Computational & Mathematical Organization Theory, 9, 91–108, 2003

Computational Laboratories for Organization Science: Questions, Validity and Docking*

RICHARD M. BURTON
The Fuqua School of Business, Duke University, 1 Towerview, Durham, NC 27708-6245, USA
email: rmb2@mail.duke.edu

Abstract

A computational laboratory is a "place" where we can: ask a question about an organization and its processes, build a computational experiment, design and conduct an experiment, and answer or comment on the question. The questions can be: *what is*, *what might be*, and *what should be*.

Validation is a fundamental concern in science; the validity of a laboratory and model depends upon the question being addressed. A laboratory for a descriptive *what is* question may not be valid for a *what should be* design question.

Docking—the alignment of two models—goes beyond validity. Docking juxtaposes two models to investigate whether they proceed in like manner or yield similar results. I argue that docking provides a guide in the use of different laboratories to address organization questions; and, further computational and non computational models can be docked to deepen and broaden our understanding of organization science.

Keywords: computational laboratories, organizational models, docking, validation, computational experiments, simulation, organization science

Introduction

When we have the real world to observe and lots of data, do we need computational laboratories in our research efforts? We have many different kinds of laboratories for organization science—including computational laboratories. Even so, data from the real world are limited in what we can see and interpret; we cannot address all questions. One interpretation is that the real world is a "big" computational laboratory where we essentially have a single run observation—a stream of data over time, which we cannot re-run. With field studies, we take a sample of smaller pieces of data, frequently assume independence, create an experiment and analyze our observations as if it were a real experiment with manipulation. Many of the statistical methods used in organization science research are based upon this experimental model. In short, we create many observations from one long stream and then proceed as if we had a real experiment. It is a powerful approach and has served organization science and social science well. Perhaps its biggest limitation is lack of experimental

*Earlier versions of this paper were presented as the keynote address at CASOS conference at Carnegie Mellon University, June 2002 and as an invited lecture at the University of Southern Denmark, August 2002. It builds upon the discussions in earlier papers: Burton and Obel (1995), Burton (1998), and Burton (2001).

manipulation; our experiment is limited by what has happened and what we can see. Except by experimental design, extrapolation, extension and inference, it does not permit us to move beyond what happened in the past. Computational laboratories permit us greater experimental variety to complement other approaches used in organization science. With computational experiments, we can manipulate experimental parameters of interest and observe outcomes—creating direct causal links.

There is a long history of contributions using computational laboratories and models to organization science: Cyert and March's (1963) behavioral theory of the firm, a new view of the firm which has been investigated further in many field studies; Cohen et al. (1972) garbage can model of organizational decision-making which complemented their observations of university administration; March's (1991) exploration—exploitation balance model which provided a new view of corporate strategy; Carley's (1992) learning model which provided a process view of learning; Cohen and Bacdayan's (1994) learning model as routines; Levinthal's (1997) rugged landscape complexity model which gave us a different notion of environment and some implications for modeling it; Harrison and Carroll's (1991) cultural transfer model which yielded insights on development of organizational culture; among many others. All of these computational models were complementary to other research approaches. Computational laboratories and models with well designed experiments can be used to explore new ideas, provide process explanations for observed relations, confirm and validate existing explanations and hypotheses, and generally complement other approaches.

In this paper, I argue that computational laboratories provide a place to address the broad array of organizational questions of *what is*, *what might be*, and *what should be*. They complement other laboratories as human subject laboratories and experimental studies which use the real world as a laboratory for data. Organization science questions can and should be studied in multiple places using different approaches. Each laboratory has its advantages and limitations (McGrath et al., 1982; Plott, 1982). A computational laboratory can be quite efficient in the application of resources for process mechanisms to develop sufficient explanations for real observations. Further, we can investigate *what might be* questions in computational laboratories—questions that would require real world manipulations that are not possible in real world contexts. For *what should be* questions, we investigate issues of organizational efficiency, effectiveness and viability before implementing the solution in actual practice. We can plan and design better solutions as well as prevent mistakes that we observe in practice.

Validation is a fundamental concern in science, and organizational laboratories are no different. The validity of a laboratory and model depends upon the question being addressed. A laboratory for a sufficiency explanation for a *what is* question may not be valid for a design recommendation for a *what should be question*. The *what is* question may require only a very simple model of explanation, but would not indicate how to intervene in the organization in a normative manner.

Docking—the alignment of two models—certainly enhances validity, but goes beyond; docking juxtaposes two models to investigate whether they proceed in like manner or yield similar results. I argue that docking can guide us in the use of different laboratories to address organization questions; and, further computational and non computational models can be docked to deepen and broaden our understanding of organizations.

Next, I define a computational laboratory followed by a categorization of organization questions as: *what is*, *what might be* and *what should be*. Computational validity is examined in terms of the research question for which the laboratory is used. Finally, I explore the concept of docking and how it can enhance the study of organizations.

Computational Laboratories

In organization science, we have many questions about: the behavior of individuals, the nature of organizational processes, the implications of alternative organizational structures, and the influences of context of the organization, among others. We investigate the nature of the organization in many venues: observation and description of the world around us, analytical modeling, controlled experiments involving humans as agents, and computational laboratory experiments.

For an organization, a computational laboratory is a "place" where one can:

– Ask a *question* about: the behavior of individual agents or teams, the information processing of the agents and communications in an organization, the outcomes of various organizational processes, the implications of alternative structures, the influence of the organizational context and environment; where each question can be stated in specific terms.
– Build a *computational model* that relates to the question.
– *Design and conduct an experiment* that yields outcomes that can be used to address the question.
– And, then answer, or at least comment on the question.

I begin with the question of how to represent the phenomenon and describe behavior, propose a process and test a hypothesis, explore beyond what we know, test and offer advice on new ways to meet goals, and generate new theory by offering and investigating new plausible explanations. The different questions may require different laboratories and models; no one model can answer all questions; no one model can answer any given question definitively. For any given laboratory, there must be a balance among: the question, the model, the experiment, and the analysis to address the question. Earlier, Burton and Obel (1995) argued that balanced simple models and parsimonious explanations are preferred to complicated models, provided simple models address the question. Computational laboratories need not be—indeed should not be—exact replica of reality; the match of the question with the model indicates the nature of the model. The central issue is to help us learn about the world and to act upon it in a good manner.

There are many computational laboratories that are representations of the organization and how it works. The laboratories can be categorized as:

- Procedural models of the time order of events, e.g., Cohen et al.'s (1972) garbage can model and Van Alstyne's (2003) Indigo,
- Agent based models of individuals in teams and organizations, e.g., Levitt's (Jin and Levitt, 1996) VDT/SimVision, Prietula's (2002) TrustMe, Burton and Obel's (1980) organizational decomposition models and Lenox's (2002) rent producing resources model,

- Equation based models of organizational processes without closed form solutions, e.g., Harrison and Carroll's (1991, 2002) culture transfer models,
- Rule based models, e.g., if-then rules or heuristics which state processes or characteristics of the organization, e.g., Burton and Obel's (1998) OrgCon diagnosis and design model, and Cohen and Bacdayan's (1994) routines model,
- Intelligent agents who follow a program or a procedure to do something for someone, e.g., Prietula and Carley's (1998) webbots.

Most of these laboratories were devised to address specific questions about organizations and the individuals in them. Cohen et al.'s garbage can model depicts decision making processes in universities; Levitt's VDT/SimVision model aims to improve project management; Harrison and Carroll explore the meaning of organizational culture and its management; Burton and Obel diagnose fit/misfit relations for design; and, Prietula and Carley's Webbots are computer generated "critters" which help people with tasks.

Earlier, Cohen and Cyert (1965, p. 308) categorized computational models as: descriptive computational studies, quasi-realistic studies, normative computational studies for designing organizations, and man-machine computational studies for training. Each model type addresses a different question and the validity of the model depends upon the question. Carley (1995) suggests four categories: organizational design, organizational learning, organizations and information technology, and organizational evolution and change. Again, each model type addresses particular questions.

In their new book, Lin and Carley (2003, pp. 30–31) discuss the multiple advantages of computational methods:

(1) Using simulation, we can more fully explore ranges of stress, organizational design, and task environment and their effect on performance (Masuch and LaPortin, 1989). (2) We can conduct balanced simulation experiments, and control certain factors to examine the effect of other factors, while imposing no damage to the existing environment. (3) We can consider both successful and failed firms. Thus results will not be biased by looking only at successes. (4) Simulated organizational have been shown to resemble the real world organizations in an idealized way (Carley and Lin, 1997). The performance characteristics of simulated organizations are under certain conditions comparable to the performance characteristics observed in the real world (Lin and Carley, 2001). (5) Researchers have also shown that organizational performance is affected by factors such as organizational design (Lawrence and Lorsch, 1967; Houskisson and Galbraith, 1985), task environment (Drazin and Van de Ven, 1985), and stress (Anderson, 1977). Hence, simultaneous examination of these factors on organizations can help address the issue of what really constitutes organizational performance. By using simulation, we can get insight into these important factors with less cost than conducting human experiments or field studies. Once the dominant factors are examined, human experiments or field studies can be done to test the theoretical results.[1]

They see computational modeling as a precursor for other approaches and as a complement to experiments and field studies.

Human subject laboratories are conceptually close to agent based computational laboratories. Human subject laboratories permit us to vary conditions within limits to test the effect of these manipulations on the behavior of individuals in the laboratory. We have greater understanding of the conditions and what has been changed to permit inferences about cause and effect relations. In essence, the computer agent, as in agent based models, is replaced by an individual who plays the role of the human agent. Burton and Obel, (1980) confirmed the M-form hypothesis using a multi-agent model without invoking opportunism. Later, they (1988) tested the M-form hypothesis where potentially opportunistic individuals replaced some of the agents. They had the advantage of the experimental control as well as the actions of complex individuals. We can only infer the underlying processes which generated the outcome—albeit, it is usually the question of interest. To obtain clearer inferences, we replicate the experiment to affirm the results under related conditions. Nonetheless, it is a continuing concern whether individuals in the laboratory behave as they do in their natural environment, or whether agents, human subjects or computational, capture the relevant procedural complexity to address the question

Any particular computational laboratory is limited in the questions and issues that it can address well. As with any laboratory, no one laboratory can address all questions well; yet, many computational laboratories can be used to address more than one question (McGrath et al., 1982). There are tradeoffs—advantages and disadvantages in the matching of the question with the laboratory and model. And, no one question can be completely addressed within one laboratory. Later, I will argue that docking, or model alignment (Axtell et al., 1995) permits us to use the laboratories in concert to pursue the science both in greater depth and breadth.

Computational laboratories have the advantage of more complete specification of conditions and the processes which yield the outcome. (We will discuss this more detail later in a comparison of the forward and backward problem.) It is a misnomer to call these experiments virtual experiments; they are quite real. To the contrary, created experiments of real world observation are artificial—created by the observation method itself. Nonetheless, computational laboratories have limitations of external validity and correspondence or closeness with the real world. How close is close enough? In brief, it depends upon the question, the unit of analysis and the nature of the answer we want (Burton and Obel, 1995). It may take a very different laboratory to confirm a hypothesis than a laboratory to offer advice on a specific managerial question. Questions drive organizational science; some questions can be addressed well in computational laboratories—others, not so well. In the next section, we examine three kinds of questions and their implications for computational laboratories: *what is*, *what might be* and *what should be*.

Types of Questions We Study

Organizational questions can be categorized as positive science or normative science—*what is* or *what should be*. We add explicitly the question of *what might be*. Here, we want to examine the implications of these three different, but related questions for computational laboratories. I begin with "what is," then examine "what might be," and finally "what should be." Scientists are usually more interested in *what is* questions; practitioners are more interested in *what should be*.

What is

The question of *what is* is basic in positive science. We use laboratories to describe and understand the world around us. Whether the laboratory is computational, human subject, or field, we create an image of the world where we attempt to mimic, replicate, reflect and in general, attempt to see the world as it is—ever through the glass darkly. Each answer generates a new question—a need for a better description. But we want to go beyond description to explanation and a deeper understanding of why the world works the way it does. If we propose a theoretical model, we want to test it. If we observe a particular behavior, then we want to explain why it obtains. The traditional method is to infer the reason why—we make guesses, hopefully good ones, about why.

Our usual positive science approach is to examine what we know and what we do not know—the question. We then create hypotheses as potential answers to our question. Using the real world as the computational laboratory and model for the experiment, we observe and collect appropriate data, and analyze these data to answer the question. If we confirm the hypotheses, we have a sufficient explanation. That is, we have a single explanation that fits the question and explains the results. We work hard to eliminate alternative explanations, or other plausible sufficient explanations. Rarely do we ever have a necessary explanation to the question. But we have greater confidence in our understanding if we have one sufficient explanation, and can eliminate a number of other explanations, which might be true, but are not. Further, no question can be definitively answered in one laboratory or with one model. We have greater confidence in the answer if we can triangulate and confirm the answer in multiple laboratories—for example, the real world laboratory, human subject laboratories and computational laboratories.

Let's examine some of the challenges to make computational laboratories good laboratories for *what is* questions. Computational laboratories for *what is* questions need to be close to the real world and have external validity, but how close depends upon the question. The model need not be an exact replica of the real world; it should be very close on the dimensions directly related to the question, less close on others which create the context and can be completely silent on issues beyond the boundary of the question. Cohen and Cyert (1965) posit:

> *...even though the assumptions of a model may not literally be exact and complete representation of reality, if they are realistic enough for the purposes of our analysis, we may be able to draw conclusions which can be shown to apply to the world.*

The research question tells us what is "realistic enough for the purpose." At a first level, we examine how well the model duplicates the phenomenon of interest, i.e., can it mimic the behavior in the real world or can we show an equivalency of the model with the real world phenomenon; if it can, then we have validated the model. But no model mimics, replicates, mirrors or describes the real world totally; that is, it mimics a small part imperfectly. It is important to understand what is included and what is ignored, or in other words, the boundary of the model.

Then, what can we learn about *what is* from computational laboratories? Computational laboratories can add a good deal of to our understanding of organizations. Lomi and

Larsen (2001) contrast two kinds of research methods: the forward problem, "given a set of assumptions about individual decision rules and problem-solving procedures, can we determine (predict) the aggregate properties of a system generated by the repeated interaction among individual units adopting these rules and procedures?" (p. 4), and backward (inverse) problem, "given the observable regularities in the behavior of a composite system (e.g., an organizational field or a market) can we specify a set of rules or procedures that—if adopted by all the elementary units—induce and sustain these regularities?" (p. 5). In an earlier discussion, Gutowitz (1990) focusing on cellular automaton stated:

> The forward problem is: Given a cellular automaton rule, determine (predict) its properties. The inverse problem is: Given a description of some properties, find a rule, or set of rules, which have these properties. These problems are obviously strongly interrelated (p. vii). On the experimental side, the forward problem amounts to the design of good simulations which adequately reveal the behaviors under investigation (p. ix). (for the inverse problem), to develop a set of techniques which allow one to find a rule, or set of rules, which quantitatively reproduce some set of observations of the physical system (p. x).

Briefly, for forward problems in organizational studies, we specify the rule, the process or the mechanism and then observe the outcomes; for backward problems, we observe the behavior or the outcome and infer the process or mechanism. "These two are obviously related" (Gurtowitz, 1990). For the last century, social science and organization science in particular has been dominated by the backward (inverse) approach; we observe the real world and infer the mechanism which generated it, which is usually stated as propositions or hypotheses concerning behavior. We have elaborate techniques on observation and analysis for confirmation.

Many, if not most computational models investigate the forward problem. The earlier list of procedural, agent, equation, rule based and intelligent agents: all are forward models. An organizational process model specifies how the organization will work; it is run and outcomes are observed. The forward model process is then a sufficient explanation for the results. From this perspective, computational laboratories are complements to real world observations; the computational model yields process data, just as the real world yields process data. The inverse problem is to infer the nature of the process. Beginning with a backward approach to a real world observation, we can then use the forward approach in a computational laboratory to replicate a process or build a plausible process model. But more importantly, we can also confirm that alternative processes or explanations do not yield the outcomes, thus eliminating them as possible explanations. And, further we can propose new explanations or new theories about the outcomes and examine whether they are feasible explanations or not; theory generation, if you will. Here we are going beyond *what is* to *what might be*.

What Might be

What might be questions go beyond *what is* to examine what we have not observed. In organization science, *what is* questions have been our focus. And, careful researchers are restrained not to go beyond the limits of their data and speculate on "what might be." There is good reason for this caution. Usually, we venture into "what might be" by simple extrapolation, which has grand underlying assumptions. Thus, we limit our "what might be" to be close in variations on "what is"—not unlike the neighborhood of partial derivatives in a Hilbert vector space.

Computational laboratories with a specification of the underlying procedural rules, or processes permit us to go beyond the nearby linear extrapolation to investigate "what might be" in the extremes. We can manipulate conditions and parameters widely within the model which may be nonlinear and observe the outcomes. The "what might be" space is larger (but contains the what is), and we can investigate a broad array of alternative explanations and alternative theories.

Levitt (Jin and Levitt, 1996) and his associates in their investigation of the implications of time crunching large scale projects using the virtual design team (VDT) laboratory found two important outcomes: first, the managers themselves were quite good at local extrapolation, i.e., reducing project time by 3 percent, but their intuition was quite limited about the ensuing difficulties created by crunching the project time to one half; and second, the VDT laboratory and the project model predicted correctly the actual problems encountered in reducing the project time by one half. In brief, the managers were quite good at the "nearby," but were not very good at predicting the implications of major changes. They confirmed that we are rightly cautious in extrapolating without an understanding of the underlying processes.

The real world has an obvious success bias, and rightly so; we would like to understand what is successful. But a more complete understanding demands that we understand what is likely to be unsuccessful and why. Computational laboratories can be crash laboratories for organizations—similar to flight simulators for airplanes. We are not only free to crash at very low cost, we can learn a good deal by creating failure. We can utilize computational laboratories to test new ideas, i.e., hypothetical organizations that we have yet to observe, but can be specified in sufficient detail to construct a model of how it might work. For example, we can construct a model of an organization where all members can talk with each other on any topic at any time, and then observe how it will behave. Further, we can test the implications of adding new members to the organization. It is lower cost than "just experimenting" and see what happens; and, we are likely to observe problems in the laboratory before we experience them in a trial run of the organization. We want to crash in the laboratory—not in the real world.

What Should be

Design, or *what should be* goes beyond *what is* or *what might be* to incorporate what is "good." *What should be* is the combination of *what might be* and what is "good"—or, more correctly what is "better" as we frequently make relative comparisons. In positive science, we are reluctant to say what is good, or *what should be*; in normative science, it is the question of interest.

Simon (1981, p. 133) stated the *what should be* or the design issue well:

> *Design, on the other hand, is concerned with how thing ought to be, with devising artifacts to attain goals. We might question whether the forms of reasoning that are appropriate to*

natural science are suitable also for design. One might well suppose that introduction of the verb 'should' may require additional rules of inference, or modification of the rules already imbedded in declarative logic.

The question of what is good can be examined on many levels. Usually, we operate at the utilitarian level or preferences on efficiency within moral constraints. For agriculture, more corn is good; for engineering, a stronger and lighter standing bridge is good; in business, more profits are good; in public policy, more open democratic processes are preferred—even if less efficient. But the problem is more involved. Most outcomes are multi-dimensional involving incommensurable tradeoffs. Frequently, there are many stakeholders in an issue, and it is well known that individual preferences do not simply add up. One approach is to test the implications of different preferences for the design to obtain a better understanding of the tradeoffs—gains and losses among various stakeholders. Computational laboratories are particularly adept at generating the implications of the various alternative preference functions and indicating what is gained and loss, and by whom. But that is only at the first level of relative good, not absolute—a more destructive atom bomb is not necessarily a good thing.

What should be questions are more difficult and may require different laboratories. First, there is the technical issue of understanding whether the new design of *what might be* is feasible or possible. Second, the question of *what should be* may require a different laboratory where *what should be* can be explored, i.e., the *what should be* question suggests that the artifact or design model should be valid—a different validity than for a *what is* question. Third, experts on *what is* are not necessarily experts on *what might be*. It can fall outside the boundaries of what they understand. Fourth, *what should be* requires a judgment on what is good. Physicists who know how to make an atom bomb are not necessarily experts on whether it should be built. Today, we are confronting such issues on genetics and ethics in corporate behavior.

Computational Laboratory Validity and the Research Question

As argued above, we begin with the question: *what is*, *what might be* or *what should be*. Earlier, Burton and Obel (1995) argued that the validity of a computational laboratory depends upon the purpose—a matching of the question, the computational model and the experimental design. Their research question categories become: *what is*, or descriptive of behavior, hypothesis testing and alternative explanations of behavior; *what might be*, or exploration and theory generation; and *what should be*, or normative model, advice and training. The correspondence of the laboratory or model with the phenomenon, or the level of realism depends upon the question.

What is questions require a degree of realism. But realism does not mean complete correspondence; simple models can address focused questions with insight. Axelrod's (Axtell et al., 1995) cultural model applies rather simple mechanisms of neighbor "touching" to give us insight into community integration—or, lack thereof. For his purpose, it is important to strip away the complexity of the real world to obtain insight and a parsimonious explanation. In real world experiments, we have observations where community integration occurs and where it does not, but sorting out cause and effect is very difficult; it is much clearer in the computational laboratory experiment where confounding realism is stripped away. In this computational laboratory, we have one sufficient and parsimonious explanation for the observations–a test of a hypothesis about the mechanism of community integration; there may be others as well. And, we understand more by examining the question in other laboratories–applying triangulation (McGrath et al., 1982) to the question.

Other *what is* purposes suggest a closer correspondence with the real world phenomenon. Levitt and his associates (Jin and Levitt, 1996) validated VDT against *what is* criteria by comparing the computational model with actual projects retrospectively, i.e., comparing the model results and processes with actual projects. (In personal correspondence, Ray Levitt suggested that Walt Disney called this "hindcasting.") For projects that are incomplete, they validate the model both retrospectively and also prospectively by predicting what will occur and comparing the predictions with what occurs. The combination gives a good deal of confidence that the VDT computational model incorporates a high degree of realism—both outcomes and processes. More generally, process models and in particular agent based models contain a trace of what occurred and it can be reviewed for validity and understanding.

For *what might be* questions, validity requires a confirmation that the laboratory and model make sense in this larger world of possibility: exploration, idea generation, alternative explanations, and testing. I argued above that a simple linear extrapolation is limited; but, an understanding of more fundamental mechanisms is more likely to tell us *what might be* possible in this larger space. Computational models which incorporate the organizational processes, or explain relations and connections among the variables and parameters are more likely to be valid in a larger space than models which explain a point observation of the past. Again, Levitt's VDT model for projects has a high level of realism in a rocket development study (personal correspondence); it yielded better insight into *what might be* than the manager's intuition based upon experience. That is, it is both valid for *what is* and for *what might be*. Burton and Obel's (1998) OrgCon model of relationships among organizational characteristics and properties provides insight into what will work well, and what is likely to lead to a performance diminishing misfit.

The validity for *what should be* laboratory adds the dimension that preference choices can be revealed among design alternatives and realized by the design solution and its implementation. We want to go beyond the feasibility of *what might be* and illustrate the desirability of *what should be*. For advice or design, we need to demonstrate the efficiency or effectiveness or higher profits for the proposed organizational solution. In Levitt's VDT model, the project outcomes are time, cost and project quality. First, the project activities and implementation process must be mapped onto these outcomes. Second, these outcomes must be relevant as goals for the decision-maker, i.e., time and quality are meaningful goals to choose among. Thomsen et al. (1999) developed a trajectory approach to a dynamic validation of VDT and more generally *what should be* models to assure confidence in the normative recommendations.

It is important to note that a valid model for one purpose may not be a valid model for another. The validity of the model needs to be established anew for each purpose or question. Levitt's VDT model is valid (Thomsen et al., 1999) for many purposes as suggested above. However, Cyert and March's (1963, pp. 128–148) department store model is valid for *what is* but it is not clear whether it is valid for *what might be* or *what should be*.

The validity of a laboratory depends upon the question. No one laboratory can address all questions; further, no one laboratory can completely answer one question. Any one question must be addressed in numerous ways to obtain insight and understanding. Thus, we need to investigate a question in multiple laboratories which can be related, or docked. As developed below, the docking of two or more laboratories or models enhances the validity of each; we know more about the capabilities and limitations of each laboratory to address a question; but more importantly, we have deeper understanding of the science.

Docking

Axtell et al. (1995), define docking—a metaphor borrowed from the docking of two dissimilar spaceships—as the "alignment of computational models." They docked Axelrod's cultural model (ACM) and Epstein and Axtell's (1996) Sugarscape by showing a distributional equivalency, i.e., the two models can produce the same results—a tough test of alignment. They defined two levels of equivalency: distributional and relational. Distribution equivalency requires the same (numerical) results; relational equivalency requires equivalent internal relation results. ACM and Sugarscape are similar in that both are models of a community and how it works—albeit at different levels of abstraction. The two models are similar in that both have agents who use relatively simple rules of behavior and a geographical space for action. Both are forward models addressing *what is* questions to explain the community behavior from relatively simple micro behaviors. They are different in that ACM is relatively focused and small; where the Sugarscape model is very complex and large with a rather complete description of a community. The purpose for the ACM is to posit simple behaviors of individuals which explain the community integration result—a parsimonious explanation; the purpose for Sugarscape is the creation of a broad set of relatively simple behaviors for a community that replicate a complex society. The docking experiment investigated whether the two models could yield equivalent distributional results with respect to the integration of the community. The equivalency between the two models provides a validation for each model—what questions can be answered well, and, those questions which cannot be answered well. We have a better understanding of what each model captures and what each model ignores. Docking gives us a much richer understanding of each model and the touch points between the two. It can be thought of as validating the models beyond the validation of each model separately. It yields greater insight into the phenomenon, and most importantly, it enhances the science.

In this broader sense, docking is not new. Cyert and March (1963, pp. 128–148), in their store buyer behavior study used the distributional equivalency criterion in comparing their model and the real world data. They docked their computational model with the real world—two *what is* models. The equivalency yielded validity to the computational model where the purpose was to show that the computational model could be used to suggest real world decision-making processes and outcomes. In a similar fashion, the garbage can model

(Cohen et al., 1972) was docked with their observations of university administration processes. Nelson and Winter (1982) replicated history using routines as the organizational processes.

Science progresses by building upon the work of others. The research question comes from what we know or think we know, and what we do not know, but would like to know. Our confidence in what we know is based upon the scientific process of investigation—and, the validity of the laboratory. Docking reveals the scientific process and makes it more explicitly integrated so that we can build upon the research of others.

Axtell et al. took a restrained view of docking as the equivalency of models. We want to generalize the concept to incorporate others forms of docking. In thinking about docking possibilities, there are several dimensions:

Laboratories, databases, field studies, human subject, and computational;
what is, what might be, and what should be; and,
forward and backward models.

The number of combinations is too numerous to investigate in total; we need an organizing criterion for what to dock. Again, the driver for any docking study is the question we want to answer: validation of the models, explanations, confirmation of results, mimic the processes as well as the results, testing alternative sufficiency explanations, disconfirming a plausible sufficient explanation for a phenomenon, developing insights about the organizational processes, and any question where the comparison of two or more models deepens our understanding of the science.

Informal Docking

Informal docking of two or more *what is* models is well known and widely practiced; in fact, it is the norm in organization science. Journal articles are filled with informal docking. The usual approach is to position the question within the research literature; good science builds upon what we know and the previous work of others. Generally, the question itself addresses what we do not know, but would find interesting to know more about. The literature review sets the stage for the research; the development of hypotheses are supported by reference to the theory, or what we think we know; the experimental design gives us confidence in the research procedure and in the results; and, the laboratory, whether it is a data set, field study, or computational model yields the data for analysis and the answer we seek. Usually we do not directly confront different models, different experimental designs, or different data sources—as Axtell et al. did. It is still quite rare to have a replication of an experiment and demonstration that the results are equivalent in either the distributional or relational sense. Our usual approach is to consider the experiments as independent tests, and downplay the connections. Our docking is largely suggestive by juxtaposition–"docking lite", if you will.

Nonetheless, informal docking is extremely important—a beginning, but limited step. We can extend the docking notion further for greater validation and deeper understanding.

Docking en Large

Axtell et al. (1995) docked two *what is* forward computational models using a distributional criterion of equivalency. Docking en large extends, or more fully develops Axtell et al.'s idea to explore the connectedness among the questions, the laboratories and the experiments. The general idea is to investigate whether two models "touch" more deeply than we are able to do by informal docking. Here we want to investigate some possible docking variations.

We want to explore the possibility of docking more generally along the dimensions:

- equivalency: distributional and relational,
- computational models, field studies, large database studies and human subject experiments, and,
- *what is* models with *what might be* and *what should be* models.

There are innumerable experiments and questions of equivalency. In the discussion above, we have suggested some possible docking experiments. Here we want to explore some docking questions and related experiments:

Is There a Forward Procedural Explanation for a Backward Explanation?

As discussed earlier, for forward problems in organizational studies, we specify the rule, the process or the mechanism and then observe the outcomes; for backward problems, we observe the behavior or the outcome and infer the process or mechanism. "These two are obviously related" (Gutowitz, 1990). Sastry (1997) developed a forward procedural explanation for the observations in Romanelli and Tushman's (1994) punctuated equilibrium model. Sastry's forward explanation was much less elaborate than Romanelli and Tushman's, suggesting that a more parsimonious explanation is sufficient for the results. Recently, Ocasio (1999) presented a sophisticated backward explanation of CEO succession. There are hypothesized procedural mechanisms which are plausible, but not tested. A procedural model of the CEO selection procedure could be created to dock the backward inference with the forward procedural model, and further to the test the robustness of the explanation as discussed below. As I have argued, our confidence in the hypothesized explanation would be enhanced.

Are the Models Alternative Sufficient Explanations of the Phenomenon?
or *Is There a Confirmation of a Model Using an Alternative Explanation?*

Many of our studies yield one, or a very few explanations of the results. We have a theoretically based hypothesis which we confirm. Any pair-wise docking and confirmation gives us greater confidence in the explanation hypothesized. In the Axtell et al. (1995) study, they generated equivalent results, and thus, have alternative mechanisms, but related explanations to explain the results. It is an interesting and open question whether there are additional explanations—probably so. Our understanding is richer and more satisfying if we have a few alternative explanations—or better yet, when we have a few sufficient explanations and can eliminate a large number of plausible alternative explanations. We have greater confidence in an explanation if we can triangulate (McGrath et al., 1982) among computational and human subject experiments and real world observations.

Can We Eliminate a Plausible Explanation That is not Feasible?

We have a stronger theory and greater confidence in the explanation if we can eliminate alternative possible sufficient explanations. Our understanding is enhanced when we can eliminate a plausible explanation which is in fact not true. E.g., if we have a plausible hypothesis for which we have a backward test, and then we construct the related procedural model and find it does not yield an equivalency, then we can reject that explanation.

What are the Boundaries or Limits of the What is Explanation for What might be?

When we have a *what is* explanation, we would like to explore the limits of the explanation to understand better *what might be*. There may be a tendency to overstate the universality of our results. Many of our tested hypotheses are more general than the evidence or data we use to support the results. It is critical to understand the limits of our conclusions. Within a single model, the robustness of the results is important. One approach is to systematically vary parameters in the model which reflect important assumptions. If the results obtain for wide variations in these parameters, then we have greater confidence in the universally of the results. If the results are fragile, we have a narrow sufficient explanation. More generally, we may want to test the boundary limitations by purposefully “crashing” the model to understand when it does not obtain as hypothesized.

If we have two or more models where we have tested the boundaries, it is not likely the boundaries will be the same. The conservative interpretation is that the intersection of the explanation interiors is the region of explanation. The intersection of the exteriors is an excluded region. The region of explanation by one model and not by the other creates a contingent region of explanation.

Is One Model a Special Version of the Other?

The special version is usually created by holding some variables constant and then examining the results. Axtell et al. (1995, p. 124) note that Newton’s concept of gravity is a special case of Einstein’s theory; Newton’s gravity works well on the earth, but not beyond. In social science, we are frequently examining a simple version of a theory to test whether it remains a sufficient explanation for a phenomenon; and if so, we usually prefer the simpler, parsimonious explanation. ACM is not a special version of Sugarscape in a strict sense in that there are some parameters in Sugarscape are turned off to create ACM, but the ACM is clearer simpler model than Sugarscape; yet, both explain the diversity in the community within the distributional equivalency criterion. Further, we are interested when the simple model is not a good explanation, as Einstein was. We have suggested above that all laboratories—computational, human subject or real world observations–are special versions of a deeper reality. Our real world laboratory and the computational one are different special versions, and further, the computational model can be a special version of the real world model. We want to know the special conditions when they are equivalent and when they are not—and whether equivalent in a distributional or relational sense. As discussed above, the two models may have differing regions where each is valid.

Are the Models at Different Levels of Analysis?

Two models can be at differing levels of analysis for the same phenomenon, e.g., the individual, the organization and the society. Organizations can be analyzed in terms of properties and characteristics: structure, centralization, formalization, etc, (Burton and Obel, 1998), or as agents or individuals who communicate and make decisions (Jin and Levitt, 1996; Carley and Lin, 1997). The two models are then at differing levels of analysis, containing some concepts in common, e.g., centralization. For docking test, the concept of equivalency would need to be further developed. Similarly, ACM is an agent based culture model; Harrison and Carroll (1991) also consider organizational culture, but at a more aggregate level of hiring, firing and socialization processes, but not at the individual level of behavior. Here, these models are quite different; a new equivalency criterion is needed before we can consider what docking might mean—culture has two definitions here.

Using institutional arguments, Dimaggio and Powell (1991) propose that organizations look alike or are isomorphic due to coercive, mimetic and normative forces in their organizational fields. The underlying micro mechanisms are intuitive and make good sense, but are not actually tested as to their feasibility. Mimetic isomophorism suggests that organizations copy successful organizations without being explicit about the micro mechanisms and possible limits. Rivkin (2000), in a computational study on a rugged landscape, found that unless an organization copies the total strategy of another organization, it is unlikely to obtain the same high level of performance. E.g., a near perfect copy of Southwest Airline's successful strategy is not likely to yield the same high profits. Informal docking suggests that there are severe limitations to mimetic strategies, but we do not know much about the conditions and boundaries. Benchmarking is advocated as a universal solution to improve efficiency and profits, and it is practiced in piecemeal fashion. Rivkin's study suggests there are limiting conditions for benchmarking to be successful. More generally, the institutional isomorphoric mechanisms need clarity in the possible micro mechanisms that yield these outcomes, and a better understanding of the boundaries of applicability.

Can the Models be Linked or Related?

There are a number of ways that two models can be linked:

- elements—variables and relations—in common, but with some different variables. As mentioned above, centralization was a variable in two different models. Here you can also link field with computational studies.
- the output of one model can be the input for another e.g., agent based organizational models and experiments can yield design heuristics for rule based expert systems. One possibility is the development of expert system design rules (Burton and Obel, 1998) from agent based models which test organizational properties, say on decentralization and yield results which can be stated as rules. Another possibility is a cascaded parameterization process where the first model determines the value of a variable that then becomes a parameter value in the second model. Here, we could reverse the direction where the design model would suggest an appropriate level of decentralization and the agent based model would take the level of decentralization as a given or input parameter.

These variations on docking are not exhaustive, nor mutually exclusive. And, further equivalency needs to be specific to the question and the computational laboratories. The concept of equivalency is fundamental, i.e., on what basis are we making a comparison.

Summary

I have argued that computational laboratories can be used to further organization science and our understanding of how organizations work. We can address questions of: *what is*, *what might be*, and *what should be*. And further, the docking of laboratories is an excellent way to enhance organization science.

Briefly, I would like to restate the main ideas and arguments:

- Computational laboratories are complementary laboratories to human subject laboratories and field studies as laboratories. Each has its advantages and its limitations. Computational laboratories permit us to go beyond and explore *what is* possible in other venues; *what might be* is a larger world of possibility to explore, test ideas and design new organizations for *what should be*.
- The validation of a laboratory and a model is dependant upon the question under consideration.
- Validity is enhanced through docking and our understanding of the science is deepened.
- Triangulation is well accepted in science and the computational laboratory gives us another place to do research. No question can be answered definitively in any one laboratory; differing laboratories are needed.
- Sufficiency explanations are more completely specified and we have a better understanding of what we are testing where the forward problem specification of the organizational process rules out a number of alternative explanations, or at least, suggesting an alternative mechanism is required.
- Docking, where two or more laboratories or models are connected in non trivial ways, deepens our understanding of organizations and how they work, and further the science of organization.

Computational laboratories have been venues for fundamental contributions to organization science. As computational laboratories become more widely available to organizational scholars, we will see continuing, and perhaps growing use.

Acknowledgments

My thanks to those who offered constructive comments on presentations at CASOS conference at Carnegie Mellon University, The University of Southern Denmark, and to colleagues who reviewed draft versions and offered constructive comments: Tim Carroll, Corinne A. Coen, Quintus Jett, Carl R. Jones, Michael Lenox, Ray Levitt, Min Li, Will Mitchell, Borge Obel, Michael Prietula, Michael Roach, Scott Rockart, Henry Sauermann, Scott Turner, Aparna Venkatraman, among others. I have incorporated many very helpful suggestions and ideas. Obviously, the shortcomings remain mine.

Note

1. The references in this quote are not included in the paper references at the end.

References

Axtell, R., R. Axelrod, J.M. Epstein and M.D. Cohen (1995), "Aligning Simulation Models: A Case Study and Results," *Computational and Mathematical Organization Theory*, 1(2), 123–142.

Burton, R. (1998), "Validating and Docking: An Overview, Summary and Challenge, Simulating Organizations: Computational Models of Institutions and Groups," in Michael J. Prietula, Kathleen M. Carley, and Les Gasser (Eds.), Menlo Park, CA: AAAI Press/The MIT Press, pp. 215–228.

Burton, R. (2001), "Afterword," in Alessandro Lomi and Erik R. Larsen (Eds.), *Dynamics of Organizations: Computational Modeling and Organization Theories.* Menlo Park, CA: AAAI Press/The MIT Press, pp. 439–444.

Burton, R.M. and Børge Obel (1995), "The Validity of Computational Models in Organization Science: From Model Realism to Purpose of the Model," *Computational and Mathematical Organization Theory*, 17(1), 57–71.

Burton, R.M. and Børge Obel (1988), "Opportunism, Incentives, and the M-form Hypothesis—A Laboratory Experiment," *Journal of Economic Behavior and Organization*, 10, 99–119.

Burton, R.M. and Børge Obel (1980), "A Computer Simulation Test of the M-form Hypothesis," *Administrative Science Quarterly*, 25(3), 457–466.

Burton, R.M. and Borge Obel (1998), *Strategic Organizational Diagnosis and Design: Developing Theory for Application*. Kluwer Academic Publishers, Boston.

Carley, K.M. (1995), "Computational and Mathematical Organization Theory: Perspective and Directions," *Computational and Mathematical Organization Theory*, 1(1), 39–56.

Carley, K.M. (1992), "Organizational Learning and Personnel Turnover," *Organization Science*, 3(1), 20–46. Reprinted in *Organizational Learning*, Michael D. Cohen and Lee S. Sproull (Eds.), 1996, Thousand Oaks, CA: Sage.

Carley, K.M. and Z. Lin (1997), "A Theoretical Study of Organizational Performance Under Information Distortion," *Management Science*, 43(7), 976–997.

Cohen, K.J. and R.M. Cyert (1965), "Simulation in Organizational Behavior," in James G. March (Ed.), *Handbook of Organizations*, Chicago, IL.

Cohen, M.D., J.G. March and J.P. Olsen (1972), "A Garbage Can Model of Organizational Choice," *Administrative Science Quarterly*, 17, 1–25.

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

Cyert, R.M. and J.G. March (1963), *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ.

Dimaggio, P.J. and W.W. Powell (1991), "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields," in W.W. Powell and P.J. Dimaggio (Eds.), *The New Institutionalism in Organizational Analysis*, Chicago: University of Chicago Press, pp. 63–82.

Epstein, J.M. and R. Axtell (1996), *Growing Artificial Societies: Social Science from the Bottom Up.*

Gilbert, N. and K.G. Troitzsch (1999), *Simulation for the Social Scientist*. Open University Press, pp. 273.

Gutowitz, H. (1990), "Cellular Automata: Theory and Experiment," *Proceedings of a Workshop Sponsored by The Center for Nonlinear Studies*. Los Alamos, The Hague, North-Holland, pp. 7–14.

Harrison, J.R. and G.R. Carroll (1991), "Keeping the Faith: A Model of Cultural Transmission in Formal Organizations," *Administrative Science Quarterly*, 36, 552–582.

Harrison, J.R. and G.R. Carroll (2002), "The Dynamics of Cultural Influence Networks," *Computational and Mathematical Organization Theory*, 8(1), 5–30.

Jin, Y. and R.E. Levitt (1996), "The Virtual Design Team: A Computational Model of Project Organizations," *Computational and Mathematical Organization Theory*, 2(3), 171–196.

Lenox, M. (2002), "Organizational Design, Information Transfer, and the Acquisition of Rent-Producing Resources," *Computational and Mathematical Organization Theory*, 8(2), 113–131.

Levinthal, D.A. (1997), "Adaptation on Rugged Landscapes," *Management Science*, 43(7), 934–950.

Lin, Z. and K.M. Carley (2003), *Designing Stress Resistant Organizations: Computational Theorizing and Crisis Applications*. Kluwer Academic Publishers, Boston.

Lomi, A. and E.R. Larsen (2001), *Dynamics of Organizations: Computational Modeling and Organization Theories*. AAAI Press/The MIT Press, Menlo Park, CA.

March, J. (1991), *Organization Science*.

McGrath, J.E., J. Martin and R.A. Kulka (1982), *Judgement Calls in Research*. Sage Publications, Beverly Hills and London.

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

Ocasio, W. (1999), "Institutional Actions and Corporate Governance: The Reliance on Rules and CEO Sucession," *Administrative Science Quarterly*, 44, 384–416.

Plott, C.R. (1982), "Industrial Organization Theory and Experimental Economics," *The Journal of Economics Literature*, xx(4), 1485–1527.

Michael, P. and K.M. Carley (1999), "Exploring the Effects of Agent Trust and Benevolence in a Simulated Organizational Task," *Applied Artificial Intelligence*, 13(3), 321–338.

Prietula, M.J. (March 12, 2002), "TrustMe 1.0: A Social Simulation of Trust, Advice and Gossip," Goizueta Business School, Emory University.

Prietual, M. and K.M. Carley (1998), "WebBots, Trust, and Organizational Science," in M. Prietula, K. Carley and L. Gasser (Eds.), *Simulating Organizations: Computational Models of Institutions and Groups*, Boston, MA: AAAI Press/The MIT Press.

Romanelli, E. and M.L. Tushman (1994), "Organizational Transformation as Punctuated Equilibrium: An Empirical Test," *Academy of Management Journal*, 37(5), 1141–1166.

Rivkin, J. (2000), "Imitation of Complex Strategies," *Management Science*, 46(6), 824–844.

Sastry, M.A. (1997), "Problems and Paradoxes in a Model of Punctuated Organizational Change," *Administrative Science Quarterly*, 42(2), 237–275.

Simon, H.A. (1981), *The Sciences of the Artificial*, 2nd ed., MIT Press, Cambridge, MA.

Thomsen, J., R.E. Levitt, J.C. Kunz, C.I. Nass and D.B. Fridsma (1999), "A Trajectory for Validating Computational Emulation Models of Organizations," *Computational & Mathematical Organization Theory*, 5(4), 385–401.

Van Alstyne, Marshall (2003), Indigo.

Weick, K.E. (1965), "Laboratory Experimentation with Organizations," in J.G. March (Ed.), *Handbook of Organizations*, Chicago: Rand McNally and Company, pp. 194–260.

Rich Burton is Professor of Business Administration at The Fuqua School of Business, Duke University. He is also Professor of Management at the EIASM (European Institute for Advanced Studies in Management) in Brussels, and Honorary Professor at the University of Southern Denmark.

His research focuses on organizational design, and particularly using computational approaches. With Professor Obel, he has authored numerous articles, and books. Their Strategic Organizational Diagnosis and Design: Developing Theory for Application is in its third edition. With the associated software, OrgCon, the book provides an integrated theoretical and practical approach to organizational design with a strategic focus.

Rich teaches executive MBA courses in organizational design, and international management. In the Ph.D. program, he teaches the theory course in organization theory, and an advanced course in Computational Organization Theory Applications.

He has DBA from the University of Illinois, as well as BS and MBA. He is active on a number of editorial boards. He is Senior Editor for Organization Science, and has been Department Editor for Strategy, Organizational design and Performance for Management Science. He has published some sixty articles on organizational design and management science, and seven books.