Durham Research Online

Deposited in DRO:

01 December 2017

Version of attached file:

Published Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Polos, L. and Adner, R. and Ryall, M. and Sorenson, O. (2009) 'The case for formal theory.', *Academy of management review.*, 34 (2). pp. 201-208.

Further information on publisher's website:

<https://doi.org/10.5465/AMR.2009.36982613>

Publisher's copyright statement:

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

INTRODUCTION TO SPECIAL TOPIC FORUM

THE CASE FOR FORMAL THEORY

RON ADNER
Dartmouth College

LÁSZLÓ PÓLOS
Durham University

MICHAEL RYALL
University of Melbourne

OLAV SORENSON
University of Toronto

This special topic forum contains seven papers that illustrate many of the ways in which management researchers can use formal tools—mathematical methods, simulation, and formal logic—to develop management research. Here we offer an overview of these methods and their advantages as tools for theory building.

A perusal of the pages of the *Academy of Management Review* over the past decade might lead one to conclude that formal methods have no place in the development of management theory. For example, we could find only one article published in this journal between 1998 and 2007 that actually used a formal approach—analytic methods, a simulation, or formal logic—to build its propositions. One therefore might naturally wonder why formal approaches have not been more prevalent in management research.

One possibility is that formal methods have little to contribute to management theory. That explanation seems unlikely to us (and not just because the articles in this issue serve as examples that strongly contradict it). Consider the disciplines on which management theory most commonly draws: economics and sociology. In nearly every theoretical article over the last decade that appeared in the *American Economic Review*, the authors used formal methods (almost exclusively analytic models). Even in sociology, during the last ten years the *American Journal of Sociology* (*AJS*) featured at least ten articles in which the authors built theory using formal approaches (divided fairly evenly across the methods featured in this special topic forum). Given that *AJS* publishes far fewer articles

per year than *AMR* and that more than half of those articles are empirical, the ratio of ten to one would actually appear to understate the extent to which sociology relative to management favors the use of formal methods. Seeing as how closely related social sciences have gained greatly from formal approaches, one might expect that management theory could similarly benefit.

A second possibility is that management scholars have not invested in acquiring the skills required to use formal methods. That explanation, too, appears to fall short, especially when one considers the response to the call for this special topic forum. We received a total of seventy-five submissions—a wealth of riches. Almost every one of these submissions, moreover, came from someone in a business school, and more than 90 percent of the authors hold doctoral degrees from business schools (rather than from, say, economics departments). We also benefited from the advice of an extremely high-quality pool of reviewers, many of whom had never before been involved with *AMR*. This overwhelming response surprised even us, the STF guest editors. But it nonetheless reveals that a large segment of management scholars have the training to pursue formal approaches to building theory, and, thus, a scarcity of

skills cannot account for the infrequent use of these tools in management research.

Although it falls beyond the scope of this introductory essay to determine precisely the reasons for the paucity of formal theory building in management research, we speculate that at least two factors come into play. First, we suspect that readers, reviewers, and editors may not fully appreciate the relative advantages of formal versus informal approaches. They may therefore feel that the investment involved in reading and understanding these formal approaches is too dear. Second, those who do use formal methods may not consider *AMR* as a potential outlet simply because it has published so few articles using these methods in the past. At least partly in support of this second point, we note that *Management Science* and *Organization Science* regularly publish papers that use analytic methods, simulations, and formal logic to build management theory.

With this STF we hope to stimulate both the demand for and the supply of management research using formal approaches. By describing formal approaches' advantages and by demonstrating them through example, we hope that readers and reviewers will become more attuned to the value of research that adopts a formal approach to theory building. And by demonstrating that *AMR* has an interest in publishing these papers, we hope to encourage management researchers to use them more often.

THE BENEFITS OF BUILDING THEORY FORMALLY

Formal approaches to theory building come in three main flavors, defined primarily by the methods used for articulating assumptions and then moving from those assumptions to propositions. In mathematical models the researcher begins by outlining a set of mathematical conditions that describe the phenomenon of interest and then uses mathematical proofs to demonstrate propositions that follow from them. In simulations (also known as computational models) the theorist similarly writes a set of conditions that represent the assumptions of the model, but instead of validating their implications via analytical proofs, he or she generates outcomes computationally across ranges of values for the parameters that determine those outcomes.

With formal logic, meanwhile, the researcher translates natural language assumptions into statements in symbolic logic and then uses methods such as truth tables to prove the propositions implied by those statements. Although these approaches vary somewhat in their strengths and weaknesses for addressing particular theoretical questions, they nonetheless share at least three strengths relative to verbal theorizing (the natural language approach most commonly pursued in management research): (1) precision and transparency, (2) logical consistency, and (3) an ability to identify unanticipated implications.

Precision and Transparency

Debates over the meaning of certain terms in the management literature have become notorious. What constitutes a "resource" in the resource-based view? What is a "dynamic capability"? What does it mean to "learn" or to be "boundedly rational"? Although these disputes occasionally converge toward one accepted definition, they can just as easily continue without any apparent resolution, despite decades of debate. Management, moreover, is not alone in this situation; all of the social sciences have had vigorous (and frequently unresolved) arguments over the meanings of key concepts.

In large part, these debates over terms and the applicability of theories reflect the imprecision of natural languages, such as English. Natural languages have evolved over time and continue to change. For a variety of reasons, many, if not most, words in natural languages therefore have multiple meanings. Listeners and readers attempt to infer which one of these meanings the communicator intended by considering the context in which they find it. But this process leaves ample room for ambiguity and subjective interpretation. One listener or reader might decide that a phrase means something very different from what another does.

Although one can attempt to communicate ideas clearly through natural languages, the very flexibility and instability of these languages make doing so an uphill battle. In contrast, mathematics and formal logic have been designed with precision in mind. Because of this precision, translating a natural language idea into a formal, symbolic representation can be challenging. Whereas one can state an assump-

tion or a proposition in natural language without even fully understanding the meaning of the statement, formal languages require that the communicator understand the concept deeply. In fact, we assert that much of the value of using formal methods stems from the fact that it forces the researchers using them to think thoroughly through the concepts they invoke. Once the theory has been translated into a formal language, everyone who sees it knows exactly what it means. Although the informal interpretation of the terms of the formal theory might vary from reader to reader, the formal characterization ensures that all of these interpretations have a common core.

That precision has at least two advantages in management research. First, it facilitates the accretion of knowledge. If different researchers use the same term to refer to several different concepts when they develop theoretical propositions, then it becomes exceedingly difficult to understand whether and how their arguments interact. Sometimes researchers duplicate effort by forwarding fundamentally equivalent ideas under new names. At other times progress comes to a halt when researchers appear to arrive at contradictory conclusions (perhaps because they use the same natural language terms but imbue them with different meanings).

Second, empirical researchers can more easily and accurately test theoretical propositions when little ambiguity surrounds the meaning of—and, hence, the appropriate measurement of—the central concepts. Progress in empirical research comes primarily from the ability to falsify theories. But without precise propositions, one often cannot say whether an empirical test refutes a theory or whether it simply fails to operationalize it adequately.

Logical Consistency

No one intentionally forwards a logically inconsistent argument, but it is nonetheless surprisingly easy to do. For example, one might not recognize an implicit assumption and therefore might fail to identify some critical condition necessary for the theory to hold. Or one might miss the fact that two (or more) assumptions act in opposite directions. Or, less crucially, one might simply impose an unnecessary and possibly redundant assumption. As theories become increasingly involved and incorporate more and

more factors, the opportunities for errors of omission and commission inevitably increase.

But even if the assumptions have been correctly specified and stated, people still regularly come to the wrong conclusions about the implications of those assumptions. The problem is that when relying on verbal theorizing, people rarely think logically through each stage of the chain of an argument; rather, they rely on their intuitions to identify the implications. Our intuitions, however, often prove wrong. Consider the Monty Hall problem (Selvin, 1975). A contestant on a game show can choose from curtain A, B, or C. Two of the curtains have goats behind them and one has a car; if the contestant picks the one with the car, he wins it. Let's assume that he initially chooses curtain A. The host then opens curtain B, which has a goat behind it, and asks the contestant whether he would like to choose a different curtain. Should he switch? Most people believe that each of the remaining two curtains has a 50 percent chance of having the car and, therefore, that the contestant should be indifferent between staying with his original choice and shifting to curtain C (Granberg & Brown, 1999). But mathematics demonstrates clearly that the contestant should switch. In fact, his original choice, curtain A, has only a one-third chance of having the car behind it, whereas curtain C has a two-thirds probability of being the winning curtain (because the host intentionally chose a curtain that does not have the car behind it, thereby revealing information about the other option). One could therefore easily imagine that someone developing a natural language theory of the Monty Hall problem might posit an intuitively appealing but incorrect proposition. More generally, researchers who rely entirely on verbal theorizing have little to prevent them from arriving at errant conclusions.

All formal approaches share the advantage of having tools available for ensuring logical consistency. Analytic models and formal logic use proofs. Simulations meanwhile calculate the implications of a variety of scenarios numerically. In any of these approaches, researchers must explicate all of the necessary assumptions, because without them they could not derive the results. Redundant and superfluous assumptions, moreover, become relatively obvious when writing proofs (although those using simulations cannot spot them so easily). Proofs and

computational methods, furthermore, ensure that the results do follow from the assumptions. In the Monty Hall example, one could have used any of these approaches to arrive at the correct answer. The researcher might make an error in the application of these methods, but, because the rules are clearly defined, readers and reviewers alike can easily audit the assumptions and the chains of logic leading to the propositions.

Unanticipated Implications

The availability of these tools has another added benefit: one can use them to derive unanticipated implications of the assumptions. Most exercises in verbal theorizing involve recognizing a relationship between two factors—say, resources and profitability—and then trying to determine a set of assumptions that would lead to such a relationship. In essence, one moves backwards from propositions to assumptions. A clear disadvantage of this approach is that one cannot find the unexpected (though it may later rear its head in empirical research testing these propositions). In contrast, formal approaches can help the researcher identify even the most surprising and counterintuitive implications that might follow from a set of assumptions.

Consumers of research using formal approaches often cannot fully appreciate this ability to identify unanticipated implications, for two reasons. In some cases, when these implications clearly run counter to reality, they may force the theorist to revisit the initial premises of the model and, therefore, may appear only as an intermediate (and unpublished) stage of the research. In other cases, by the time they appear in print, not only have these implications been derived but the intuition underlying them has been explicated and therefore the results—though originally surprising—may even seem somewhat obvious (if only in retrospect).

But the value of this ability to explicate more completely the implications of a set of assumptions can be seen in the breadth of propositions derived from relatively simple models. Readers new to formal approaches are often surprised by the complexity of the behaviors that emerge from relatively few assumptions. Simulations of the behavior of flocks of birds, herds of animals, and schools of fish, for example, have demon-

strated that a very small number of rules (as few as three) can generate realistic simulations of groups of animals in motion (Reynolds, 1987). In management research, Postrel (this issue) also illustrates the wide range of implications that one can derive from a small set of assumptions.

Beyond simply generating additional interesting propositions, this ability to identify unintended implications also accelerates progress in understanding a phenomenon. Additional propositions offer empirical researchers more targets—more opportunities to falsify the theory. If the propositions follow from the assumptions and empirical analysis falsifies one or more of these propositions, then this suggests that one or more of the assumptions must not hold as well. Researchers must then revisit their theories.

ANALYTIC MODELS, SIMULATIONS, AND FORMAL LOGIC

Despite these advantages, each approach to formal theory building is not without its limitations. The use of mathematical methods, for example, sometimes requires unappealing assumptions. Simulations have the disadvantage of only being able to prove the existence of a relationship between a set of inputs and an outcome; they cannot determine the necessity of a set of conditions. Formal logic, meanwhile, does not as readily accommodate functional forms of relationships and, thus, cannot always adjudicate between the expected effects of countervailing forces.

Each of these limitations nevertheless applies only to a particular formal approach. These trade-offs, therefore, might constitute a reason for choosing one formal method over another, but none of them suggests that one would want to eschew formal methods altogether. Although it is beyond the scope of this introduction to review each approach in detail, let us expand a little on the relative strengths of each.¹

¹ For those who wish to delve more deeply into one or more of these approaches, we recommend Harrison, Lin, Carroll, and Carly (2007) for simulation and Hannan, Pólos, and Carroll (2007) for formal logic. Mathematics includes such a vast array of approaches and tools that we cannot recommend any single survey.

Analytic Models

Some see analytic models as the “gold standard” of formal approaches. The translation of assumptions and definitions into mathematical statements and then the use of proofs to validate propositions mean that these methods impose a strict discipline on clarity and logical consistency. The use of mathematics, moreover, allows one to model the interactions of many, many “moving parts”—pieces of the model that influence its behavior—with the confidence that one’s conclusions do indeed follow from one’s premises.

Both the greatest strength and the greatest limitation of analytic methods stem from the need to produce a model with analytically tractable solutions. On the positive side, mathematical methods encourage the theorist to distill a situation to its essence. Consider, for example, the prisoner’s dilemma game. This scenario elegantly captures the problem of shirking in team cooperation. It does not describe all of the complexity of a real-life encounter, but it does capture the problem’s essential features. By abstracting away from the details of a particular case, it provides deep insight into an entire class of situations. Indeed, the art of analytic analysis appears in the ability to identify, to abstract, and to match to a mathematical approach the essential elements of the situation being modeled. But sometimes the depiction of a situation does not fit well with any analytical framework. Attempts to force it into one may therefore result in a theory with little face validity.

Simulation

In situations that elude mathematical methods, simulation may provide an attractive alternative. Before moving forward, we should note that although simulation research sometimes refers to particular sets of inputs as “experiments” and the results of these experiments as “data,” simulation is almost always an exercise in theory building; it is not empirical research.

Simulations run the gamut, from the simple to the complex. The primary advantage of these computational models relative to analytic methods is that they allow much greater flexibility in certain sorts of assumptions made by the theorist. For example, simulations can accommodate

any functional form linking two (or more) variables that a theorist might imagine. But simulations also have their own constraints. The investigation of some classes of problems—such as games with an infinite number of periods—remains computationally intractable, even with access to colossal computer power. Whether mathematical methods or simulation offers the preferable approach, therefore, depends on the problem.

Relative to analytic models, simulations have at least two disadvantages. First, because the method itself imposes relatively less discipline on the theorist, these models can become so complex that their inner workings end up being opaque. One can still observe how outcomes vary with changes in the inputs, but the interactions become so convoluted as to prevent the researcher from being able to determine *why* these relationships exist. When the individual components of the simulation are well understood, however, even such highly detailed and complex simulations can be useful. For example, physicists often build incredibly complicated simulations to determine the expected behavior of large-scale systems. They understand the behavior of each component—for example, a molecule or an elementary particle—extremely well, and they also have detailed knowledge of how small numbers of these components interact, but they must nonetheless rely on simulations to predict the behavior of hundreds, thousands, millions, or even billions of these components.

Second, simulation does not constitute a proof. Simulations sample a finite, although potentially very large, number of cases from an infinite set of possibilities (at least when the model includes continuous parameters). One therefore can say that some specific set of assumptions leads to a particular outcome, but one does not know whether any of those assumptions are necessary, or whether some unanalyzed set of assumptions might also prove sufficient to yield the results. Sampling the parameter space as widely and with as fine a resolution as possible can help to allay these concerns, but it does not alleviate them.

In many cases, theorists can combine simulation methods with analytic techniques, capturing a little bit of the best of both approaches. A researcher might, for example, build a simplified model that does not incorporate all of the

interesting features of a phenomenon but that does fit within the restrictions of some mathematical framework. After characterizing that model, he or she might then exploit computational methods to relax some of the assumptions or to extend the features of the model. Such an approach nicely exploits the ability of analytic methods to identify both necessary and sufficient conditions while still allowing the consideration of those cases that stymie analytic analysis (but that remain amenable to simulation).

Formal Logic

In many ways, formal logic is similar to mathematical methods. Indeed, mathematics uses first-order logic. Since many of the applications of formal logic to management and organizations also rely on first-order logic, such as Péli's (this issue), formal logic papers differ little in their relative strengths and limitations from papers using mathematical methods.

The biggest difference between formal logic and both analytic and simulation methods arises when researchers select or define alternative logics (i.e., other than first-order logic). In management research, for example, Hannan et al. (2007) have developed a nonmonotonic logic—a formal language—for theory building (for an example of the application of this logic, see Kuilman, Vermeulen, and Li, this issue). This logic includes a number of modifiers and relationships that allow the researcher to represent more accurately the incompleteness of our empirical understanding of a phenomenon. For example, it allows quantifiers like “normally” and “presumably” that first-order logic cannot accommodate (and defines a consistent set of rules for dealing with these quantifiers). Thus, it can provide a more realistic basis for developing and combining theory fragments describing phenomena that still require a great deal of empirical exploration.

Although the ability to develop the language used to build theory might seem to reduce dramatically the discipline imposed by formalization on theorists, logicians have developed a strict set of rules that all formal languages should meet to ensure their precision and completeness. All formal logics therefore have the rigor of mathematical methods, requiring precise assumptions and definitions and allowing

one to prove propositions and to identify necessary and sufficient conditions.

Formal methods provide powerful tools for developing theory, but we would not claim that verbal theorizing—that is, the use of natural languages to build theory—has no place in management research. Whereas formal methods operate deductively, natural languages more readily accommodate induction and intuition. Hence, when an interesting phenomenon is first identified, conjectures, expressed in natural languages, can provide a useful starting point. These conjectures can help to guide empirical researchers in terms of what data they should collect and what questions they should attempt to answer. As the phenomenon becomes better understood, researchers can then revisit the theory with formal tools to develop a more precise, transparent, and logically consistent description of it.

THE PAPERS IN THIS TISSUE

As noted above, we had an embarrassment of riches from which to choose. We selected papers for this special topic forum not just for their high quality and the importance of their managerial implications but also for their ability to highlight the variety and value of formal approaches.

Because of the breadth of papers in the special issue, we could have ordered them in multiple ways. We ended up choosing a rough grouping by subject, but the papers have many other points of contact.

The first three papers deal broadly with the issue of entrepreneurial entry, but each approaches it using a different formal method. Alvarez and Parker model the decision problem of how a pair of entrepreneurs should allocate the rights to control a new venture. Although perfectly rational actors would give control to the higher-ability member of the team, these authors interestingly demonstrate that, under conditions of subjective rationality, the team will allocate control rights to the more optimistic of the two.² This misallocation of control rights can

² Subjective rationality refers to the idea that actors must behave in a way that appears optimal, given their beliefs and observations of the environment, but that may not be rational for the omniscient actor (Ryall, 2003).

lead start-ups to perform poorly. Given the well-known propensity of entrepreneurs to be over-optimistic, this result strikes us as particularly interesting.

Ganco and Agarwal, meanwhile, construct a simulation so as to develop a better understanding of the performance differentials between start-ups and diversifying entrants, and how those differentials depend on the environment. They build their simulation on a variant of the NK model that has been widely used in management research (e.g., Rivkin, 2000). Their model, extending the counterintuitive idea that adaptation actually hurts organizations in turbulent environments, can account for all of the existing empirical regularities that have been found and suggests several new propositions that researchers should explore.

Finally, Kuilman, Vermeulen, and Li consider how the nature of the environment at the time of founding influences the life chances of entrants. Empirical research has usually found that firms that enter during periods of intense competitive crowding suffer higher mortality rates. But some studies have also found the opposite pattern: lower mortality rates among those that enter during times of higher density. The articles reporting these findings appear to arrive at contradictory conclusions, but Kuilman et al. demonstrate that one can incorporate them into a single theory by introducing the stage of the population's history as an additional factor. Whether density at the time of founding constitutes an advantage or a disadvantage depends on whether the population remains in the legitimating phase or whether it has achieved legitimacy and interactions have become primarily competitive. In essence, the authors introduce a scope condition.

We should note that the application of the nonmonotonic logic in this paper is particularly instructive. Whereas the logic that Kuilman et al. use would have allowed them to resolve these conflicting findings simply by assuming that legitimacy takes precedence over competition (or vice versa), they avoid this easy solution because they have no good substantive reason to introduce a precedence ordering between competition and legitimacy. Therefore, they wisely avoid such an ad hoc assumption.

The next two papers investigate, again broadly, the problem of coordination and incen-

tives. Postrel introduces the novel idea of goal ambiguity into the well-understood effects of substitution versus complementarity in joint production. He shows that goal ambiguity can lead to a failure in the allocation of effort, even when both actors have incentives only in terms of their joint output. Although this failure looks like a coordination problem, standard solutions to coordination failures, such as sequencing actors' decisions, may not solve this problem and may even exacerbate it. The managerial implications are clear: coordination and goal clarification have synergistic effects on team performance.

Makadok and Coff similarly explore an issue of clear managerial importance: How can a principal (an owner) ensure that agents (employees) do not free ride? In developing their theory, they point out that the manager actually has at least three levers under his or her control in achieving this outcome: authority, incentives, and ownership. All of these vary from being more firm-like to being more market-like. Managers therefore have eight different forms of hybrid organization at their disposal, and these authors demonstrate that at least one of them should generate the optimal allocation of effort across the agents for any type of task.

Finally, we end with two pieces that provide new insight into macroorganizational behavior. Cowan and Jonard develop a model of alliance formation in which successful partnership requires that two firms have moderately similar characteristics. They demonstrate that this single assumption can generate a surprisingly wide variety of network structures, including a tendency for repeated interaction and small-world structures. Their results therefore raise the possibility that the need for somewhat similar partners, rather than social factors, may explain a number of interorganizational structures. Since one could imagine the development of a parallel model at the level of the individual, their results may also provide insight into interpersonal networks.

Péli, meanwhile, attempts to reconcile the idea that inertia and imprinting imply that population-level change should occur through the selection [failure] of organizations with the intuition that populations sometimes shift through the adaptation of their individual members. To resolve this apparent inconsistency, he suggests

that environmental change may either be quantitative or qualitative in nature. The first, a quantitative change, might force firms to adjust their intensity of various activities, but the second, a qualitative change, requires them to do entirely different things to survive. The former therefore allows adaptation, while selection reigns supreme in the latter regime. Péli's approach to resolving this conundrum recalls a similar solution that he posited for resolving when broad versus narrow scope would benefit organizations (Péli, 1997).

We applaud the authors included in this special topic forum for their creative and compelling contributions to management theory and for their able applications of formal methods. We also appreciate the efforts of all of those who took an interest in our call. We hope that the papers that appear here become but the beginning of a long line of interesting and insightful investigations that use formal approaches to develop management theory.

REFERENCES

- Granberg, D., & Brown, T. A. 1999. The Monty Hall dilemma. *Personality and Social Psychology Bulletin*, 21: 711-729.
- Hannan, M. T., Pólos, L., & Carroll, G. R. 2007. *Logics of organization theory: Audiences, codes and ecologies*. Princeton, NJ: Princeton University Press.
- Harrison, J. R., Lin, Z., Carroll, G. R., & Carly, K. M. 2007. Simulation modeling in organizational and management research. *Academy of Management Review*, 32: 1229-1245.
- Péli, G. 1997. The niche hiker's guide to population ecology: A logical reconstruction of organizational ecology's niche theory. *Sociological Methodology*, 27: 1-46.
- Reynolds, C. W. 1987. Flocks, herds and schools: A distributed behavioral model. *Computer Graphics*, 21: 25-34.
- Rivkin, J. W. 2000. Imitation of complex strategies. *Management Science*, 46: 824-844.
- Ryall, M. D. 2003. Subjective rationality, self-confirming equilibrium and corporate strategy. *Management Science*, 49: 936-949.
- Selvin, S. 1975. On the Monty Hall problem. *American Statistician*, 29: 134.

Ron Adner (ron.adner@dartmouth.edu) is associate professor of business administration at the Tuck School of Business, Dartmouth College. He received his Ph.D. from the University of Pennsylvania. His current research interests include demand-based approaches to strategy, diversification, and innovation ecosystems.

László Pólos (laszlo.polos@durham.ac.uk) is professor of organizational theory at Durham University. He received his Ph.D. from the Hungarian Academy of Sciences. His current research focuses on translating the theory fragments of organizational ecology into formal logic and using that foundation as a basis for extending the theory.

Michael Ryall (m.ryall@mbs.edu) is associate professor of strategy at Melbourne Business School, University of Melbourne. He received his Ph.D. from UCLA. His research focuses on the formal theoretical foundations of strategy.

Olav Sorenson (olav.sorenson@rotman.utoronto.ca) is the Jeffrey S. Skoll Chair in Technical Innovation and Entrepreneurship and professor of strategic management at the Rotman School of the University of Toronto. He received his Ph.D. from Stanford University. His research interests include entrepreneurship and economic geography, organizational learning, organizational ecology, and the management of science and technology.

Copyright of *Academy of Management Review* is the property of Academy of Management and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.