

Charles Sabel

COLLEAGUES:

What you have before you is the skeleton of a research agenda directed at what I hope are common concerns. It is in the form of an introduction to a paper, because we are at a conference, and papers are the conventional form of communication at conferences. I hope the form and the incompleteness of the paper will not distract you from the substance of the research agenda. To reduce the chances of misunderstanding, let me just state the core propositions at the outset:

Faced with a set of new challenges, the administrative state in the US, the EU and Latin America is evolving away from hierarchical and towards what we can loosely call decentralized, network forms of organization

If these new types of administration are to be democratically accountable, they will have to respond to a new kind of public sphere that differs in two ways from the traditional, Habermasian public sphere that emerged in association with the nation state.

Where the Habermasian public sphere aims to influence legislative decision-making, a public sphere that holds network administration accountable will have to be capable of monitoring and participating in decision-making, as much as influencing the general legislative frame within which that decision making occurs.

And where the Habermasian public sphere is directed to national issues and decisions, the new public sphere will have to link local monitoring and participation with higher, or more encompassing engagements, including the ability to revise decision-making frameworks.

There are fragmentary and fugitive signs of the emergence of such a new public sphere in the US and the EU.

But it is in Latin America that the new type of democratic public seems most likely to develop, because of the current co-evolution there of the public sphere and of public administration.

Thus countries such as Brazil are developing a potentially distinctive public sphere in response to the democratization of the oligarchic state, not the formation of the nation, even while state administration is increasingly “networked” as national economies, and the communities they shape, integrate into world markets.

But whether the post-oligarchic state will indeed be capable of effective, networked problem solving, and whether the emergent public sphere can indeed hold that state democratically accountable by monitoring and participating in decision making across a sufficiently broad range of policy arenas, while linking the relevant scales of action in each, are two very open questions.

So let’s get to work, taking a close look at developments in Brazil from this perspective and using what we learn from this investigation to prod discussion of these themes in the US and the EU. One place to start might be the role, if any, of neighborhood educational association in school, reform, or of locally based watershed monitors in pollution control.

But these suggestions are meant only as illustrations of research projects that would test the general claim that new accountability mechanisms and new problem-solving institutions are co-evolving. The real questions, of course, are whether networked administration has the problem solving capacities many, myself included, attribute to it; whether the post-nation public sphere can indeed democratically domesticate the administrative networks, and whether that possibility, if more than a pipe dream, is being, or could plausibly be realized in Latin America or anywhere else at present.

The Introduction to the following “paper” note presents a fuller version of the same argument as it emerges from a comparison of current discussion in the US, the EU, and Latin America. It also contains pointers to recent work of my own that develops some of the key assertions. The next, and for now last, section of the “paper” presents a compressed overview of the operating mechanisms and foundational principles of the network or pragmatist organization as it is currently diffusing in both the private and public sectors. I include this because I think that “network administration,” though loosely used, does in fact name a general and fundamental organizational innovation that can be precisely characterized. This characterization suggests, I argue, a reorientation in our understanding of governance, which may in turn direct where we look, and what we look for in a “new” public sphere. Although the material included raises the governance question and goes some way towards characterizing pragmatist or network governance, it breaks off

inconclusively, without raising, much less answering these final, crucial questions regarding the character of the new public sphere. Even if I could presently do more than wave my hands at these questions—which it far from certain—any serious attempt to answer them would give you the easy job of criticizing my insufficiencies rather than the hard one of advancing what do very much hope is a common project. I look forward to the discussion.

I. INTRODUCTION

The administrative state is challenged today in new ways. It must regulate exchanges with networked industries such as telecommunications and electric power, whose structure and possible evolution is often opaque to the primary actors themselves, and among networked firms, whose boundaries and obligations to each other are constantly in flux. It must protect the public against diffuse and rapidly shifts threats from, for example, the “non-point source” pollution that runs off countless farms, homes, and construction sites, and from new, rapidly mutating pathogens spread quickly through food supply chains by global commerce. As economic volatility—call it globalization—makes the risks of unemployment and other threats to citizen well-being effectively uninsurable, the administrative state must restructure such traditional public services as schools, continuing education programs, and job placement agencies, substance abuse and day care programs to “capiccate” individuals and families—especially those who fare poorly with the “standard” offerings—to find their way not just in the economy, but in life.

There are sharp, revealing divides in the US, the EU and Latin America as to whether the administrative state can meet these challenges, and, if so, whether its eventual successes will increase or (further) diminish the democratic accountability of public administration. In the United States authoritative voices (Jerry Mashaw, Bruce Ackerman—for present purposes the Yale School) suggest that the New Deal state is unlikely to adjust to the new circumstances. Their empirical touchstones are the Environmental Protection Agency and the National Highway Traffic Administration of the Department of Transportation, entities created in the late 1960s and early 1970s, at the height of Congress’s infatuation with top-down, command and control methods. From this vantage point the administrative state will fail, and fail unaccountably, at the new tasks for the same reason is largely failed, with little accountability, at the old: It depends on the expertise of officials who know vastly more than citizens, legislators, outside auditors and judges about what needs to be done, but still not nearly enough to solve the problems at hand. And it is subject to many, conflicting forms of legislative, judicial and managerial oversight that exacerbate, or at least make it impossible for agencies to correct the very problems the oversight authorities identify. Given these limits the best and perhaps

only way to secure social justice in a world that admits of neither effective social insurance nor the provision of capacitating services is to endow disadvantaged citizens with resources they can use to purchase privately goods and services that better their prospects.

Other, equally authoritative voices, such as Richard Stewart's, hold on the contrary that the New Deal state is not, and perhaps never was, as ossified as the Yale School and others maintain. Looking both to procedural innovations that might have been (in response, for example, to pressure from the judicial hard look) and at current developments at the EPA and elsewhere, Stewart sees the US administrative state as potentially able to serve new ends by adopting some combination of permit trading regimes to minimize the cost of social compliance with general regulations, and introducing non-hierarchic, networked types of decision making to allow general rules to benefit from, and in turn inform, local experience. My own work on the reorganization of public schools, long regarded as a paramount example of the unreformability of public bureaucracies, goes in the same direction. But Stewart, and other equally close students of recent developments of the US administrative state wonder whether the increased efficacy of networked administration may not come at the price of reduced public accountability, not least because the network arrangements blur the distinction between governed and government, allowing the latter to slip the control of the former.

These different assessments of the potential for administrative reform are reflected in the different reactions of the high federal judiciary on the one hand and the lower federal bench and State courts on the other. In areas such as school reform, fair (low-income) housing, and the reform of mental health institutions State courts have routinely assumed the feasibility of and insisted on organizational reform. On the authority of State constitutions or legislation they have required poorly performing institutions to meet general goals, such as the provision of an "adequate" education to poor, minority students as well as to rich, non-minority ones; to create systems for measuring progress towards those goals; to report the results of periodic measurement to the public, in ways that allow comparison of how much—or little—similarly situated institutions better the condition of at-risk groups; and to ensure that low performers learn from more successful ones. This form of judicial oversight, often undertaken in cooperation with legislative reform of the accountability system to which the affected institutions respond, seems to grow out of and encourage the diffusion of "network" administration. (on this see the paper with Bill Simon on destabilization rights just published in the Harvard Law Review and available in draft at <http://www2.law.columbia.edu/sabel/papers/Chayes20clean.doc>)

The high federal bench, on the contrary, is fixated on establishing the accountability, and limiting the discretion of administrative agencies, to the exclusion of any consideration of administrative efficacy. Textualist canons of interpretation, limiting

agencies to classes of actions expressly authorized in legislation and insistence that judges defer to administrators are acting only with the intent of making law are only two of the many doctrinal strategies deployed to this end. Such judicially prompted clarification of the sovereign pedigree of administrative decisions—express and legitimate delegation by Congress to an agency to make legally consequential rules, followed by the express declaration of a politically accountable administrative official that a decision is to have this effect—could well have the effect of centralizing decision making. But such is the indifference of both the high federal judiciary and the academics who comment on their decisions to question of administrative structure that is impossible to know whether this is an intended outcome, or simply the side effect of a general view that, despairing of any possibility of reform, simply seeks to minimize the harm that administration can do to democracy.

In the European Union the view of the administrative renewal as feasible, even likely, but at the cost of democratic legitimacy, is more nearly a mainstream position than a countercurrent. Renewal seems feasible because in creating the Single Market while protecting the citizens of member states against many, changing threats to public health and safety, the EU has *already* established a web of novel, hard to characterize institutions—“comitological” technical advisory committees, networked observatory-agencies, the “open method of coordination” for comparing and revising national practices in many policy areas—whose overall, and generally acknowledged effect is to allow just the contextualization and periodic revision of general rules that might address the new challenges to the administrative state. EU experience broadly seen has been that the most satisfactory way to “harmonize” the regulatory systems of the member states is by instigating a process of mutual learning and monitoring in which the discovery of new possibilities changes the actors’ preferences and leads to outcomes different from—and generally superior to—those that would have been produced by negotiation alone. Indeed the fact that the “winning” coalition in regulatory matters typically does not get what it initially thought it wanted is a good part of what makes the whole complex process politically sustainable. (For more on this, and reading of the EU developments as deliberative polyarchy—a type of network governance—see the paper with Josh Cohen at <http://www2.law.columbia.edu/sabel/papers/Cohen-Sabel%20final.doc>)

But this welcome increase in technical problem-solving capacity is widely regarded as a potential threat to democracy in the EU. It is unclear to what democratic public the problem solvers are to respond, and the currently compelling ways of constructing such a public seem at odds with basic features of the emergent system of network governance. Thus the EU famously lacks a *demos*: a natural community of language, sentiment, and historical experience coherent enough to embody itself, and its solidary commitments, in a parliament. On one reading the story ends there, because natural communities are the precondition of nation states,

which in turn are the precondition for political and social democracy. But a second, and historically better informed reading argues that the democratic nation state in fact arose from a circular relation between traditional contexts—substantive *demoi*—and a nascent public sphere: a space for the open discussion of the broad purposes of public power among equal persons ever more free of overriding group solidarities and everyday practical arrangements. Opinions formed in this public sphere direct political controversy, which shapes legislation, which orients the administrative apparatus, which operates to suffuse the original, traditional attachments with an aspect of civic or political loyalty and so broaden and deepen the public sphere. “Modern democracy and the nation-state have developed in tandem,” Habermas, the leading exponent of this view, writes:

Both have jointly produced the striking innovation of a civic solidarity that provides the cement of national societies. National consciousness emerged as much from the mass communication of formally educated readers as from the mobilization of enfranchised voters and drafted soldiers. It has been shaped as much by the intellectual construction of national histories as by the discourse of competing parties, struggling for political power.

To “conserve” (Habermas’ word) the achievements of the democratic nation-state at the EU level it, is from this vantage point necessary to construct a European public sphere by refocusing media attention on “European” developments—so citizens in the various member states are discussing the same things at the same time—and by endowing the EU with political institutions similar to those of the nation-state—so citizens, from inveterate civic habit, will know (how) to address their public concerns to them. Brussels becomes the notional, if not the physical capitol of the EU, towards which the influence of the public sphere is directed, and from which the power of law, coursing through administration, flows.

But this historically attractive way of creating the democratic “principal” in the EU immediately encounters two closely related objections. The first is that the distinction between public deliberation and practical or technically informed action characteristic of this notion of the public sphere requires just the kind of cascading delegation of sovereign authority—from democratic people to legislature, from legislature to administration—which US experience demonstrates to be unworkable. Even if we can assume, the Yale school notwithstanding, that the future of US administration is yet to be written, there is little in its past that warrants the assumptions that elected officials, acting on behalf of citizens, can present clear rules and standards to agencies, and that those agencies will confine their activity to the efficient implementation of those rules and standards, subject to external monitoring. More generally, ‘agency’ problems are the notorious weak

link in the minimalist views of democracy that accept any duly elected legislative authority as a legitimate ‘principal’.

More concretely—and here the second objection—EU regulation, and political deliberation broadly conceived, is *already* so distributed in administrative processes beyond the reach of conventional legislative oversight that the Habermasian conception of the public sphere cannot serve as a practical guide to reform. Presciently or inadvertently, as just noted, the EU has created a system of harmonizing administration that works as much “horizontally” as “vertically,” and more by federated influence of lower-level units on higher ones than by top-down determination of ends and means. If the EU public sphere turns out to be the public sphere of the nation state writ large, it is unclear to whom it is to address itself. Intuiting this dilemma, but unable to resolve it, the practical architects of political reform, convened in the Convention for the Future of EU have been careful not to write the new problem-solving institutions out of the draft constitution, without finding a way of establishing their democratic legitimacy by writing them in. Conversely, those familiar enough with the accomplishments of the new problem-solving institutions to count as their advocates are agnostic, verging on pessimistic, with regard to their democratic aptitude. As in the US I believe that close attention reveals a more subtle and promising interplay between administrative adjustment and effective public oversight. But for now the stark contrast between efficacy and legitimacy dominates public debate.

In key Latin American countries such as Brazil, Mexico and Argentina the prospects for accountable administrative reform are different again. If the administrative state is widely seen in the US as maladapted to current tasks but perhaps accountable to a vigilant judiciary, and in the EU as adaptable but unaccountable to any plausible public, in these countries the administrative state and the public sphere are evolving together in ways that suggest the possibility—but hardly the inevitability—of an adaptive administration responding to a novel public. Consider the case of Brazil as rendered by Avritzer and other participants in this meeting. On the one side the clientelistic, patrimonial state bequeathed from colonial times, and engorged with populism and management of import substitution, must open itself to democratic supervision while reforming to meet challenges different in degree, not kind, from those faced in the advanced countries. Hence a profusion of new, more or less “autonomous” entities at various levels of government, at times explicitly modeled on US examples, but often “networked” on lines that the EU debate tends to acknowledge even if the US debate does not. On the other side is a potentially new public sphere, or, more exactly, constellation of new publics, born not in the struggles of nation building but rather in the long fight to democratize oligarchic national states. The collection of publics or emergent public sphere, born of social movements that mistrusted the day-to-day clientelism at the street level of the old state

and reviled the anti-democratic isolation of its peak institutions, has two characteristics that mark it off from the Habermasian model. First, it is concerned at least as much to monitor specific state actions as to influence general policy; and when decision-making is decentralized, as it often appears to be, the monitoring shades into a form of co-decision making, or at least a kind of continuing, informed contestation of technical expertise by citizen users of public services. Second, the new publics tend to be local, or at least sub national, without being isolated in a particular setting. Municipal participatory budgeting as Avritzer presents it illustrates both features: citizen assemblies, with members drawn from different neighborhood-based associations, review the performance of municipal road construction agencies and water utilities, helping to determine the sequence in which districts are connected to the existing networks, and in some cases the balance of investments among road building, expanded water service, and other classes of municipal activity.

Could the new post-national public sphere that the democratic domestication of network governance seems to require arise from such a matrix? To answer that question we need to get a better idea of just what network organization is, and how it might effect regulation and the public sector generally. The next section goes some way towards addressing these issues. Then we need to step back from this account of the emergent form of problem solving to reflect on its possible implications for democratic accountability and democracy generally. Finally we could begin to assess Brazilian and Latin American developments more broadly in the light of our revised conception of what governance, and democracy, need to do. The “paper,” as noted, does not address these issues. But I very much hope the conference will. It is with that goal in mind that I invite you to read on.

2. THEE TYPES OF PRAGMATIST ORGANIZATION

In the century before 1980 the canonical organization was hierarchical and closed. Headquarters set goals, parsed complex tasks into simpler, more manageable ones; organizational routines specified how to parse, execute and review the compliance of tasks. Subordinates were rewarded for complying with instructions; middle managers checked the conformity of activity to plan. After 1980 the canonical organization is federated—decisions of higher units are shaped by lower ones and open—the lower units can be formally outside the organization. Or, to capture the idea that information in the new organization flows up and down as well as sideways, they are said to be networked. General designs are set provisionally by the highest level and revised in light of proposals by lower level units responsible for executing key modules or subsystems. The organizational routines define methods for choosing provisional, initial designs and revising them in the light of more detailed, partial proposals. Collaborators are rewarded for achieving broad goals according to

standards defined as part of the process by which the goals themselves are set. To get a more precise idea of how these new organizations work, and of the principles informing their design, consider three variants of the networked organization as it appears in both the private and public sectors.

Take first the *deliberately innovative organization*. This is the canonical case of the new, networked organization: the firm—it almost always is a firm—whose very purpose is to produce innovative products or services. Perhaps the best documented example is of design and production in the automobile industry. The process starts when, say, the new-van design team in an automobile firm sets the general performance characteristics of the vehicle by benchmarking the best features of current vans and exploring which innovations under development can be incorporated in its design. To benchmark the potential of developmental work the team may ask for engineering simulations of possible outcomes, “flash market” a product embodying a potentially valuable feature, or otherwise try to test the actual reaction of buyers to some approximation of the design they are exploring. Assessing the results of these probes, and again guided by reference to leading examples and comparison of possibilities, the team next provisionally subdivides or, to take a term from cognitive science, “chunks”¹ its general goals into subtasks—the design of an engine, or heating, ventilation and air conditioning system—and chooses a specialist team from inside or outside the parent company to realize the initial specifications. (It may seem unduly fussy to refer in this connection to “chunking” rather than the more familiar “modularizing,” especially since the latter is often used loosely in the automobile industry and elsewhere to mean the former. But modularization, more strictly speaking, aims at the creation of fixed, black boxes whose performance is durably defined in an interface listing the output returned for any input. Where the corrigible provisionality of design choices being describe here is useful, modularization *strictu sensu* is, as we will see below, risky, even impossible.) After this initial chunking separate project teams elaborate all the provisional subsystems concurrently, applying to that task the same kind of evaluation of competitors’ successful efforts and developmental possibilities used in the van team’s first round of benchmarking. In addition, they benchmark the production processes central to their eventual products to ensure that the methods employed will meet or surpass the efficiency of their most capable competitors. Engine plants, for instance, will have to produce engines that are at least as cheap and warrantable as those of competitors making similar engines in comparable volumes. Then the initial overall goals are modified by the methods of simultaneous or concurrent engineering, e.g. the engine-design group may find a way to better its target specifications or to cut its manufacturing costs if it can persuade other component groups that design characteristics should be modified accordingly. Refinement of this iterated co-design continues once production begins by means of just-in-time and the error-detection and

correction methods associated with it. In just-in-time production, parts are supplied to each work station only as needed: ideally, one at a time. Hence disruptions are immediately visible. A breakdown at one station halts production by stopping the flow of parts to downstream operations. To assure the flow of production, therefore, the source of disruption must be identified. This typically requires tracing long causal chains back to improbable origins by insistent questions sometimes called the “five why’s”. For example:

Why is machine A broken?	No preventive maintenance was performed.
Why was the maintenance crew derelict?	It is always repairing machine B.
Why is machine B always broken?	The part it machines always jams.
Why does the jam recur?	The part warps from heat stress.
Why does the part overheat?	A design flaw.

Thus error-detection and correction, like benchmarking and simultaneous engineering, reveal possibilities for improvement in unexpected (mis-) connections among the parts of complex endeavors; and the cumulative effect of these results is captured in improvements in the benchmark standards for various production processes.

Benchmarking, simultaneous engineering and error-detection methods like the “five why’s” are counter-intuitive from the vantage point of classical organization theory. In classical theory, routines are questioned only in exceptional circumstances. It is simply oxymoronic from this perspective to try to do what the new institutions do as a matter of course: routinely question the suitability of current routines for defining and solving problems. We can think of these new institutions as pragmatist: They systematically provoke doubt, in the characteristically pragmatist sense of an urgent suspicion that routines—habits gone hard, into dogma²—are poor guides to current problems.³

As a second instance of such pragmatist institutions consider the *High Reliability Organizations (HROs)*. HROs are well, but incompletely designed to perform without fail such extraordinarily demanding tasks as generating electric power through nuclear fission, launching and recovering jet aircraft rapidly from and back onto pitching, greasy flight decks, launching and recovering space shuttles, or fighting forest fires as they race through rough terrain. The designs are necessarily incomplete because the conditions under which the specified tasks are to be accomplished are changing continuously in more or less subtle ways. If the organization is not adjusted accordingly, it fails, catastrophically. In the pragmatist terms just introduced, HROs become disastrously unreliable if they assume that routine, (nearly) invariant success is the result of following invariant routines; and the organizational challenge is to avoid accidents day after day without imperceptibly making this assumption.⁴

Key to this are error-detection and root-cause analysis disciplines of the sort broadly familiar from the production/operations level of the innovative organization. In the setting of HROs, the most important and characteristic of these is Near-miss reporting and analysis. Near misses of course are accidents that only accidentally didn't happen. So the near misses—and beyond that “out-of-control” sequences that nearly produced near misses—are the urgent analog in the HRO to the line stoppages in a just-in-time system. Both trigger root-cause analysis meant to uncover not only the proximate cause of the incident, but to eliminate, through redesign of the organization if necessary, the background conditions which generated the immediate source of danger.

Note here a circumstance to which return later when considering governance—the root-cause analysis regimes are often regarded as a key element in the governance of HROs, and when that is the case they are subject to regulatory scrutiny. For example nuclear power plants in the US are rated in part on how effectively they respond to reports to reports of anomalies at other plants that might be of relevance given their own set-up, and so on.

New Public Services are third case of pragmatist organizations. This type of institution is little discussed in the organizational sociology literature, but is receiving increasing attention in various domains of policy studies. Well documented examples are schools in Texas, Kentucky, North Carolina and elsewhere in the US that actually teach poor children of color to read and do mathematics with proficiency comparable to that attained by rich, white pupils. The new organizations are like HROs in being designed to achieve a single, complex task—teach children to read and use mathematics—and in relying on error-detection to compensate for design deficiencies. But there is an important difference. Where the design gaps are in HROs are, roughly speaking, small in relation to the overall structure (that is why we say it is well designed) the new public service the situation is reversed: there is lots of “gap” and little by way of consolidated design. No one, in other words knows how to build a school system that produces the desired result. Instead the institutional solution is to build an organization that uses error-detection and correction at the lowest levels to find out what works, and then adjust the higher level structures to generalize that behavior and encourage more refined error detection, and so on. With regard to reading, for example, all students learn by some idiosyncratic combination of decoding strings of letters/phonemes (phonics) and derivation of the meaning of words and sentences from context (whole language method). Teachers identify the strengths and weakness of each student's mixture of strategies by sampling their skills in brief, daily sessions, and suggest improvements. (This might be called first-order error detection and correction.) The performance of students in the same grade is measured periodically state wide by a standard test, allowing for the comparison of the performance of teachers within schools, schools, and districts. (second-order error

detection.) The job of principals in this system is to create conditions in the school for generalizing the successes of the most successful teachers, and The job of the job of the principals' superior—the district supervisor—is to create conditions for diffusing the successes of the most successful principals.

In this way the reformed school is invented through the piecemeal, but eventually comprehensive improvement of a crude, but serviceably provisional starting structure that supposes only the broadest, non-vapid agreement on goals—educational achievement (by mainstream measures) should not vary across groups in culturally salient hierarchies, and gaps should be closed by leveling up, not down—and methods—teachers must aid students to improve their individual bundles of learning strategies, and administrators must aid teachers and other administrators in doing this.

Accepting now as a stylized fact that the new organizations operate successfully by application of these methods, and leaving for another time questions concerning the morphology of and methods for constructing the networks presumed in what follows, we return to the classic questions of bounded rationality, efficiency and governance and try to provide a consistent and conceptually plausible account of this success.

3. BOUNDED RATIONALITY, EFFICIENCY AND GOVERNANCE IN THE PRAGMATIST ORGANIZATION

RETHINKING BOUNDED RATIONALITY

The centrality of search routines to the new organizations suggests that there is an alternative to the decomposition of tasks as a solution to the problem of bounded rationality: Under volatile conditions, when it is given that no one can know the whole answer to the design question we face, the way to respond to bounded rationality is to find actors who are already solving (part of) the problem we will turn out to be trying to solve. To be workable any such problem-solving search (regardless of where in the overall design or production sequence it is conducted) will, at a minimum, have to Disentrench enquiry, directing attention to unfamiliar solutions, while it Produces sufficient information about advantages and disadvantages of rival possibilities—the solution space—to become Self limiting. If it doesn't meet these requirements the search will be uselessly redundant—when you don't know what you are looking for it is especially unproductive to look where you always look—or outright destructive, sending the actors off on an endless fool's errand.

Benchmarking and error detection do indeed oblige the actors to search for solutions in a novel, initially open, but ultimately bounded space of possibilities: the set of best current or potential designs (benchmarking) or the activity chains that

might have caused a particular breakdown (error detection). The initial canvas of design solutions is of necessity novel and open because it must consider responses that are “like,” but potentially better than current practice on at least some of many dimensions, and “like” has no determinate meaning in advance of a particular search. Indeed, to search for likes is to invite surprises. The point of the search is precisely to uncover an unsuspected but highly informative resemblance.

Consider by way of example this telling episode in the evolution of the modern paintbrush. A company is trying with little success to produce the traditional house painter’s brush with artificial bristles. Even an accomplished artisan can’t apply a smooth coat of paint with the synthetic bristles. The question arises: what is a paintbrush like? If a trowel, then the performance is determined by the shape of the bristles’ tip, and the company’s problem is to find an artificial equivalent for the split ends of natural hairs. If, however, paintbrushes are like pumps, performance is determined by the contour of the channels formed by the bristles, and the next step is to reposition the artificial bristles to channel paint correctly.⁵ In this case the brush was more like a pump than a trowel—but this result, of course, is provisional. In more artfully inclined hands a brush may need to be less a pump than a trowel, sling-shot or eyedropper.

As in this example, each possibility considered in metaphoric benchmarking provides a vantage point from which to evaluate the strengths and weakness of the others. Hence the search produces something “like” a provisional taxonomy or map of accessible solution strategies in relation to each other. In this way it is self-limiting: Once you have a serviceable map of the solution space, you stop doing cartography and decide where you want to go.

ROBUSTNESS AS HIGHLY DYNAMIC EFFICIENCY

Before we can ask about the efficiency of organizations in high-volatility environments, however, we have to clarify the question. Efficiency is a static concept, or rather a concept for stable environments. The efficient solution is just the least-cost means of reaching a known end. When the ends are various, and by some measure closely related, it can still be meaningful to speak of “an” efficient solution or a set of these. But when the ends become more various and less well specified it no longer makes sense to speak of optimum or efficient solutions. Instead we can ask whether our problem-solving technique is robust, meaning that it can be expected to produce workable answers turbulent task environments. Think of robustness as highly dynamic efficiency.

To underscore this point let me note here the misleading incompleteness of Simon’s well-known story about two watch-makers, Tempus and Hora, whose competition is intended to show the robustness (resilience in the face of disruption) of hierarchical modularization. Tempus and Hora, you will recall, made similarly

complex watches of about 1,000 parts each. But Tempus, the craftsman, assembles each watch from 1,000 pieces, while Hora, the mass producer, makes sub-assemblies of 10 pieces each and then combines these into the final product. Simon shows easily that unless the probability of disruption per unit time is extremely low, Tempus almost never completes a watch, and the typical disruption destroys an extremely valuable, nearly finished piece. Hora completes much more work and losses much less in each disruption, so modularization seems robust.⁶ And it is in comparison with craft production of large numbers of extremely complex machines.

But modularization is plainly not robust in our current world, where key module makers—meaning here sole-source suppliers of indispensable components—are routinely destroyed by natural catastrophes or the complete devaluation of a core capacity through a competitor's innovation. The robust strategy in this world is one where there is no difference between setting up the production system and repairing or replacing it in case of disruption; or, put another way, where the same principles that generate the initial set-up also generate innovative transformations of that first design as circumstance require. These are the principles that our robust producer, Quaesitor, follows when she searches for collaborators who are good at searching for collaborators all the way down. In this system, if any of the producers “breaks,” the broken one and its collaborators find at least one replacement or substitute.

The destruction, by fire, and the almost instantaneous, collaborative reconstruction of the capacity of a key Toyota supplier illustrates the robustness of Quaesitor-type re-chunkers. When it burned down in 1997, Aisin's Kariya plant 1 in Kobe was the extraordinarily efficient, sole-source supplier to Toyota of P-values, a relatively simple, but high-precision component of an anti-skid break-control mounted on all of the assemblers' makes and models. True to its just-in-time discipline, Toyota had only two-days' worth of P-values in inventory on the day the fire broke out; when that stock was exhausted, the assembly lines stopped. But by 3 days after the fire a congeries of a few Aisin suppliers and many other firms in the Toyota group—very few with prior experience in manufacturing this particular part, all presumably motivated by solidarity and the self interested desire to return to normalcy and distinguish themselves on the way—were producing more than a hundred types of P-valves in their own facilities. By 5 days after the fire two of Toyota's Kobe plants were reopened; 3 days later car production was at more than 90 percent of the pre-fire level; a week later there was no shortage of P-values at all. Ultimately 62 firms became emergency producers of P-values; these producers themselves relied on more than 150 emergency suppliers. So devastating had been the fire that none of the participants had access to the precision tooling that Kariya plant 1 used to make the original part. Aisin did contribute to the recovery by installing thousands of additional phone lines to respond to requests for information. But neither it nor Toyota attempted to direct the reconstruction effort.

The key to the extraordinary success was the participants' vast, common experience of the Toyoda variant of iterated co-design: Aisin could characterize the part and the production process in general terms—chunk them; the emergency producers could devise, starting from their own experience, many different ways (some quite innovative) of achieving these ends. Then they could chunk the processes they were contemplating so that their suppliers could do the same with respect to subtasks. So, as Quaesitor would have expected, the same disciplines that generated a network of “module makers” continuously searching for ways to improve also generated the capacity to search out alternative solutions if one of the network's module-making nodes fails. In a malign world this distinction between efficiency and robustness would translate into an impossible choice: we would have to either increase efficiency or increase robustness, without knowing from one day to the next which is called for. In a more benign world we would expect that increased robustness would increase efficiency as well, so that measures that help us enhance performance in a volatile environment also improve operations under more stable conditions. This is in fact what we observe with the new organization. I suspect that the overlap between robustness-enhancing and efficiency-enhancing mechanisms explains a good deal of the explosive diffusion of this new institutional type.. The consider as an illustration of this overlap benchmarking and the metaphoric learning associated with it. In a volatile environment we expect tomorrow's conditions to differ sharply from today's. So minimal prudence demands that we continuously benchmark new developments to avoid disastrous surprises. But we just saw that benchmarking searches inform us about *all* the possibilities in the—expanded—solution space. That means that so long as the environment has not cooled down to the organizational equivalent of absolute zero, where benchmarking itself is unnecessary because conditions never change, these comparisons tell us new things about solutions we already know even as they reveal solution-strategies we had not foreseen at all. This is surely an important part of the reason that comparison of additional variants often improves the reliability the eventual designs rather than degrading it.

Something similar is probably at work with the second mechanism: concurrency. Concurrent development—in which “upstream” and “downstream” steps proceed simultaneously, each taking account of the (changes in the) requirements of the other—is almost unavoidable when rapid change puts a prohibitive penalty on missing, or being late for a market, But like benchmarking, concurrent development sheds new light on familiar designs and practices even while illuminating new environments. In particular concurrency calls into question taken-for-granted assumptions about the relation among components in any subsystem by prompting chains of what-if questions about how perturbations in the environment would reverberate through the inner structure of the complex. So in all but the most stable settings,

some concurrency probably leads to efficiency-improving discoveries about shortcomings in familiar arrangements while also increasing robustness.

GOVERNANCE

Governance in the classic organization, we saw, is on the principal-agent model, or its stake-holder variant. The agent is punished for using her discretion to make self-interested use of gaps or imprecision in the principal's plan. She is rewarded for faithfully following that plan, or, better yet, using her discretion to fill gaps and correct mistakes in a way that realizes its master goal. Principal-agent governance is plainly unworkable in the pragmatist, search organization, because in such organizations final goals are the product, not the starting point of organizational activity. In this network setting the very distinction between principals who set plans and agents who execute them is blurred. "Agents"—meaning simply those whose relatively narrow responsibilities put them "lower down" in the organization—certainly can't be held to account for realizing the "principal's" original plan or complying with rules specifying their tasks because there was no plan, and hence nothing from which to derive detailed task assignments.

How, for starters, can accountability can be established if not by looking backwards? One alternative is to by looking in effect sideways and forwards, or rather sideways at how well others are looking forwards. An actor in a pragmatist organizations is acting accountability when she learns at least as much as other collaborations about the unfolding possibilities for advancing (and re-defining) the joint project, and discloses what she learns. Accountability thus becomes a matter of peer review: All the actors, or some relevant subset of them, assess the capability and intention of each to engage in iterated co-design or its production-level equivalents. Accountable behavior no longer entails compliance in the sense of rule following, but rather provision of a compelling explanation for choosing, in the light of fresh knowledge, one way of achieving the common goal over others. Hence the accountability systems of pragmatist institutions will be characterized by declaration of broad goals or performance standards that are explicitly treated as provisional and subject to revision through the actors' peer reviews of their own disciplined learning.

Such a rolling review of purpose and performance will only be possible if at least some of the information needed for the substance of collaborative problem solving in particular cases can be used for benchmarking the abilities and probity of current and potential partners. By making each party's facility with shared and highly revealing problem-solving techniques transparent to the others, pragmatist organizations make current collaboration (or the exploration of possible joint work) richly informative about the potential and risks of partnership. A firm that easily ferrets out the source of errors with five-why methods will presumably do better at solving new problems, and hence be more reliable collaborator, than one that gets lost

in the maze of its own confusions. Similarly, the mutual transparency that results from co-design disciplines makes it possible to detect, and attempt to correct, potential problems before they become disasters. Call the fusion of substantive and evaluative knowledge that allows an organization to learn as it monitors accountability (and vice versa) learning by monitoring. The explosive diffusion of pragmatist disciplines and the success of the new organization generally strongly suggest that such learning by monitoring is in fact workable at the level of individual projects or the operation of relatively small units, where “workable” here means only that show-stopping opportunism and inaptitude are detected, not that all interests are perfectly aligned, eliminating every trace of power. At the limit, if participation in the new disciplines did *not* facilitate detection of dangerously incapable or opportunistic partners, the spread of pragmatist organizations would have been quickly stopped by some variant of the hold up—the collaboration-destroying use of collaborative discoveries—that classical organization theory long took to be a fatal threat to just the sort of vertically disintegrated co-design commonplace today.

But the very connection between problem-solving substance and partner evaluation that makes learning by monitoring an instrument of governance on the project or small-unit level seems to rule it out as a mechanism for governing large and complex institutions such as multinational firms, public service providers or regulatory systems that encompass many different kinds of projects. Deciding to allocate scarce resources among these different, perhaps competitive projects, or to add entirely new ones, is of course a central task of such organizations. The ability of a project group or unit to demonstrate that it and its collaborators are highly capable of developing new tasks from current ones plainly bears on, but cannot alone decide, such questions. How, if at all, can such encompassing organizations select projects and maintain accountability by looking forwards and sideways. Put another way, is there some higher-level, project-spanning device, compatible with and perhaps analogous to learning by monitoring, for disciplining choices about the direction of collaboration and the identity of collaborators with the help of current, cooperative experience?

We can begin to outline the features of such a governance mechanism by listing the constraints it will have to respect in operation, just as we did in describing what turns out to be the closely related mechanism of dis-entrenching search. First, the forward-looking accountability mechanism must of course produce categories for comparing, and eventually ranking, different projects. Otherwise it is useless as a device for project selection. But second, in establishing this comparability, the accountability mechanism must encourage the exploration of new strategic goals, not restrict it by only rewarding search for better means to one or a limited range of given and unchallengeable outcomes. Recall that the aim of the metaphoric benchmarking key to the pragmatist institution is precisely to encourage the search for new solutions or projects outside the realm marked out by routine. A governance

mechanism that created disincentives to just the kind of searches the organization was designed to perform would plainly be self defeating. Thus the first two conditions just rephrase, and apply to the level of strategic choice, the conditions on dis-trenching search: accountability and search mechanisms in pragmatist organiza-tions must encourage investigation of novel goals, while using mutually illuminating comparisons among the results to limit the search and establish criteria for choos-ing a course of action from its results.

A third and closely related condition is that the pragmatist governance system domesticate, not generate or exasperate power struggles in the organizations to which they are applied. This seems like it should go without saying: a governance structure that invites opportunism is an institutional equivalent of a doctor who sys-tematically violates the Hippocratic oath. But notice that governance in pragmatist organizations is likely to be especially susceptible to such counter-intuitive abuse because of the characteristic circumstances that give rise to opportunism in them. In hierarchies the occasion for opportunism is the agents' informal—unobserved and discretionary—revision of formal rules.

Since the unrevised rules are simply unworkable (working to rule is a famously effective way to paralyze a plant), the informal adjustments that make the organiza-tion function create opportunities for opportunism as well. The governance system with all its incentives and reviews is in some sense (and however flawed) a correc-tive or supplement to the self-limiting distribution of decision-making authority hard-wired into the hierarchy itself.

In the pragmatist organization as stylized here, in contrast, the actors routinely exercise discretion, generating or modifying the organization structure even as they do. We can say either that the emergent or informal network of decision making is the (new-model) formal structure. Or—in view of the way flexible formalization makes explicit much tacit knowledge and discretionary action without reducing it to a traditional rule set—we can say there is neither a formal nor an informal struc-ture in the traditional sense. In this setting governance, in imposing procedures for high-level, strategic choices, creates new occasions for opportunism rather than merely providing a (perhaps faulty) means of limiting existing ones. Consider, for instance, the case where the governance rules ignore the second constraint and insti-tutionalize preferences for particular means or ends. In that case the units and proj-ect groups are tempted to manipulate the center into selecting the rules that show their own efforts in the best possible light. Indeed the ensuring struggles for power, and the attendant disincentives to guess wrong about what will be rewarded, is sure-ly one of the mechanisms by which a misgoverned pragmatist organization chokes off the very searches it is intended to pursue.

The fourth and final constraint is simply transparency: the governance mechanism must enable stakeholders—regulators, the public, employees, investors, suppliers,

creditors—to look into the organization, and into the system of governance itself. The background assumption here is simply that governance arrangements, like all other aspects of the pragmatist organization, are provisional, and designed to be improved in view of the experience they help generate. A governance system whose key purpose is to facilitate strategic choice among competing projects should allow stakeholders to go some way towards corroborating the eventual choices by looking in the relevant detail at the projects under consideration. Or, if the governance system fails to provide such corroboration, it should at least reveal enough about its own operation to allow improvement the next time around.

In posing the problem for pragmatist governance the preceding list of conditions has also telegraphed a general solution: A pragmatist organization can govern itself subject to these constraints by applying the core principle of iterated codesign to the choice of strategy or goals itself. As a first step in a highly stylized version of such a process the center—acting, perhaps on behalf and with the help of representatives of the individual subunits—metaphorically or openly benchmarks its overall objectives, looking for goals “like” the current ones, but arguable better on some dimension. Each of the federated units then does the same with respect to the broad task for which it is responsible; and the general institutional goals are, if necessary, revised in the light of the interim results. At least in theory, in metaphorically benchmarking their goals the overall organization and its subunits make the formulation of their purposes, and what kinds of projects should accordingly be rewarded, as open as, and therefore fully responsive to, pragmatist searches at the project level. Through iterated pooling of the benchmark goals the organization and its units set the general priorities with which to rank projects. So long as the goal-setting process remains open in this way there is little incentive to try and seize control of it: If currently successful projects will routinely be compared with a (limited, but ex ante unspecified) list of potential alternatives, no group can institutionalize selection criteria that are reliably favorable to itself. Hence there is no incentive to struggle for control of the goal benchmarking process (except, of course, to replace it with some method that does allow the powerful to fix purposes their way). Brief accounts of a range of pragmatist governance mechanisms from the private and public sectors indicate how these general principles can be brought to bear in particular contexts. Take first the case of the Illinois Tool Works, a mid-size conglomerate with \$9 billion dollars in annual sales and nearly 50,000 employees world wide. One of its divisions makes plastic and metal components, fasteners and assemblies, industrial fluids and adhesives, fastening tools, and welding products for the construction, automotive and consumer durables markets. Another makes consumer and industrial packaging, as well as product identification and quality assurance equipment for the food retail and service, food and beverage, and capital goods industries. The firm is organized as a federation of some 600 units. These are financed by headquarters,

to which they for the use of corporate research and development facilities. But within these limits the units are largely independent. Each is responsible for setting its own goals, and is held accountable for outcomes. Because of this focus on components and small subsystems, and the flexibility of its individual units, ITW has flourished where mega-module makers, in their rigidity, have floundered. The key governance rule for goal setting in ITW is derived from Pareto's rule of thumb that only a small fraction—about 20 percent—of all the causes of an outcome account for some 80 percent of the total effect. The 80/20 governance rule obligates the units to regularly re-determine *which* 20 percent of their activities account for (roughly) 80 percent of their profits; to build up these disproportionately beneficial activities; and to spin off or simply abandon the rest. But each unit's sales to its customers bundle many different kinds of goods and activities: particular co-design services; unusually short product-development cycles; innovative use of a particular materials or the processes by which they are worked. So, following the rule, the unit's first task is in effect to do a root-cause analysis of the grounds of its own current successes, and re-direct its strategy according to the results. That done, the unit must metaphorically benchmark the new strategy by exploring potential uses of its newly re-characterized capabilities that are "like" the currently successful ones. Because judgments about the viability of strategies discovered in this way always contain a speculative element, a subsidiary governance rule provides that managers are not penalized for being wrong (once or twice), but face immediate sanctions if they are caught pursuing strategies that have not been disclosed and justified to headquarters. This second rule gives managers an incentive to (temporarily) immunize themselves against the risks of incorrect decisions by increasing the transparency of their decision making to headquarters, and so to other units in the firm and outside stakeholders.⁷

A second illustration is the governance system of Cisco Systems, the leading maker of network routers, switches and interface devices. To remain competitive in one of the worlds' technologically and commercially most volatile markets, Cisco invests of course heavily in research (roughly 17 percent of sales in recent years). But the distinctive aspect of its strategy is a policy of acquiring technologies or products pioneered by other, usually much small firms, and then working with the managers of acquired firms to develop them. As of mid 2001, XXX years after its founding, Cisco had incorporated some 75 units through A & D. This policy depends in turn on two governance rules. The first, dis-entrenching rule obligates business units to conduct a "make or buy" review when preparing their annual business plan. This requires each unit to compare the strengths and weaknesses of its current produce or service, and closely related variants of these, to those alternatives under development or already produced by competitors. Because Cisco's headquarters has rich knowledge of the changing needs of end users through its sales force, and many ties to the research community, business unit managers have every reason to identify and

evaluate potential acquisition targets before they come to general attention. In case of an acquisition, a second governance rule provides that the inside managers are rewarded for retaining the managers of the target firm, and integrating them into Cisco. Together the two rules not only encourage (as at ITW) regular reassessment and occasional brusque change of strategy and the business—unit level, but also, by making outsiders into insiders, increase the cognitive diversity of management generally, and so facilitate the next rounds of assessment and change.⁸ And of course business-unit changes can in combination lead to large changes in overall strategy.

The public-sector equivalents of these dis-entrenching governance rules take the general form of an obligation that each unit measure itself against some general goal or performance standard—reliable and safe operations, continuous improvement in service provision—and correct shortfalls revealed by comparison with the performance of others facing similar situations. A straightforward regulatory application is the requirement to undertake near-miss analysis in the US nuclear power-generating industry. Utilities in the industry must report disruptions in their operations to the Institute for Nuclear Power Operation (INPO), an industry-funded entity ultimately responsible to the Nuclear Regulatory Commission. INPO officials sift these reports to distinguish harmless disruptions from dangerous ones. Thorough analyses of the causes of the dangerous disruptions, and ways of preventing them, are then circulated as Significant Operating Experience Reports, or SOERs. Industry Operating Experience Reviews then periodically assess the ability of particular plants to effectively use the SOERs and other means to improve their own affairs. For purposes of this review, a team of specialists in a variety of areas evaluates the plant's troubles since the last INPO inspection, paying particular attention to the plant's own reports on how it has responded to SOERs.

Where minute variations in daily operations are less likely to signal the possibility of substantial, hidden risks, the government can require the regulated entity to scan periodically for possible hazards and present a plan for mitigating those that it identifies. A regulatory oversight body then evaluates the adequacy of the plans, and the steps to realize them, against the benchmarks set by the best performers.

The shift in the US in the 1990s from poke-and-sniff (organo-leptic) methods of ensuring food safety to the hazard analysis of critical control points (HACCPs) shows the drift of developments.⁹ In the organo-leptic method an inspector from the Federal Safety and Inspection Service examines every head of cattle or chicken being dis-assembled in a slaughterhouse for quality defects and especially signs of pathogens. The limitation of the method of course is that some pathogens may not be detectable by the usual examinations, so that meat products leaving the processing plant are not assuredly safe; and even if they are, pathogens introduced at later stages of the food supply chain would notice. Under the HACCP regulations introduced by the US Department of Agriculture, meat and poultry processors have to

identify all the points in their production processes where pathogens are likely to be introduced; detail how they will reduce these risks and verify, by testing, the success of the adequacy of their measures. The HACCP plan must be complemented by a Standard Sanitation Operations Plan detailing the plant's regular housekeeping measures. The role of federal inspectors shifts from direct examination of animal carcasses to verification of the processors' hazard reduction systems. This verification starts with assessment of the adequacy of the HACCP plan (and the companion Standard Sanitation Operations Plan) and includes review of the plants test results as well as independent testing by the inspectors. Eventually the HACCP system is to cover every link in the food supply chain from farm to plate.¹⁰

The US Securities and Exchange Commission is shifting the regime governing financial disclosures by publicly traded companies in an analogous direction. Having re-discovered in the recent stock-market bubble that complex rule systems are easily gamed by managers with powerful financial incentives to do so, the SEC is requiring firms to practice "critical accounting": In reporting their financial results, firms will have to identify the critical accounting issues where their choice among arguably legitimate but rival methods of valuing results made a material difference to their statement of overall performance. For each of these issues the firm will have to document the alternatives considered; the valuation that would have resulted from the application of each; the reasons for the choice of the method actually used, and the key participants in the decision-making process that produced the final result; and even an assessment of the risks to the regulatory system as a whole that would eventuate if the firm's preferred method were broadly adopted.

With regard, finally, to complex public services, the dis-entrenching governance rules typically requires operating units to formulate and periodically revise strategies for increasing rates of improvement towards a general end. In the case of pragmatist public school reform discussed above, for example, the general goal is to reduce and eventually eliminate the difference in performance in key subjects such as reading and mathematics between affluent, white students and poor students, who are also often of color. In Texas the State governance regime accordingly requires period testing in these subjects by means of sophisticated standard test that (now) rewards the ability to conceptualize rather than rote learning. Each school must report the results of these tests disaggregated by economic and ethnic groups, and the State pools the data so that parents with children in a particular school can compare the rate of improvement of improvement of the relevant subgroup in that school to the 39 other schools in Texas demographically most similar to their children's. Further dis-aggregation at the district or school level then yields information about the performance of particular teachers and administrators that can guide further reorganization.

In elaborating these general features of pragmatist governance and illustrating their practicality the aim is to make plausible the claim that the new search networks

are in principle governable, but not to suggest that they are already well governed. On the contrary: Leaving aside cases of egregious financial wrongdoing, it is striking that there is not a single, generally recognized example of sustained, organized, and well characterized good governance in a multinational firm, including of course the many indisputably successful ones. Worse yet, many, if not all, the cases of corporate governance of innovation that have been presented as exemplary soon proved seriously defective.

REALITY CHECK—GRAPH ON CURRENT GOVERNANCE MESS

Final word on the hope for democracy in the revolt against routine. No one should believe that democracy is about to be reborn just because there is pervasive pressure to redistribute “improvisation” rights. But only a congenial pessimist can be unmoved by the possibility that the deepest fundamentals of our despair about democracy and self determination are being quietly shaken by a revolution in organized problem solving.

NOTES

- 1 See below, p.
- 2 John Dewey, *Human Nature and Conduct*, is careful to distinguish habit, as the creative disposition to respond, almost unselfconsciously, to new situations on the basis on long experience, from routine, as habit frozen into a compulsion to repeat prior responses despite novelty in the situations that typically trigger them.
- 3 For examples of how NUMMI (the GM-Toyota joint venture) handles the tension between standard operating procedures and continuous improvement, see Adler et al, 1999.
- 4 Weick, Sutcliffe, and Obstfeld
- 5 Schoen, D. A. “Generative Metaphor: A Perspective on Problem-Setting in Social Policy”, in *Metaphor and Thought*. Ortony, A., ed. New York: Cambridge University Press, 1979. pp. 254-283, pp. 257-259.
- 6 Simon, *The Science of the Artificial*, pp. 90-92.
- 7 See Herrigel’s essays, and his interview with ITW officer. Notice that in casting about for an organizational solution to the innovator’s dilemma of managing “sustaining” or routine lines of business without being blindsided by “disruptive” threats, Christensen seizes on the example of a company—Nyro—whose structure is similar to ITW’s, but whose governance procedures are in crucial ways less formal and effective. What ITW is to metal components and small systems, Nypro is to plastics: a federation of many small units providing specialized products in various industries. Nypro, like ITW provides central research and development services to its largely independent business units. The key difference is that Nypro has no equivalent to ITW’s 80/20 dis-entrenching rule. Instead the firm uses periodic meetings

among unit managers to spread ideas and develop a cooperative culture despite their rivalry for the headquarters' approval. In addition the CEO keeps close touch of innovations that diffuse among several units, and requires that the most successful of these be incorporated into the firm's standard operating procedures. There are, however, no systematic incentives, as there are at ITW, to scrutinize present practices for clues about new markets. This, for Christensen, remains the job of the CEO. (innovator's Solution, 271-275, "institutions can't disrupt themselves" is from p. 274)

8 Mayer, David, and Kenney, Martin F. "Economic Action Does Not Take Place in a Vacuum: Understanding Cisco's Acquisition and Development Strategy," BRIE Working Paper 148 September 16, 2002, esp. p. 24 for the account of the "make or buy" review.

9 See US Department of Agriculture, Food Safety and Inspection Service, 9CFR Part 304, et al. Pathogen Reduction; Hazard Analysis and Critical control Point (HACCP) Systems; Final Rule, Federal Register, Thursday, July 25, 1996, 38805-38850.

10 For the limits to the current implementation of the HACCP, see Petersen, Melody, and Drew, Christopher, "New Safety Rules Fail to Stop Tainted Meat," New York Times, Oct 10, 2003. pg. A.1. But for an illustration of the ways this system produces the information needed to pinpoint deficiencies in its own implementation, see US Department of Agriculture, Office of Inspector General, Great Plains region Audit Report, Food Safety and Inspection Service Oversight of Production Process and Recall at Conagra Plant (Establishment 969), report No. 24601-2-KC, September 2003.

Charles Sabel

MAURICE T. MOORE PROFESSOR OF LAW AND POLITICAL SCIENCE
AT COLUMBIA UNIVERSITY