly related to my earlier one) on the relation of routines to learning. One possible view is to consider routines as the "fixed point" of a learning process. To make it a little extreme, this view would oppose non-routine behavior (the dynamic path) to routine behavior (the "equilibrium"). Another view (more "imperialist") would contend that at the beginning there were routines which get modified, selected, discarded in the learning process (itself driven by "learning routines") and give rise to other routines. I certainly tend to be sympathetic with the latter view, as the former seems to have a lot in common with some neo-classical approaches to learning.

BC and GD A crucial step when trying to bridge the evidence from cognitive psychology with organizational routines, involves an explicit account of the double nature of routines, both as problem-solving action patterns and as mechanisms of governance and control. The issue is discussed in Coriat and Dosi (1994) with reference to the emergence of Tayloristic and Fordist routines in modern corporations. In brief, what we argue is that the painstaking establishment of a new set of routines (say those associated with Taylorism) involved both a different social distribution of knowledge but also a different distribution of power and control among individuals and social groups. And abundant witness of all that is the record of social conflict which accompanied the diffusion of such routines. There is a general issue here which pertains to the very nature and function of organizations. A good point of departure is March and Simon's definition of organizations as "systems of coordinated actions among individuals and groups whose preferences, informations, interests and knowledge differ. Organization theories describe the delicate conversion of conflict into cooperation" (1993, p.2) In a similar fashion, Nelson and Winter, after having defined routines also as truces among conflicting interests, go on emphasizing that: ". . . here we recognize the divergence of interests among organizations" and that "for some purposes it is important, for some essential, to recognize the conflict of interest contained in and reflected by organizational behaviors" (1982). The analysis of the purely cognitive aspect of organizational routines misses the point that they emerge and operate in

a universe of (at least potential) conflict and diverging interests. Relatedly, one would like to see also experiments that are set in ways which account for both dimensions (for example a modified Cohen-Bacdayan-Egidi experiment involving distributional problems). Note also that the conflict-of-interest dimension is likely to be particularly important when analyzing the emergence of organizational routines (as opposed to their steady-state characteristics).

Moreover, it might not be totally irrelevant to revisit, in the new light of our discussion, older experiments which had a much stronger behavioristic flavor (Milgrom (1974) on obedience to authority). And finally, see the interpretation of the nature of routines by Postrel and Rumelt (1991) who argue that they have a fundamental function of "impulsiveness control". One might not fully agree with such a quasi-Pavlovian view, but they certainly hint at the *lato sensu* political dimension of routines.

In this respect, one could analyze as a sort of extreme archetype the nature of routines in such organizations as armies (where one is likely to find a significantly different balance between their cognitive and their political contents as compared to, for instance, research teams). The general point here is that most organizational routines involve also the legitimation of an asymmetric distribution of power. Hence, together with the cognitive aspect of routines and their characteristics of truces, it is worth exploring their domination elements, whereby, paraphrasing Max Weber, the content of the command becomes a natural and automatic response of "ruled" agents (Weber (1968), Kemp(1993) and also Dosi (1995)).

We're not experts in experimental design, but we think that an item high on the agenda should be a better understanding of how the political relations among actors shape cognitive and behavioral patterns.

GD The issue of automaticity somewhat overlaps with the perspective from which we look at routines, i.e. do we consider them in their "steady state" form or from the point of view of origin? As to the former, I'm convinced that varying degrees of automaticity are an intrinsic feature of routines (maybe one of their evolutionary values). Conversely, it might be more difficult to identify invariant patterns in the process of discovery/development of routines. In (Dosi and Egidi, 1991) we suggested that the problem-solving process is essentially an "intelligent" search whose success yields the discovery of robust rules that can be quasi-idiotically applied thereafter. Note that if this holds, the same action might be a routine for someone and involve a highly deliberative activity for someone

else. So a good deal of what we apparently see as "expert knowledge" (cf. for instance Lane et al., 1995) might in fact be the outcome of highly routinized and automatic rules of evaluation and decision which we simply do not understand, not being experts ourselves.

5. How can evolution be seen in relation to action patterns?

The meeting participants agreed that evolutionary accounts usually exploit a dualism of representation and expression. In the biological instance this occurs as the alternating roles of genotype and phenotype, in which a genome stores the pattern-guidance needed to reproduce its (evaluatable) phenotypic expression, and modifications of the genetic representation are the source of evolvable variation.

In the setting of human action patterns rather than organisms, the location and nature of the representations that guide enactment is not so sharp as the DNA on which so much current biology is focused. Nonetheless, it seems to make sense to distinguish representations (such as learned, skilled responses of machinists) from the action patterns (say, machine assembly steps) that occur when those representations are expressed.

* Some felt this distinction deserved considerable stress and saw representations as something to study in addition to the actions they might generate. (This was heartening for those interested in longitudinal field studies, where past expressions are not observable, but some historically preserved representations may be.) Others, while acknowledging the distinction, found it of lesser priority, preferring to focus on the observable form of the actions as expressed. (This is more consonant with "ethnographic" field observation.)

There was agreement that selection forces in this setting are quite various, beginning with reduction of cognitive burdens on actors, but also including many other powerful organizational forces, (e.g., issues of control, status, economic cost, shared mental models, and general cultural context). Also selection operates at multiple levels, both within firms and in markets where individual firms are born, grow and die.

* There was contention over the relative research priority to be given to cognitive aspects of selection. Some see them as dominant factors, others as important along with many others or as varying in importance across situations.

Commentary

on evolution, representation and expression **RB** For an evolutionary analogy to be applied to something called ‘routines’, a distinction must be preserved between units which are selected and behavior produced by these units which results in their selection. Mere action patterns are not enough, since their focus on description becomes only a restatement of the behavior by which an organization succeeds or fails. Evolution on a biological model assumes that the behavior of an organization, including any capacity to adapt, is constrained by elements that limit the range of actions and responses it can make. When behavior is controlled by such elements, change occurs either by change and reconfiguration of the elements themselves, or by selection (continuation, growth, or collapse) of the organization that contains them.

The evolutionary analogy to organizations will always be relatively weak, since creation, diffusion, and configuration of behavior-producing routines occurs on a continuous basis throughout an organization’s lifetime rather than the fixed inheritance by which an organism lives or dies. Additionally, deliberate design and replication of structures to control and generate behavior is at least attempted by many organizations. Nevertheless, an evolutionary perspective would seem required to temper any presumed influence of rational and deliberate factors.

The challenge is to identify the various forms in which capacities of action (including decision-making) become embedded in organizations, to remain relatively inaccessible or impervious to change and thus a basis only for replication and selection. These forms may or may not resemble anything conventionally regarded as representations, but they may still be regarded as kinds of genetic units on which evolution proceeds. Cognitive learning is certainly one leading candidate for such lock-in of behavior, but patterns of learning and responsibility distributed across an organization, together with explicit policies and procedures, physical artifacts and arrangements, and social and political factors, can all be expected to play significant roles.

MC The distinction between representations of routines and their realization strongly appealed to me. Among a number of nice features, an especially appealing one is that it offers an insight into the difference between routines and standard operating procedures. These terms are often used interchangeably, but the distinction makes it natural to consider SOPs as one kind of representation, a formalized statement of what actions should occur. This separates SOPs from the actions occurring as routines are expressed in context. That is desirable since

there is so much literature showing that real behavior diverges substantially from formalized SOPs. Indeed, "working to rule" has proven an effective labor tactic for bringing an organization to a halt.

MW The representation-expression dualism points to a key ambiguity in the DNA metaphor in Nelson and Winter's book: at which level do routines belong? If one takes the behavioral-level characterization of routines which is implicit in Nelson and Winter, routines clearly belong to the expression-side. This has a deep implication: what is often reproduced is not the routine itself, but some kind of "coded knowledge" which usually implies a mix of linguistic representation, rules, and artifacts (although "apprenticeship" processes may directly replicate the routine itself through forms of learning by examples: see, e.g., Lave and Wenger 1991).

Consider two concrete examples:

"**Best practice**" diffusion has been a major focus of recent organizational management. In a way, it is built on the assumption that actual work experience generates successful routines to be reproduced and diffused - thus it seems to reinforce the "routine" view. But if one looks at the actual reproduction process of such successful routines, one invariably finds a large effort to set up a "technology of replication" that usually implies (a) learning a language within which to code successful routines (b) creating cognitive artifacts that can be diffused (through flow charts and other replicable representations) (c) translating the high level description contained in the cognitive artifact in actual practice, generating a new routine adapted to the new context (see Hutchins and Hazelhurst 1991, for a similar conceptualization of learning in the cultural evolution process). The awareness of this reproduction cycle is well expressed in the fact that managers mostly invest their attention in the "techniques and tools for replication" rather than in actual routines to be replicated, although they know that what finally generates value is the routine (the expression) rather than the code (the representation).

* This is even more apparent in "**business process re-engineering**" (BPR), which in evolutionary terms might be thought of as a sort of organizational genetic engineering. In the BPR case, existing routines are used as materials to be manipulated mainly through deletion and recombination of their elementary components. Again, in order to make possible such manipulation, one has to code the routine in cognitive artifacts (e.g. work-flow graphs) amenable to engineering.

Manipulation then happens at the representation level, and finally its outcomes are "brought back to the field" generating new expressions (routines).

These examples point to the complexities of the process of reproducing patterns of actions, and to its inherent cultural nature, relying on our ability to code experience, to communicate it, and to decode it in ways that may adapt it to new contexts.

GD In a paper that will appear in *Industrial and Corporate Change*, (1996), Benedicte Reynaud distinguishes between rules-to-be-interpreted and (interpreted) rules-to-be-executed. The distinction seems quite similar to the dichotomy between formally stated SOPs and actual behavioral patterns. Indeed, I totally agree that the relation between the two levels, however defined, is a promising area of investigation. However, I would not go as far as drawing an analogy with the genotypical-phenotypical distinction in biology. First, I am not convinced that representations are equivalent to some genotypical level. I would find that at least equally plausible to think of actual action/behavioral patterns as the genotypes, which reproduce via apprenticeship, and acculturation mechanisms within organizations, and imperfect imitation across organizations. Relatedly, it is not implausible to think of representations as sorts of imperfect and mostly ex-post rationalizations of what has been done, should be done, one wished ought to be done . . .

Second, I would suggest a strong "Lamarckian" interpretation of the coupled dynamics between behavioral patterns and their interpretation without any a priori hypothesis on the relative speed of change (how many times do new CEOs redefine SOP's just to find that people continue to do the same things renamed in different ways?)

GD and LM We note also that in the two examples suggested above by MW - the way we interpret them - it is quite clear that the object of replication are routines themselves. However the distinctive features of these examples of replication is that in the social domain there is no straightforward equivalence to sex, mitosis, etc... Therefore - unlike the biological models - we search for coding procedures which might indeed be more transient and conditional on particular contexts than so-called "expressions" themselves.

6. How can simulation contribute to better theory in this area?

A significant recent development in computer science is the bifurcation of the artificial intelligence community into subschools, sometimes labelled "symbolic" and "behavioral". Papers by Brooks (1991, 1991a) and Maes (1993) give the flavor of the behavioral approach. It is strongly connected to robotics, to models of intelligence that can be instantiated in systems that take real actions in a messy real world. It assigns much of the responsibility for keeping track of details about the (changing) world to the world itself, rather than to an elaborate internal model of the world maintained by the software. Performance occurs as a modest number of simple rules interact with each other via their effects on the system's environment.

Roger Burkhart suggested that a number of the features of the "behavioral" movement in AI were strongly consonant with some aspects of research on routine. For example, work on routine often sees artifacts and spatial arrangements as forming an essential part of the organizational memory from which the routine performance is drawn. It is argued that some organizational routines may have a distributed, implicit quality, so that it is not possible to find an account of the whole routine — or even an awareness that there is a routine — located in any one document or any one person's head.

The discussion of simulation modelling echoed important dimensions of the discussion of cognitive aspects of routine. In simulation too there is a question of finding an appropriate blend or balance among tacit and deliberate elements. There was an extended discussion of efforts to model the learning of routine action patterns seen in the experiments conducted by Egidi (1994) and by Cohen and Bacdayan (1994). Luigi Marengo and Michael Cohen both reported that it was possible to build simple systems of condition action rules, such as Holland (1995) classifier systems, that would reproduce a substantial portion of the actual moves made by subjects. However, they both agreed that it was very difficult to build a system that would learn such a set of rules with an amount of experience similar to the amount subjects have in the course of the experiment. (Reasonable competence is achieved after about ten hands of the card game task, on the order of thirty moves by each player.)

Marengo called the group's attention to work by Riolo (1990) on introducing foresight ("lookahead") into classifier systems. It is possible that such a capability would speed the learning process considerably. Bringing in lookahead gives

a model elements of response to both experienced past and expected future. This seems similar to observations that show routines to contain both tacit and consciously deliberative elements.

Lookahead lends itself to a learning dynamic driven by "surprise", deviation of outcomes from expectations. This is increasingly an important factor in models of animal learning (Walker, 1987), and it is significant that Samuel's checker player, one of the greatest achievements of machine modelling of learning, is also driven by surprise. In the Samuel model, as in the experimental data, learning occurs on each action step, not just when a large-scale episode, such as a full game, ends in success or failure.

Cohen pointed out that the experimental task (a card game) was chosen because it drew on substantial knowledge that all subjects could be presumed to have. Times required for computer models to learn are often measured based on random initial knowledge. This is "unfair" to the models. But once one begins to put knowledge into the models there are delicate problems of what to put in, and questions about how to assess the quality of the model's learning. The stable behavior at the end of learning by experimental subjects is relatively homogeneous and easy to model, but the important learning has occurred. However, the transient behavior of early stages of subject learning contains many "mistakes" and false paths, and it would be very hard to get a model to reproduce such sequences accurately.

The discussion of simulation was broken off reluctantly with the shared observation that simulation of (systems of) routines could itself be the topic of a valuable conference. Beyond the issues mentioned above, several other themes were raised during the discussion:

1. Efforts to simulate routines of groups (as in the experimental card game task) suggest that the requirements of coordinating the actions taken by individuals place severe restrictions on how complex individual roles can be within a larger system. It is important to develop simulation tools that allow deeper explorations of such issues.
2. There is variation in what builders of models take as their fundamental objective: to model a routine in its "steady state", for which a hand-crafted system of rules often suffices; or to model the learning process by which the steady state emerged, which clearly requires a commitment to some learning mechanisms of the individuals or an adaptive process acting on the relations

among the individuals. The former approach is generally more manageable than the latter, and correspondingly less satisfying. (Steady state systems of rules tend to be brittle in the face of variation, while the routine functioning they model tends to be fairly robust.)

3. Much of the interest in routines centers on what happens as systems of actors learn action patterns. However, there are a confusing variety of methods available for modelling individual learning. Giovanni Dosi argued strongly for work on the classification of these alternatives into "types", with better understanding of the consequences for a modelling exercise of using one type or another.
4. Handling of time has been a serious deficiency of much previous simulation of organizational action. Many reported phenomena would be unlikely to remain if time were modelled in another, equally plausible way. (Huberman and Glance, 1993) (E.g., Conway's famous "Game of Life" exhibits almost no interesting structure if updating of cell states is not perfectly synchronous (Page 1995).)
5. Theoretical development will be accelerated if it becomes easier to compare simulation models to each other (Axtell et al., 1995), and if there are more opportunities for reuse of computer code — which will itself facilitate comparison.

Thus a computing architecture that would facilitate comparison and reuse could be of considerable value to the field. The Swarm system (Langton, Minar, Burkhart, 1995) appears in demonstrations of early versions to offer many of these needed advantages. If it does not prove sufficient, it might be valuable for a team of organizational researchers to devote themselves to creating another framework that would adequately facilitate the required comparison and sharing.

Commentary on simulation

LM I would mention here a very stimulating idea which is coming from simulation work: the approach which considers organizational routines as emergent properties of the interaction of distributed learning and adaptation processes. This casts also some light on the distributed nature of organizational routines (an aspect on which perhaps we didn't insist enough). (And see related comment on "cognition" above).

GD We did not have much time to discuss in detail the relative values and limitations of modeling instruments and styles. here are some relevant questions:

1) Models of learning by artificial agents have been using a variety of techniques (neural nets, genetic algorithms, genetic programming, classifier systems, . . .). Is there any theoretical reason to choose one or the other. Under what circumstance do they yield similar results?

2) This leads to a more general question of the status of simulation models. I believe none of us is naive enough to think that our models represent how the mind or the organization actually works. But then one faces what we call in (Dosi, Marengo, Bassanini, and Valente, 1995) "weak isomorphism": under which conditions can we say that learning patterns in the model capture something isomorphic to learning individuals and organizations? Are we happy just with the fact that we are able to represent some mechanism of mutation and recombination among elementary cognitive and behavioral traits? And what about the way learning is driven in most models? More generally, thinking of our discussion on simulation models, I remain with a persistent feeling of an underlying epistemological ambiguity. Is the "artificial reality" of our models understood as a sort of analogue to experiments and empirical investigations? I really hope this is not so: simulation models (as well as economists' theorems) are ways of producing coherent and non-intuitive conjectures, but are a far cry from validating the conjectures themselves.

3) A direction of future inquiry that I consider promising is to build highly constructive models a la (Fontana and Buss, 1995) and thus begin to account for the endogenous emergence of higher level entities.

BC and GD It follows from our earlier comments on the "double nature" of routines that we consider it urgent to model an explicit co-evolutionary origin of routines, i.e. to model evolutionary processes nested in multiple selection environments. For instance, think of action/behavioral patterns that evolve as driven by both their problem-solving features and, at the same time, by the conflict among agents over who controls what. If we were able to do that, then we would also be able to compare the ensuing action patterns with those emerging in a "pure problem-solving coordination mode" (such as in the current Cohen Bacdayan Egidi experiment) and, at the opposite end, those emerging in a "pure adversarial no-problem solving mode" (in the spirit of prisoner's dilemma type adaptive learning such as Lindgren (1992)).

MW A greater psychological accuracy of models of individual learning we employ for simulating organizational processes is of course highly desirable. However, in my view there are other directions in which substantial progress is at least as much needed.

Direct vs. vicarious learning. Most models of machine learning of the kind exploited in organizational simulation are basically models of direct, experiential learning. However, a substantial part of the rate of progress in any social system is related to processes of diffusion of knowledge, which imply vicarious learning, or the ability to acquire knowledge from other people's experience (March, et al., 1991). Incorporating explicitly vicarious learning and diffusion mechanisms in our simulations might generate significant returns. In a pioneering paper, Hutchins and Hazlehurst (1991) have suggested some basic kinds of learning that a model of such a process needs to include: (a) direct learning of an environmental regularity; (b) mediated learning about an environmental regularity from the structure of artifactual descriptions of it; and (c) learning a language that will permit a mapping (in both directions) between the structure of the natural regularity and the structure of artifactual description of it. I would add to that list that, in organizational contexts, some attention needs to be devoted to how structural features constrain or bias the way individual experience is diffused; organizational learning sometimes implies evolving those structural aspects, giving rise to a sort of "architectural learning" (see Henderson, 1992, on organizational structures, the "architectural knowledge" they embody, and organizational learning processes).

Learning at the organizational level. How well do machine learning models represent organizational vs individual level processes of learning? For example, one might argue that "genetically inspired" models like genetic algorithms or genetic programming may leverage on processes which have little psychological plausibility at the individual level, but can be adapted to better match some organizational level dynamics (for example, Cohen (1981) advances important remarks on organizational systems as "parallel heuristics"). Furthermore, it would be desirable to model learning processes occurring simultaneously both at the individual and at the organizational level - a kind of multiple-level, nested dynamics for which the SWARM system seems to provide promising modelling tools.

Learning and evolution. Finally, understanding interactions between learning and evolution should be a primary concern for an evolutionary theory of the firm which claims that learning plays a fundamental role in shaping organiza-

tional life. There have been important advances in simulating these processes in the field of artificial life (see for example Ackley and Littman 1991; Parisi et al., 1995). These developments significantly parallel the rising interest within organization theory in complementarities between learning and selection (Levinthal 1990), although much more structure should be put into the relationships among learning entities to provide significant insight in organizational matters. Again, the SWARM system may prove useful in supporting such efforts to simulate interactions between learning and evolution in organizational contexts, thanks to its ability to manage simultaneously multiple dynamics occurring at different levels and with different clocks.

7. What are examples of various approaches to empirical research that reveal key problems or results?

We explored, more briefly than we had hoped, the fit between alternate empirical styles and the various substantive concerns. Among the observational methods reviewed were experiments, ethnographic field studies, longitudinal empirical studies, and statistical/econometric approaches.

Discussion of experiments occupied considerable attention, probably for the simple reason that several experimentalists were present. To say the least, they were not a representative sample of all students of organizational routine. However, they were able to contribute especially detailed example material based on controlled multiple observations, and this rare asset was useful to the discussions.

There was general agreement that experiments can show gross evidence for routinization very clearly, and can provide data that supports "micro-genetic" analysis (Siegler and Crowley, 1991). This observation incorporates a distinction offered by Massimo Warglien, between the detailed content of an actual routine, which may be hard to measure, and the overall level of "routineness" in the action, for which one may be able to construct workable indicators.

Explaining in detail the exact routines that emerge for a particular experimental work group can be very hard, however. This is so even when one has the advantages of extremely detailed data on actions and exact knowledge of the task environment. Still, it has been possible to define broad indicators of routinized action and measure their presence.

With respect to field studies, there is a possibility of exploiting this same distinction, looking for indicators of routineness, rather than — or along with — documenting actual contents of routines. Learning curve studies (Argote and Epple, 1990) provide a good example of evidence at the level of routinization without detailing the contents of the underlying routines. Others might include: implicit versus explicit representation of task requirements; methods of teaching skills to new members (formal schooling vs apprenticeship, Orr 1990); and amounts of attention available for analysis of other, more novel tasks.

A further observation on field studies of the content of routine — not new, but very telling to those who have done it — was that data collection was very expensive, and progress requires "an army of ethnographers." The work generates richly suggestive results, and accumulation of such accounts, as in anthropology, provides essential grist for theory development. But the "costs" of observations are also at the extraordinarily high levels typical of anthropology.

Thus while field observation will always be of major value, the high costs there suggest that some effort should go into other approaches. An alternative is what were called longitudinal empirical studies: work on routine elements of action in an organization over years — or even decades — of existence. Here the research method is not direct observation, but rather reconstruction from the organization's written and oral histories, and perhaps from preserved artifacts such as products or machines. Thus in microelectronic firms, different devices may embody the same design competence, and they are usually classified in a same "process family" despite being destined for very different usages. Therefore by tracking the rise and fall of such process families, one can learn a lot about the diffusion of specific design competences.

Finally, there is an approach that tries to extract evidence for routines, or other recurring patterns, via statistical analysis. Examples, would include work on typical rates of mark-up or R&D investment, characterized by industries and eras (beginning with Cyert and March, 1963).

Commentary

on MC observation Video data from field sites presents many interesting possibilities for research on routine. At the moment, it is very time-consuming to analyze, but technical improvements relying on computers are beginning to emerge. (Suchman and Trigg, 1991; Carnegie-Mellon automatic segmentation project, <http://www.informedia.cs.cmu.edu/>) And the accumulation of video data sets will

eventually offer possibilities for secondary and comparative analysis of field observations, which will be a major advance in our research options.

MW Two remarks on "routineness" and longitudinal studies:

On "routineness": an advantage of measuring routineness rather than the actual routinized behaviors is that it puts less stringent constraints on our modelling strategies when trying to simulate learning processes. While machine learning models can be poor in reproducing observed learned behaviors, many of them generate patterns of routinization that can decently match experimental or field observations (such as aggregate learning curves). In many cases, this can be enough for reasonable accuracy to be reached when the actual behavior is less relevant than its impact on organizational performance.

On "longitudinal studies": a merit of longitudinal analysis carried over decades is not only that it sometimes avoids ethnographic costs (although it raises dramatically the costs of archival research), but also that it allows us to look at evolutionary dynamics (which are by their nature long term ones) in ways that more accurate but shorter term ethnographic studies wouldn't allow. In particular, longitudinal studies may allow us to go through serious econometric (or econobiometric?) analysis and explore long terms patterns of innovation and diffusion of organizational replicators. If evolutionary theories of the firm want to take up the challenge of the population ecology of organizations, this seems to me an unavoidable step. Then the usefulness of the representation approach to categorizing suggested in the evolution section becomes apparent. Patterns of action and their diffusion are hardly trackable and measurable over long time horizons, while at least some kind of representations are usually recorded in organizational archives. They may not tell us all the story about routines, but like fossil evidence, they supply important cues for understanding organizational evolution (especially if we can provide an understanding of the role of such elements in the cycle of the reproduction of patterns of action). This point is further reinforced if we embrace the point of view that representations (like artifacts), rather than expressions, are what actually gets reproduced in organizational processes - in this case, representations might be not just a convenient substitute for closer observations of actual behaviors, but rather a more appropriate object of observation.

8. How might 'routine' be defined?

An agreed-upon definition was not seen as essential to progress in the meeting.

However, considering a proposal did clarify many dimensions of variation in the group, and the result may be useful. The discussion produced some striking agreement along with the sharpest disagreements of the workshop.

The proposal discussed was Roger Burkhart's. As amended from the floor, and after minor rearrangement in preparing this note, it was:

A routine is an executable *capability* for repeated performance in some *context* that been *learned* by an organization in response to *selective pressures*.

The bold-faced terms are the key "slots" of the definition. Any researcher using it would have to bring it into closer contact with a specific problem of observation or theory by giving more specialized accounts of the key terms. In its general form it runs a little risk of including too many things, but it does have the virtue of showing how both a manufacturing competence and a computer program could be thought of together as routines. Each of the slots is discussed below, giving some expanded sense of its intended meaning and examples of how it could be specialized.

A *capability* was characterized as the capacity to generate action, to guide or direct an unfolding action sequence, that has been stored in some localized or distributed form. (E.g., the ability of a group of factory workers to assemble an engine, of laboratory technicians to recognize cancerous cells, or of a computer package to return a Fourier analysis of a data stream.)

Execution was conceived as possible only in some *context* within which action accomplishes some transformation. Context was seen as a powerful form of "external memory" or representation of portions of the routine (e.g. the arrangement of tools and machines on a factory floor), and as a source of necessary inputs to actions. It provides the natural locus of attention for lines of research focused on the role of artifacts.

The characterization of routine as *learned* was noted to imply possibilities — but not certainty — of the tacitness and automaticity which have hallmarks of routine in many empirical studies. The term is used broadly to cover any of a wide variety of processes that could alter the probabilities of future enactments of the capability. (E.g., individual actors in a team might be learning motor skills necessary for a particular team performance, learning might be occurring at the organizational level via changes in personnel or in training practices that change skill mix, improved computer programs might be generated by adaptive tuning

methods, . . .) As with the other slots, the breadth of the term ‘learning’ requires a researcher being guided by the trial definition to provide a more specific account.

Selective pressures is a broad term meant to indicate a wide variety of forces that could operate to make action sequences more or less likely. (E.g. see above.)

* This was one point of major divergence in the conversations, with some arguing that cognitive loads should receive clear priority as the forces most responsible for change and/or stability, while others argued for equal significance of many other factors such as personal incentives, status concerns, or organizational culture.

Another difficulty noted for this term is the absence, in the organizational case, of a population of directly competing action sequences that might be expected by analogy to the biological case.

Some further notes on the definitional debate:

The proposed definition endorsed a view of learning and routine as very tightly coupled to each other, and hence excluded one-off performances from the category. This seemed to be generally accepted.

The proposal also acknowledged the duality of representation and expression, in an effort to situate routine in an evolutionary view of change. There was greater debate about this. * Attention to representations was seen by some as leading to research difficulties while others saw it as a welcome research opportunity. The closely related concern over distinguishing action from behavior was clearly evoked in discussing the proposal. Those advocating a "behavioral" research focus again saw dangers in bringing intent into a central role and felt it increased ambiguity of observations, while proponents of the action-not-just-behavior line argued that action viewpoint was needed to reduce ambiguity.

Winter contributed what he called his “radical” view that achieving maximum tightness in key definitions may sometimes inhibit progress.

Commentary

on definition MC Looking at our discussion of ambiguity of observations and of action vs. behavior, one might speculate that in research ambiguity of observation is conserved. What we debate is which rug to sweep it under . . .

Bibliography

Ackley D., Littman M. (1991) 'Interactions Between Learning and Evolution', in C. Langton et al. (eds) *Artificial Life II*, Reading (Mass): Addison Wesley.

Allais M. (1953) Le comportement de l' homme rationel devant le risque: Critique del postulats et axiomes de l'Ecole Americaine, *Econometrica* 21, pp. 503-46.

Argote, Linda, and Dennis Epple (1990) Learning Curves in Manufacturing. *Science*, vol 247, pp 920-924 .

Axtell, Robert, Robert Axelrod, Joshua Epstein, and Michael Cohen (1995) Aligning Simulation Models: A Case Study and Results, (Working Paper 95-07-065). Santa Fe Institute. (Also, forthcoming *Journal of Computational and Mathematical Organization Theory*.)

Brooks, Rodney A., (1991), New Approaches to Robotics, *Science*, vol. 253, pp. 1227-1232, 13 September.

Brooks, Rodney A. (1991a), Intelligence without reason, (AI Memo #1293). Cambridge, MA: MIT Artificial Intelligence Laboratory, April 1991.

Bruner, J. S., and Leo Postman (1949), "On the Perception of Incongruity: A Paradigm," *Journal of Personality*, XVIII, 206-23.

Chaitin, G. J. (1995) The Berry Paradox, *Complexity*, vol. 1, #1, pp. 26-30.

Church A. (1936) "An unsolvable problem of elementary number theory", *American Journal of Math.* 58, pp. 345-363.

Cohen Michael D. (1981) The Power of Parallel Thinking, in *Journal of Economic Behavior and Organization*, 2: 285-386.

Cohen, Michael D. and Paul Bacdayan (1994), Organizational Routines are Stored as Procedural Memory: Evidence from a Laboratory Study. *Organization Science*, vol 5, #4, pp. 554-568.

Coriat, B and G. Dosi (1994), Learning how to govern and learning how to solve problems. On the co-evolution of competences, conflicts and organizational routines, Laxenburg, Austria, IIASA working paper.

Cyert, Richard M. and James G. March (1963, 1992), *A behavioral theory of the firm*, 2nd edition. Cambridge, MA: Blackwell Publishers, 1992.

Dosi, G. and M. Egidi (1991), Substantive and procedural rationality. An exploration of economic behavior under uncertainty, *Journal of Evolutionary Economics*.

Dosi, G., L. Marengo, A. Bassanini and M. Valente (1995), "Norms as Emergent Properties of Adaptive Learning", Laxenburg, Austria, IIASA working paper.

Dosi, G (1984), *Technical Change and Industrial Transformation*, New York: St. Martins Press.

Dosi, G. (1995), Hierarchies, markets and power: some foundational issues on the nature of contemporary economic organizations, *Industrial and Corporate Change*, pp.1-9.

Egidi M. (1994) "Routines, hierarchies of problems , procedural behavior: some evidence from experiments" IIASA working Paper WP-94-58 July. To appear in K. Arrow et al. (editors) *The Rational Foundations of Economic Behavior*, MacMillan, in Press.

Fontana W. and L. Buss (1994), The arrival of the fittest: toward a theory of biological organization, *Bulletin of Mathematical Biology*, PP. 1-64.

Hayek F. A. (1952) *The Sensory Order. An Inquiry into the Foundations of Theoretical Psychology*, London: Routledge and Kegan Paul.

Henderson, Rebecca M., (1992), Technological Change and the Management of Architectural Knowledge. In Kochan, Thomas A. and Michael Useem (eds), *Transforming Organizations*, pp. 118-131. Oxford University Press, 1992.

Hirschman A. O. (1984) "Against Parsimony: three Easy ways of Complicating Some Categories of Economic Discourse" in *Bulletin of the American Academy of Arts and Science*.

Holland, John H. (1995), *Hidden Order: How Adaptation Builds Complexity*. Reading MA: Addison-Wesley.

Holyoak, Keith J. and Barbara A. Spellman (1993) 'Thinking', in *Annual Review of Psychology*, 44: 265-315.

Huberman, B. and N. and Glance (1993), Evolutionary Games and Computer Simulations, *Proc. National Acad. Sci. USA*, 90:7716–7718 (August).

Hutchins Edwin (1995), *Cognition in the Wild* Cambridge, Mass. : MIT Press.

Hutchins, Edwin and Brian Hazelhurst (1990), Learning in the Cultural Process, in *Artificial Life II: Studies in the Sciences of Complexity*, C. G. Langton et al. (eds). Reading MA: Addison-Wesley, 1990.

Kahneman D., Tversky A. (1986) Rational choice and the Framing of Decisions in Hogart R. M., Reder M. W. *Rational choice - The Contrast between Economics and Psychology*, Chicago, The University of Chicago Press.

Kemp, D. A. (1993), Authority and public policy, in R. Radner (ed.) *An heretical heir of the enlightenment: politics, policy and science in the work of Charles E. Lindblom*. Boulder Co., Westview Press.

Kuhn, Thomas (1962), *The Structure of Scientific Revolutions*, Chicago University Press, Chicago.

Lakoff, George (1987), *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*. Chicago: Chicago U. Press.

Langton, Chris, Nelson Minar and Roger Burkhart (1995), The Swarm Simulation System: A tool for studying complex systems. Draft working paper, Santa Fe Institute, <http://www.santafe.edu/projects/swarm/swarmdoc/swarmdoc.html>.

Lave, Jean (1988) *Cognition in Practice*. Cambridge, MA: Cambridge University Press.

Langton, C. G. et al. eds. (1992), *Artificial Life II*, Redwood City, Addison Wesley.

Lave, Jean and Etienne Wenger (1991) *Situated Learning*. Cambridge, MA: Cambridge University Press.

Levinthal, Daniel (1990) 'Organizational adaptation and environmental Selection - Interrelated Processes of Change'. *Organization Science*, 2:140-145.

Lindgren, R. (1992), Evolutionary phenomena in simple dynamics of prisoner dilemma, in Langton, et al. 1992.

Luchins A. S (1942) "Mechanization in Problem-Solving" , *Psychological Monograph* 54, pp. 1-95.

Luchins A. S. and Luchins E. H (1950) New experimental Attempts in Preventing Mechanization in Problem-Solving, *The Journal of General Psychology*, 42, pp.279-291.

Maes, Pattie, (1993), Behavior-Based Artificial Intelligence. *Proceedings of the 2nd Conference on Adaptive Behavior*, Cambridge MA: MIT Press.

March, James G., and Herbert Simon (1958) *Organizations*. New York: Wiley.

March J. G. and H. Simon (1993) *Organizations*, 2nd ed., Oxford, Blackwell.

March J. G. (1992), *A Primer on Decision-making: how decisions happen*. NY: The Free Press.

March, J. G., L. Sproull, and M. Tanuz (1991), Learning from Samples of One or Fewer, *Organization Science*, 2:1–13.

Marengo, L. (1992), "Coordination and Organizational Learning in the Firm", *Journal of Evolutionary Economics*, vol. 2, pp. 313-326.

Milgram S. (1974), *Obedience to authority: An experimental view*, London, Tavistock Institute.

Nelson, Richard and Sidney Winter (1982), *An Evolutionary Theory of Economic Change*. Cambridge MA: Belknap Press.

Newell, Allen (1990) *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Orr, Julian, (1990), Sharing Knowledge, Celebrating Identity: War Stories and Community Memory in a Service Culture. In Middleton, D. S. and D. Edwards (eds.), *Collective Remembering: Memory in Society*. Beverly Hills, CA: Sage Publications, 1990.

Page, Scott (1995) , A Note on Incentive Based Asynchronous Updating in Cellular Automata. Unpublished, CalTech.

Parisi, Domenico and Stefano Nolfi (1995) 'The Influence of Learning on Evolution'. In R. K. Belew and M. Mitchell (eds) *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Reading (Mass): Addison Wesley,

Pentland, Brian (1994), Reconceptualizing Routines: Grammatical Models of the Sequential Structure of Work Processes. *Administrative Science Quarterly*.

Polanyi M. (1958) *Personal Knowledge: Towards a Post-Critical Philosophy*, London: Routledge and Kegan.

Postrel S. and R. Rumelt (1991), Incentives, routines and self-command, *Industrial and Corporate Change*.

Riolo, R. L. (1990), Lookahead Planning and Latent Learning in a Classifier System. In Jean-Arcady Meyer and Stewart Wilson (eds), *From Animals to Animats*. Cambridge, MA: MIT Press, 1990.

Samuel, Arthur L. (1959), Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, vol. 3 pp.210-229.

Siegler, Robert S. and Kevin Crowley (1991), The Microgenetic Method: A direct means for studying cognitive development. *American Psychologist*, vol 46, pp. 606-620.

Simon, Herbert A. (1981), *The Sciences of the Artificial*, 2nd edition. Cambridge, MA: The MIT Press.

Suchman, Lucy A. (1987) *Plans and Situated Action*. Cambridge, MA: Cambridge University Press.

Suchman L. and R. Trigg (1991), Understanding Practice: Video as a Medium for Reflection and Design. In J. Greenbaum and M. Kyng (eds.) *Design at Work: Cooperative Design of Computer Systems*, Hillsdale, NJ, 1991.

Turing, A. M. (1936), "On Computable Numbers, with an Application to the Entscheidungs-problem", Proceedings of the London Mathematical Society, Vol XLIII, pp 230-265.

Walker, Stephen (1987), *Animal Learning*. New York: Routledge & Kegan Paul.

Weber M (1968) *Economy and Society*, Berkeley, University of California Press.

Weick, Karl (1979), *The Social Psychology of Organizing*, 2nd edition. Reading, MA: Addison-Wesley.

Weisberg R. (1980) *Memory: Thought and Behavior* , New York: Oxford University Press.

Appendix A : "Accountants and Termites", a longer comment by Massimo Egidi

One of the issues raised during the SFI meeting was to analyze the most relevant properties of routinized behaviors. While there is a largely dissonant use of the magic word 'routine' and other related keywords in different disciplines — Economics, Theory of Organization, Artificial Intelligence — and this creates confusion, I feel it useful to elaborate a small lexicon of synonyms and opposites. What follows is a reconsidered report of what I tried to say in Santa Fe, plus the lexicon.

1. Routine as set of rules

Within the theory of computation, tradition considers "routines" to be synonymous with 'programs', i.e. list of instruction in a (artificial) language. Turing wrote his celebrated article on "Computable Numbers", proposing his computing machine, in the same year - 1936 - in which Alonzo Church was completing his work on Lambda Calculus, and it soon became clear that the two systems were equivalent.

Since these two progenitors, Turing machines and Lambda Calculus, a large variety of mathematical representations of "computing machines" have been proposed: URM machines, Post systems, Production Systems, etc. All of them have been proved equivalent to a Turing machine.

Therefore a consolidated tradition exists which considers routines to be computable programs, which can be represented by a set of condition-action rules (i.e. Production Systems) or an equivalent computing device. It is obvious that in the social sciences strong resistance can arise against considering the humans involved in a routinized activity as pure automata executing a set of condition-action rules: nobody likes being compared to a Pavlovian dog. However, before looking at the differences between the behavior of machines and men, I want to emphasize that this idea was implicit in Turing's scientific program, and imbued mathematical culture at the beginning of the 20th century, when the Hilbert program was dominant: and Hilbert's idea was, in fact, to mechanize all mathematical propositions via an axiomatic method in order to be able mechanically to decide the "truth" or "falsity" of any given proposition.

Turing's approach is implicitly based on a project to construct "thinking machines" able to substitute for humans in their fundamental mental abilities. I propose to see to what extent this idea can be used to explain the features of human decisions and behaviors within organizations, my aim being to discover some of desirable features of the notion of routine in organizational and economic contexts.

Let me start by reversing the Turing approach, i.e. by considering an experiment where men imitate the behavior of a Turing machine: many individuals execute mechanically a list of instructions, following the orders of some authority which coordinates their actions to realize a given task. Since the execution of the same list of instructions can be repeated many times, when the appropriate conditions arise, if individuals mechanically follow the instructions, the same sequence of behaviors will arise corresponding to the same set of conditions.

Therefore an external observer would describe the team's behavior as "routinized", because he observes that the same set of actions is performed in response to the same conditions.

One point to clarify is that when we observe a team of workers performing a task in a routinized way, we observe a set of coordinated actions repeatedly realized over time by the team, but we do not directly observe motivations and rules. This distinction is crucial, even beyond our experiment of imitation. We follow the displacement of a set of rule-based behaviors as a dynamic process over time, at the immediate observational level : at this level, routinized behaviors can be interpreted in various ways, by referring to different models of human action. I will minimally assume that behaviors are the consequence of mental models and more stringently (but provisionally) that routinized behaviors are the outcome of the execution of set of condition- action rules, stored in individuals' long-term memories. Is this purely a caricature of human behavior or does some family tie exist between the mathematical approach and real routinized human behaviors?

2. Observational levels

Before answering this question, some preliminary problems relate to the observability and refutability of the models of mind we assume. If we suppose individuals to be capable of forming an internal mental representation of the situation, based on symbols and their manipulation, we can find experiments to confirm or refute our approach not only at the behavioral level but also - and largely - at the level of mental models, with a sophisticated use of the methods of experimental psychology.

I emphasize that this level — interiorization of rules, memorization, etc. — is not unobservable. The problems in classical Expected Utility theory began when, with the experiments of Maurice Allais, and later of Kahnemann and Tversky, it became possible directly to confirm or refute a large number of propositions claimed by the theory . The same holds for the empirical analysis of other mental activities involved in human behavior, like reasoning, problem solving etc. Therefore, the more the methods of experimental psychology advance, the easier it is to verify on empirical grounds whether routinized behaviors in the real world are or not are rule-based, and more generally to what extent individuals act on the basis of an internal model of the world.

Cohen and Bacdayan's experimental model permits to use a large array of sophisticated statistical methods to conduct analysis either at the behavioral or at the mental level : it is in fact possible to examine the sequences of collective actions over time (behavioral level) or to operate at the level of subjective beliefs , models and expectations of the players.

3. Rules and routinized behaviors

If we assume as a starting point that routinized behaviors can be rule-based, we must carefully distinguish between the sequence of collective actions realized over time, and the set of rules which generate this sequence when applied by a team of individuals. It is quite a common habit to call the set of actions or behaviors 'routine', therefore confusing the two levels of analysis. Following the computational tradition we should consider 'routines' the algorithms (the set of rules) and 'routinized behaviors' the sequence of action generated by the realization of the rules. But this is a purely nominal assumption, which I am ready to abandon provided that a different definition will preserve distinction between the two levels. Anyway, to avoid confusion in the following discussion I will not use the term 'routine' in the context of humans behaviors and mental models, reserving to this term the usual meaning of theory of computation and referring only to routinized behaviors in order to indicate complex, collective, repetitive patterns of action. Let me clarify some further elements of the vocabulary used.

- i. A "rule" - here used with the limited meaning of condition-action rule - is a relationship which allows individuals to trigger an action when a condition is realized; the triggering can be automatic, like a stimulus-response mechanism, or deliberate and conscious.

- ii. In a dynamic process of coordination among many individuals, at any time the system configuration is composed by the events of the external environment plus the actors' behaviors. This configuration is the background from which "conditions" emerge and are recognized by the actors. Once again the recognition of a condition may be automatic, like a stimulus-response mechanism, or deliberate and conscious.
- iii. Domain of applicability of condition-action rules: one important point is that in computation theory the set of rules is supposed to cover all contingencies, i.e. all the possible conditions with which individuals will have to cope.
- iv. A further point is that rules are defined a-temporally, i.e. they prescribe the action to be triggered in relation to a condition (or a set of conditions) which is defined independently of a specific date and of the system's configuration. This is a quite complex argument which relates to the "compressibility" and "representability" of knowledge and information, and I shall deal with the question only briefly: observe that a set of condition-action rules which compose a program must be completely specified if we want a machine to be able to execute it. This means that every configuration that the system can exhibit must match a "condition" in the machine's list of rules. The same applies to fully routinized actors. In consequence, identifying an appropriate sets of rules is a matter of high intelligence, because it require one to "reduce" all the complex features of the process to a few, essential traits which are described in the condition-action set. The possibility of reducing the system's features to few essential rules is not given in advance, but depends upon the nature of the problem involved. Chess is a good example, because it shows very modest compressibility: a winning strategy, which exists, cannot be represented with a reasonably limited number of rules. I suggest Chaitin's (1995) theory of uncertainty as a starting point for serious exploration of this issue.

4. The problem of identical action

I have pointed out that at the most elementary level of observation, routinization is recognized by the observer when to the "same environmental conditions" corresponds the "same sequence of behaviors" ; therefore the obvious problem is to define to what extent a sequence of behaviors can be defined as "the same". This question involves very subtle epistemological issues which I cannot raise here. I shall limit my discussion to pointing out that most of the problems in-

volved are soluble if we accept the idea of rule-based actions: if we can prove that individuals use the same set of condition-action rules during their activity (by a protocol analysis, for example), then we can identify the routinized behaviors as "the same" , even though we observe behavioral variations over time. The same set of rules can in fact generate "similar" blocks of actions if some of the condition are "similar" but not equal.

In Cohen and Bacdayan's game (1994) this aspect is quite clear: assume that each player adopts a stable sets of rules, compatible with the rules adopted by the other, and applicable to every valid game configuration. In different runs, when players must find a card which is essential to achieve the solution, even if they apply the same set of rules, they can perform sequences of actions which may be quite different, depending upon the distribution of the covered cards. In this case I would not hesitate to consider the routinized behavior as the same, because all the reasonable elements of an "identical action" are satisfied: the behaviors over time are partially different, but the generating mechanism is the same. As a counter proof, examine what would happen if two Turing machines played the game: the sequence of actions would differ accordingly to the distribution of the covered cards, even though the machines were executing the same search algorithm .

5. What happens in real organizations and with men instead of machines?

I now turn to the original question, i.e. exploration of the differences between the behaviors of men and machines in the realization of a repetitive task.

- i. When the actors are men, the rules that any individual adopts are quasi-exhaustive, i.e. they cover a large part of the possible contingencies but not necessarily all of them. Therefore a set of rules does not prescribe the action to be executed for every condition which may arise during the collective process.

"Action" and "condition" are terms which normally refer to a class of possible actions or conditions: for example the action "Search for card 2 of hearts " can be realized in many different ways, i.e. refers to a class of possible actions. To realize the condition-action a machine must have a detailed search procedure, i.e. a fully exhaustive set of condition-actions to execute the program: therefore an executable machine program must recursively contain sub-procedures until all possible contingencies are covered.

Quite different is the situation involving individuals, who normally do not memorize all the specific and detailed set of rules, because they are able to recreate parts of them in many different ways. Therefore we must accept that a micro-learning

activity is normally at work, and consider situations in which an individual's activity is fully routinized, i.e. covers all possible contingencies with memorized set of rules, as an extreme case . In these extreme situations behaviors are fully mechanized, and the learning process is inactive.

- ii. Call "coordination rules" the rules which embody mutual relationships, i.e. prescribe actions which are compatible with the actions performed by the partners. Coordination rules are largely internalized by individuals, who do not need detailed and specific orders to realize the coordinate task. If an unexpected condition arises, there follows either a conflict among rules or the lack of rules, and coordination fails. A machine cannot work in these conditions, while human reaction is typically to find a solution, i.e. start learning and exploring in the space of rules.
- iii. This suggests that a fully routinized activity is possible, or at least more likely to happen, when tasks are performed by isolated individuals. In coordination processes, in fact, purely routinized behaviors would require taking into account the reaction to the other's actions, including errors and conflicts, which require deliberation and learning: interactions normally are exponentially increasing with the number of participants. Therefore to be covered by automatic rules, individuals would have to memorize an incredibly large number of rules: a situation very similar to chess. This aspect of the problem suggests that it is most likely to observe a purely automatized behavior - to which can correspond an automatized thinking (Weisberg (1980))- at the level of isolated individual action, while this is quite rare and unstable in a widely coordinated action context.
- iv. Purely routinized collective behaviors are rather difficult to realize because, to cover any possible contingency, they require a huge set of rules governing the interactions among actors: this is the equivalent of complete markets with rational expectations in general equilibrium theory.

To activate a huge, possibly hierarchical set of rules , a large set of conditions must be recognized, and a great deal of information used and processed. Therefore a complex activity of computation is required , which is difficult to reduce to an automatic execution of rules. I claim that there are limits to the complexity of the set of rules that can be activated by boundedly rational actors. Beyond this threshold, reasoning cannot be substituted for by purely automatic behavior. Chess seems exactly to configure situations where the threshold is surpassed and players cannot reduce their activities to the execution of a set of condition-action rules.

6. Routinized behaviors

For the four reasons above I consider still unsurpassed the definition of routinized behaviors provided by March and Simon, i.e.

We will regard a set of activities as routinized, [then,] to the degree that choice has been simplified by the development of a fixed response to defined stimuli. If search has been eliminated, but a choice remains in the form of clearly defined and systematic computing routine, we will say that the activities are routinized (March and Simon, (1958), p. 142).

Here computing routine is used — I believe — in the meaning given to it by the theory of computation, as synonymous with the Turing machine or Algorithm. Following March and Simon, I shall use the term ‘routinized behaviors’ to mean behaviors which emerge as rule-based actions. As the psychological literature has emphasized, these behaviors are based on "routinized thinking", which allows individuals to save on mental efforts; an experimental way to verify whether thinking is automatized is to check if subjects are able to perform some complex mental activity, like problem solving, while they are executing a different, repetitive task. The studies on the mechanization of thinking, the so called "Einstellung effect", have a long tradition in psychology (Luchins (1942,1950), Weisberg (1980)).

If we take into account the observations on points i-iv , the above definition seems to fit accurately with most of the features assumed in Nelson and Winter’s description of routines: automaticity, triggering of actions, complexity.

The features which seem not to be grasped by March and Simon’s definition — with the previous remarks — are tacitness and awareness.

7. Awareness, division of knowledge, repair of patterns of behaviors

Can behaviors be considered routinized independently from the degree of awareness of the individuals and irrespective of the level of tacitness of the knowledge involved? To better frame the problem let us consider an extreme case of organizational shape, a top-down hierarchical organization whose members execute orders in deliberate, conscious way. I assume a hierarchical, authority guided, centrally designed activity, such as were, historically, the accountants offices in large business institutions and banks before the advent of the computer era. The everyday activity performed by large teams of accountants involved a huge computing and accounting activity, mostly deliberate. Moreover the

accounting process itself was the outcome of a deliberate design realized by a team of experts.

Can we avoid considering this process — where highly repetitive tasks were performed by the same team of employees in the same way everyday — as an example of routinized behavior? If we consider unawareness as a crucial feature of routinized behaviors, we should be tempted to say: no. The accounting process seems to be a perfectly deliberate, goal-oriented activity, realized by individuals clearly conscious of the goals they have to achieve and the means they use. The means are essentially the mathematical and financial algorithms created by "experts" (from Pythagoras and Euclid to the most advanced financial mathematics of the time) which they carefully applied.

I subscribe to this opinion, and suggest that even this kind of activity cannot be reduced to a perfectly deliberate and conscious execution of instructions. The recursive character of knowledge suggests we be very prudent with the words "awareness", and "consciousness": in fact accountants — as all humans — are able to mentally dominate only a limited area of the knowledge, even of the specific knowledge required to perform their regular everyday task.

To be persuaded, think of some basic algorithms used by the accountants: counting, summing and multiplying. They use these algorithms to compute, but they have to know only the domain of applicability of the algorithm and the rules of application: which is a minimal and very bounded part of the involved knowledge. It is easy to show some elementary cases of these limitations. For example: everybody knows how to sum two numbers, but a very restricted number of persons know why the rules which compose the summing algorithm work, if there are different algorithms to sum etc. . In a sudden world catastrophe, the surviving accountants presumably should not be able to reconstruct the theory underlying the mathematical algorithms they use in the everyday life. (see Asimov stories). Individuals directly involved in a routinized activity may not be able to repair the pattern of cooperation to which they adapt if an eventual failure happens at a very deep level; a larger set of individuals, involved in the relevant knowledge, must intervene. This is the outcome of a characteristic feature of knowledge: its partial separability or "orthogonality". Part of knowledge can be used in a perfectly independent way from other parts which are indirectly involved, i.e., in terms of problem solving approach, there exist sub-spaces of problems which are orthogonal and independently solvable.

Therefore, even in the world of deliberate actions, where a symbolic representation is developed and a common language established, knowledge is never fully transparent to the limited minds of a single individual. In consequence, we have to admit the incomplete character of knowledge involved in the so-called "deliberate" activities.

Perfect deliberativeness cannot exist and human behaviors are based on incomplete knowledge and partially opaque deliberations; this is also the reason for the cognitive delegation which characterizes the division of knowledge, and the origin of the conjectural and incomplete character of top-down hierarchical planning.

Now consider briefly the situation at the opposite extreme of the accountants example, i.e. assume that it is possible that a routinized pattern of behavior arises even if a common language does not exist, individuals have an opaque internal representation of the problem. Here we move from accountants to forms of organization where a "conscious social mind", to cite Hayek, does not exist. Individuals act in a completely automatized way, as in a termite nest. But assume that Mister Brown, a worker who every day accomplishes his routinized work, this morning has a bad headache, does not respond "normally" to some stimuli or does not recognize correctly the familiar patterns, and consequently does not trigger the correct action: a problem of mis-coordination arises, because Brown introduces a noise in the system. Does there exist a set of rules, memorized by all the workers, to solve this kind of problem? It is easy to understand that this can happen only to a limited extent, otherwise we ought to admit that there are no limits to human memory.

In general, the situation of conflicting goals, incompleteness in decisions, errors in responding to familiar conditions etc., stimulates the onset of the learning process. Therefore in the two extreme situations of purely deliberate and purely tacit systems of cooperation, learning is the fundamental force which stabilizes cooperation.

8. Limitations

Let me suggest the most important limitations in the above discussion.

- i. Routinized behaviors can emerge in a huge set of different social contexts either at the individual or at the collective level. Even if we are mainly concerned with economic organizations, we cannot avoid recognizing that highly routinized behaviors in games, warfare, soccer, etc. are a widespread

phenomenon. What seems to me common to all these situations is the existence of a challenging or conflicting environment in which coordination is required among a large number of individuals . I suggest that routinization can happen in many different ways, some of which are the following :

- a. an active search for a set of rules to coordinate a collective action
- b. a passive adaptation to orders and rules issued by an external authority
- c. a process of interiorization with a low level of comprehension , as happens in imitation and in the internalization of norms.

For this reason it is reasonable to consider —as in the title of our meeting —"other patterns of action" beyond behaviors which are cognitively routinized

- ii. One of the most interesting features of routinized behaviors is their inertia, local stability and suboptimality. The last property, emphasized by Nelson and Winter, allow us to go beyond the Chicago school approach (the optimality of an organization's behavior arises from selection), and to consider the possibility that economic institutions can remain in highly sub-optimal configurations for a long time without being able to establish new innovative internal rules. Innovation can therefore be better framed in its real difficulties and obstacles.

Local stability is a largely unexplored feature. To say that observed routinized behaviors are stable means that the rules governing behaviors do not change: this is a symptom that a localized learning process does not lead to the discovery of a new, more efficient set of rules. I have emphasized this cognitive explanation of the persistence of a set of rules within an organization. But, recalling again March and Simon's book, (1958, page 215) let me suggest that an equally relevant source of persistence (inertia) is that rules can be the outcome of a locally solved conflict of interests. A set of rules governing an organization can be stable because the team involved does not discover new, more efficient rules, as I have emphasized above. But also because of the reallocative effects involved in the change of rules: any change would recreate new internal conflicts which could not be easily solved. Consequently, even if our discussion has been based on the cognitive aspect of rule-based behaviors, I suggest that an important challenge for future research is to gain a better understanding of the links between the cognitive and the conflictual forces which give rise to the persistence of rules. From this viewpoint

rules can be to some extent considered social norms, and the internalization of norms should be considered as one of most important sources of persistence.

9. Final remarks

Turning to the problem of "routine" definition, I consider fully satisfactory the March and Simon definition of "routinized behaviors" based upon mental activity-inactivity. The exclusion of a psychological micro-foundation would have the effect of excluding any chance to confirm experimentally our assumptions on bounded rationality, tacitness and awareness: I would stress that Polanyi's analysis of tacit knowledge (1958) and Hayek's theory of knowledge incompleteness (1952) are deeply rooted in a theory of human mind and of learning processes. Moreover, without experiments at the level of mental models it would be impossible to decide which of the two explanation of behaviors - the rule based approach and the opposite explanation based on global optimality and rational expectations - can be confirmed.

The bounded rationality hypothesis and March and Simon's definition of routinized behaviors imply in my view the epistemological position that we must inevitably refer to experiments on human psychology in order to confirm our models and explain behaviors.

Everybody can agree that there are many possible models of mental activity. But whatever model we assume, routinized behaviors should - I believe - be based on the absence or the reduction of active-thinking, i.e. on the emergence, in the mental activity involved, of automatization and tacitness: this is a question which involves not only economic but also political behavior. (see Hirschman's (1984) distinction between "wanton" and "non-wanton" choices.)

To conclude, we can decide to use the magic word 'routine' as a synonym for "pattern of recurring routinized behaviors" as some of us suggested - and in this case the use of the word 'routine' as synonym for "procedure" or "program" should be excluded. Or, vice versa, we can assume 'routine' to be synonymous with procedure, i.e. a set of rules which generates the repetitive behaviors: this is a purely nominal question.

My position rests on the idea that what matters is not how to use the magic word 'routine', but how to give a better framing and understanding of the process of routinization. As I suggested, even though we limit ourself to the analysis of routinization as a cognitive process, many properties have to be explored

to better understand the rise and modification of the rules governing economic organizations.

Appendix B : Addresses of the co-authors

Roger Burkhart Deere & Company John Deere Road Moline, IL 61265 Tel: 309 765-4365 Fax: 309 765-5168 email: roger@ci.deere.com

Michael D. Cohen 407 Lorch Hall University of Michigan Ann Arbor, MI 48109-1220 313 763-0098 Fx: 763-9181 mdc@umich.edu <http://www.crew.umich.edu>

Giovanni Dosi Department of Economics, University of Rome "La Sapienza", Via A. Cesalpino 12/14, 00161 Roma, Italy, tel: 39-6-44284232 fax: 4404572 e-mail: g.dosi@caspur.it or taher@iiasa.ac.at

Massimo Egidi University of Trento Laboratory of Experimental Economics via Inama 1, 38100 Trento +39 - 461-882223 / 882203 Fax +39 - 461-882222 email megidi@risc1.gelso.unitn.it

Luigi Marengo Department of Economics University of Trento Via Inama, 1 38100 TRENTO Italy Tel: +39-461-882201/03 Fax: +39-461-882222 email: lmarengo@risc1.gelso.unitn.it

Massimo Warglien, Universita' degli Studi di Venezia, Ca' Bembo, S. Trovaso, Venezia 30123 Italy eMail: warglien@unive.it Tel: +39 41 528 9311 Fax: +39 41 520 8657

Sidney Winter, The Wharton School, University of Pennsylvania Philadelphia, PA 19104 Tel: 215 898 4140 Fax: 202 363 2759