Journal of Theoretical Politics 2 (1): 59-83 (1990)

AN INSTITUTIONAL APPROACH TO THE THEORY OF POLICY-MAKING: THE ROLE OF *GUIDANCE MECHANISMS* IN POLICY FORMULATION

Stephen H. Linder and B. Guy Peters

ABSTRACT

The premise of this paper is that the revision of the *stages metaphor* for the public policy-making process is best undertaken from an institutional rather than individualist approach to analysis. The theoretical linkage between policy formulation and information – including feedback about performance at other stages as well as the results of social research – is mediated at the level of the policy formulating institution by complex learning and adaptation processes. Together these processes constitute a guidance mechanism for managing error and making course corrections. Different types of mechanisms, however, will admit different kinds of information which, in turn, will condition how the institution adapts to changing circumstances. The analytical features of guidance mechanisms are related to knowledge utilization, normative theories of policy-making and implementation, and to policy design.

KEY WORDS: • policy formulation • organization learning • policy design

Introduction

One of the principal contributions of the discipline of political science to policy studies is the explication of a process model of policy-making (Jones, 1984; Peters, 1986). This model imposes a temporal order on policy events, punctuated by decision points that involve individual actors shaping decisions in pursuit of certain values. The actors and the values at stake are generally exogenous to the model which concentrates instead on characterizing their roles in a sequence of distinct policy-making stages. The metaphor of stages in a policy-making process has been a powerful one, not only framing the conduct of policy analysis (Dunn, 1982; Hogwood and Gunn, 1984) but also defining new subfields of policy inquiry – implementation research, evaluation research, and so on. Aside from the stages metaphor, however, the process model has also altered the way policy research treats institutions.

The model's assumptions about the centrality of individual actors and the strong influence of policy environments serve to downplay the role of policy-making institutions. Autonomy is attributed to individual actors but not to any organizational form; accordingly, individual choices and their outcomes rather than institutional action form the proper focus for inquiry.

The process model, for the most part, treats the institution simply as the arena within which policy interactions occur. The forces that drive the policy process from stage to stage are seen as more likely to emerge from outside the institution than from within it. As a consequence, institutions appear as intervening variables, typically acting as constraints on feasibility, bargaining or other forms of strategic behavior, but seldom play a distinct role in policy-making beyond that of the individual actors participating in the process. In short, the process model seems to reduce policy-making to the level of individual calculation, emphasizing micro-level processes over meso or macro ones. While this may be a reflection of a broader reductionism common to contemporary theories of politics (March and Olsen, 1984), it nonetheless inhibits the process model from accommodating organizational phenomena that are not easily reduced to the behavior of individual actors.

Analytical efforts to revise the process model over the last few years have focused on the refinement of the stages metaphor but have usually accepted its underlying reductionism without complaint. The impetus to revise the model in this way can be traced to a growing recognition of how the transfer of knowledge across stages affects policy-making. Research on the strategic importance of program implementation for the improvement of policy performance (Sabatier, 1987), work on the utilization of technical information (Weiss and Bucalis, 1980), and on the communication of feedback about outcomes (May and Wildavsky, 1978; Cobb and Elder, 1981; Hays, 1985) all point to the inadequacy of a sequential ordering where policy events flow in only one direction. These revisions establish the dynamic character of the process, featuring internal feedback loops and iterative forms of adjustment. Nonetheless, it remains a process shaped by the resources, capabilities and information-processing habits of the individual actors involved.

Our intent is to build on this and other work that strives toward a new conception of the traditional process model, but to do so from a 'neo-institutional' vantage point (March and Olsen, 1984; Zucker, 1987). Consistent with our reform-minded predecessors, we view policy-making as a developmental process involving feedback and learning, and feature the formulation stage as the focal point for information bearing on both past and future policy performance. We part their company, however, once we assume that a theory of policy formulation need not be limited to the action of individual actors engaged in bargaining, conflict and choice. Instead, we emphasize the role of organizational forms capable of purposive action, learning, adaptation and importing their own interpretation of reality. These forms are not merely production or exchange systems serving as the backdrop for the play of individual rationality, they generate and support non-instrumental processes, express symbolic value and constitute more than the sum of the preferences and cognitions of their membership at any one time (Scott, 1987). From this perspective, the organizations at

the heart of the formulation process are best viewed as 'institutions' – just as likely to shape the perceptions, belief systems and values of their members as be shaped by them – continuously engaged in negotiating and altering their environment rather than simply mirroring its influences.

An institutional perspective on policy formulation permits us to reconsider the role of information in the formulation process and to offer a richer conception of the institutional dynamics involved. At the center of our theory of formulation lies the concept of a 'guidance mechanism' that links the learning side of the process involving knowledge and information use with the adaptation side tied to policy implementation and performance. The discussion to follow concentrates on the significance of variation across guidance mechanisms of different types and examines how such variation might account for a range of formulation-based phenomena, including patterns of knowledge utilization, and the incidence of various policy designs. However, contextual explanations for the emergence of one mechanism-type over another and treatment of the complex relationship between a given type and its environmental, structural and behavioral correlates have been omitted due to limitations on space. The next two sections develop an institutional view of policy formulation and establish the basic premises of our argument. Next, the concept of a 'guidance mechanism' is introduced and linked to a rudimentary theory of the formulation process. With these tasks completed, attention turns to the theoretical import of the guidance mechanism as a gatekeeping device selectively mediating information functions and requirements.

Background

Of all the stages in the traditional process model of policy-making, formulation has received the least analytical treatment – although probably more than its share of descriptive study – and is arguably the least understood. Policy-relevant events that precede formulation (Kingdon, 1984) and those subsequent to it (Nakamura and Smallwood, 1980) continue to attract more scholarly attention than does the process of formulation itself. Apart from the analytical models of how individual actors make choices (Cohen, March and Olsen, 1972; Steinbruner, 1974), and several normative ones for improving the choice process (Dror, 1985; Braybrooke, 1985), the formulation stage of policy development remains a 'black-box', presumed to contain a varying mixture of cogitative and interactive components (Wildavsky, 1979).

Under the traditional process model, the best attempts to theorize about what goes on inside the 'black-box' of policy formulation have built on inferences drawn from the outside, focusing on either the 'inputs' or the pattern of 'outputs' for clues about the intervening processes. The empirical modeling studies of formulation (e.g. McFadden, 1975; Davis, Demps-

ter and Wildavsky, 1966, 1974) are representative of the output focus which devotes careful attention to specifying the decision rules underlying the choices of selected actors. The input focus, in contrast, typically emphasizes the resources and capabilities that the participants bring to the formulation arena rather than the choices they might make (Sundquist, 1968; Polsby, 1984).

The input focus builds on the presumption of a strong interactive component in policy formulation, while the output emphasis highlights its cogitative aspects. Neither gives very much weight to the role of institutions as actors 'in their own right'. Aside from the notion of role constraints associated with the bureaucratic politics literature – 'where you stand depends on where you sit' – institutions are generally thought to reflect the attributes and psyches of its members. The process model's apparent neglect of institutions may have reinforced the 'black-box' image of formulation by overlooking organizational phenomena that defied reduction.

An institutional approach then can be seen as way of opening up the 'black box' without having to add on layers of complexity to keep explaining things at the individual level. It also permits us to take advantage of the rich tradition of information-processing theory applied to organizational decision-making (March and Simon, 1958; Schroder, Driver, and Streufer, 1967). This theory supports two powerful claims: (1) organizations develop characteristic ways of gathering and integrating information about their environment (Ungson, Braunstein and Hall, 1981), and (2) the link between these information functions and an organization's behavior may be mediated by symbolic processes that have little to do with the rational behavior of individuals (Feldman and March, 1981). Recognizing such processes challenges the validity of purely instrumental conceptions of information use and goes beyond the assumptions of bounded rationality and limited search (March, 1987). As organizations learn, they develop an interpretive map of their environment and transmit this learning to their members in the form of routines that govern both adaptive and manipulative responses. From an institutional perspective, these routines are more likely to shape the preferences and behaviors of individuals within institutions than be shaped by them (see Wildavsky, 1987).

To the reductionist uncomfortable with institutional explanations, these assertions may appear untenable unless one is willing to assume that organizations are perfectly cohesive entities. As long as the focus remains on individual actors, cohesion would be necessary for any action not directly attributable to individual choice and calculation. From an institutional perspective, however, the whole is more than the sum of its parts; the relative cohesiveness of the actors within the organization is more likely to be a product of organizational learning and adaptation than a necessary condition for them. Moreover, there may be a rich ecology of learning and adaptive processes within an organization that survives turn-

over in its membership and yet maintains diversity within it. Institutional arguments in general are undergoing a revival in the social sciences, as recent reviews in political science (March and Olsen, 1984), economics (North, 1986) and sociology (DiMaggio, 1987) attest. They hold in common a concern with the narrowness of purely instrumental conceptions of collective phenomena that rely on rational and efficiency arguments ultimately grounded in methodological individualism. In contrast, institutional arguments view institutions as autonomous actors and typically feature non-instrumental processes thought to emerge as an organization evolves and begins to take on a logic of its own.

A second advantage of an institutional emphasis is that it readily accommodates recent efforts to build a more dynamic model of the policy process. Formulation, then, can be treated as a real-time process rather than as a series of discrete events, each with a beginning and end. The most interesting aspect of formulation from this perspective is not 'who wins and who loses' in the battle over control of allocative choices but rather how the process, as it unfolds, seems to manage its environment and respond to its own missteps and mistakes. To the extent that formulation involves some institutionalized form of error control, the interconnection of events traditionally separated by the temporal ordering of stages will assume added importance. What happens in the implementation stage, in other words, will begin to matter a great deal for subsequent rounds of policy formulation.

Under the traditional model, when existing policies are revised or new ones developed, the knowledge involved is attributed either to individual judgment or to analytical methods of the sort taught in professional schools of policy and administration. The question is framed in terms of the kind of information that individuals seem to prefer. Some contend that information from analytical methods plays a subordinate role and may be inferior to more intuitive judgments (Lindblom and Cohen, 1979). Countering this view, others emphasize the cognitive rather than instrumental value of information drawn from those analytical methods (Weiss, 1977). From an institutional perspective, however, knowledge in organizations serves a variety of strategic and symbolic functions that are typically insensitive to the predilections of individual users. Moreover, the best way to understand these functions is to ignore the necessity of any presumed connection between information and individual decision-making.

The idea that decision-making should be tightly coupled to information often serves as a benchmark for assessing the utilization of research results in policy-making (e.g. Rich, 1981). Tight coupling, however, sets too narrow a standard for judging instrumental use – although it does support a more inclusive definition of utilization failure and a greater need for remedies. While our emphasis will be on the instrumental function of information, we assume no such standard. More information does not (necessarily) produce better decisions, nor is information-gathering a sign

of unmitigated virtue. It is not uncommon, or at times even undesirable, that decisions in the public sector not accord strictly to the balance of information presented, and that more information not result in better choices (Feldman and March, 1981). Instead, we expect information – decision linkages of different types and degrees, some involving more information and others less, but nonetheless tied more to organizational phenomena than to the predispositions of individual actors.

At the level of the institution, we argue, the linkage between knowledge and policy formulation is best understood as mediated by two tandem processes – learning and adaptation – residing inside the 'black-box' of formulation. Together these basic processes constitute a mechanism for guiding the institution along its chosen path of development. Since the primary function of this mechanism is error management and course correction, it is referred to as a 'guidance' mechanism, consistent with the systems concept of self-correction (Dunsire, 1986). We interpose the guidance mechanism between the input of performance information and any compensating response emerging from the formulation process. Different types of mechanisms will admit different kinds of information which, in turn, will condition the range of plausible responses. Given the selectivity implicit in the guidance function, the type of mechanism in place is likely to leave an unmistakable imprint on the character of formulation process.

Normative theories of policy-making subsumed by the process model of policy-making invariably entail claims about the functions that we attribute to guidance mechanisms of different types. Identifying the guidance requirements posed by different theories then might serve as an alternative basis for meta-analysis which bypasses the contentious issue of how good the individual can (or should) be at making decisions. Emphasis would shift to the plausibility, and perhaps the relative desirability, of different patterns of institutional learning and adaptation. Shifting the focus of theorizing from 'the decision' to the institutional guidance mechanism not only affords a better vantage point for speculation on how policy change actually occurs, but also enables us to ask whether the policy process has the capacity to improve itself and its performance over time. The related question then becomes, 'why is one guidance mechanism in place and not another?' Rather than automatically imputing functional advantages to observed decision processes, explanation involves accounting for anomalies. For example, 'why, in some situations, are the same mistakes being made over and over.', or, more generally, 'how might we best account for mechanisms that by most lights appear dysfunctional?'

Before elaborating on the conceptual form and workings of guidance mechanisms, the issue of environmental influence needs to be addressed. From neo-institutionalists in political science (March and Olsen, 1984; Krasner, 1984) and theorists of organizational development (Child and Kieser, 1981), we find assurances that institutions seldom respond passively to their environment. Accordingly, efforts at adaptation are more

likely to involve managerial strategies designed to reduce uncertainty and dependence than simple compliance with externally imposed requirements. In other words, assuming that what happens outside the 'black-box' of formulation matters a great deal for what goes on inside does not lead inexorably to the conclusion that institutions in general are a product of their policy environment and have no claim to autonomy.

Since we are interested primarily in how guidance mechanisms function rather than in the conditions that lead to the emergence of one type over another, we will take a single specific context as given, effectively holding environmental variables constant, and examine how a range of distinct processes representing at least four types of guidance mechanisms might operate in these circumstances. Although, empirically, different mechanisms may emerge in different contexts, this should not be taken as a sign that there is one proper form to be found for each context. Unless we are willing to assume what March and Olsen (1984) call 'historical efficiency' then goodness-of-fit should serve neither as a design criterion nor as proof of environmental selection. Claims about the functional requisites of certain environments or, in the case of policy formulation, the goodness-of-fit between existing processes and selected policy domains, inhibit the consideration of institutions that, to many, appear anomalous or maladaptive but are actively engaged in reducing dependence and extending their autonomy.

The Policy Environment

Policy-making in all but a few relatively uncharted areas will encounter a crowded field of earlier programmatic efforts devoted to addressing overlapping conditions, needs and groups. Policy crowding is, at least in part, a condition common to liberal democratic regimes whose rule is grounded in demand satisfaction and a cumulative expansion of public services (Heclo, 1978; Lowi, 1979); it is also testimony to the endurance of client-based programs and their managers' instincts for survival (Hogwood and Peters, 1983). Crowding sets up complex and typically unforeseen interactions among programs, at times, offsetting any intended actions with unintended and perhaps retrogressive ones and at other times, leaving no trace of program impact whatsoever. The untoward effects of this crowding may be manifested as 'noise' in performance information, as opportunism in the behavior of clients, or simply as high risks of error in a given policy domain.

A second contextual feature compounding the effects of crowding in many policy domains is the 'turbulence' of the policy environment (Emery and Trist, 1965). A highly turbulent environment is not only unpredictable but volatile and subject to change in erratic ways. The demand and supply conditions behind the use of any given policy instrument are likely to be

unstable, and the balance among values that define problems to be addressed, easily upset by political maneuvering and the press of new events. In this context, policy-making becomes primarily an effort to accommodate changing and unpredictable circumstances in a way that, among other things, minimizes error. Here, the challenge from the formulation side might focus on charting the risks to policy performance, perhaps identifying sources of potential opposition and building-in modes of accommodation. Policy formulation can be viewed, in part, as an attempt to impose greater structure and predictability on the policy environment. This may help explain the apparent preference of many policy-makers for more intrusive solutions, e.g. direct provision or regulation, as opposed to more indirect but perhaps also more efficient means.

Operating in a crowded field, then, places a higher premium on performance information, if for no other reason than as a hedge against prospective failure. When crowding is complicated by turbulence, however, knowledge and information should play an even larger role, as should the mechanism used for detecting when things are going badly. This mechanism may be primitive and rudimentary or refined and elaborate; it may be based in group processes, rational calculation or organizational routines. In any event, we can expect to find evidence of some facility for error recognition and correction, as well as some way of taking the prospects of failure into account, whenever policy-making addresses problems in crowded and turbulent settings. In these circumstances, the mechanism for appraising information on performance becomes more salient and thus more easily identified. With greater salience, the distinctiveness of one mechanism over another becomes more pronounced as well. In effect, the adversity of the context serves as an experimental condition intended to bring out the best (and the worst) in guidance mechanisms of different types.

Once the connection between formulation and guidance assumptions is acknowledged, several other relationships move into focus. How important an influence on formulation is retrospective knowledge coming out of either prior experience or research on past performance? Is there any room for learning? How is 'success' to be distinguished from 'failure' and what bearing should this have on inferences about a policy's future prospects? Not all retrospective knowledge, for instance, information on the apparent correlates of implementation success, will be admitted into formulation consideration by all guidance mechanisms. Some mechanisms will favor prospective information, while others will respond only to direct experience. The intent here is not to advocate one mechanism over another but to develop a more inclusive view of the dynamic connection between policy experience and formulation.

The Guidance Mechanism

The warrant for a guidance mechanism is based on the claim that, for most policy-making institutions, the relationship between information and policy formulation is not a strong and direct one but tends rather to be weak and proximate. In other words, even when information is available in abundance, it may not help in reaching decisions. The corollary is that it need not help. Information-gathering need not be directed exclusively at decision-making but may serve other purposes just as vital to the organization. Relatedly, not all information-processing focuses on decisions, and not all decision-making is grounded in the available information. This leads to the paradox described by the organization theorists, Feldman and March:

> Most organizations and individuals often collect more information than they use or can reasonably expect to use in the making of decisions. At the same time, they appear to be constantly needing or requesting more information, or complaining about inadequacies in information (Feldman and March, 1981: 174)

If we assume a strong link between information and formulation decisions then the only way to explain this paradox is by claiming that either most information or most decision-makers are seriously deficient. Presupposing a weak link, on the other hand, admits both a range of alternative functions served by information and an information-processing capability on the part of most organizations that may or may not show up in any particular decision. It becomes impossible then to infer a pattern of information use from the outcomes of particular decisions. An incremental change, for example, can just as likely be the product of comprehensive information, exhaustively analyzed, as of information restricted to rough comparisons at the margin.

In other words, when information enters the 'black-box' of formulation, it need not elicit any compensatory response; conversely, responses are not easily traceable back to a particular set of information inputs. Note that these claims are inconsistent with the assumption of a two-step, stimulus-response linkage of information and decision found at the core of rational calculation and its variants, but equally at home with the cybernetic view of formulation-as-servomechanism (Steinbruner, 1974). In its place, we propose a four-step linkage that may or may not be completed for any given input. The two extra steps that intervene between input and response are taken up with the workings of the guidance mechanism.

A simple diagram of the four-step linkage appears in Figure 1. The two-step linkage that bypasses the guidance mechanism is shown by the dashed lines labeled (a) and (b); this is the traditional, strong-linkage view with the individual decision-maker as the focal point of the process. The four-step linkage is indicated by the solid lines, numbered 1 to 4. The first two

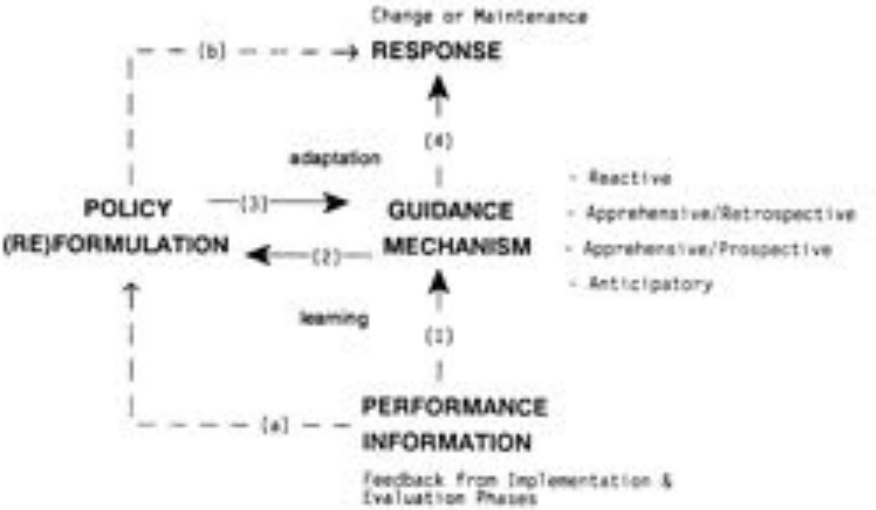

Figure 1. The Linkage Function of the Guidance Mechanism

steps constitute the learning component of the linkage and the remaining two, the adaptation component; when the linkage is complete, the guidance mechanism is performing both a learning and an adaptation function. While these two functions are separate, we can expect some forms of learning and adaptation to appear in combination more often than others. These pairings provide a coherent basis on which to identify and classify distinct types of guidance mechanisms. Our discussion is limited to four such mechanisms, although a larger number of pairs are introduced.

In place of a four-step linkage, which represents learning and adaptation as distinct events, the path from (a) through (b) collapses them into a single process based on the presumption of a strong linkage between performance information and action. Here, adaptation becomes the mirror image of learning and vice versa; both represent a common response to a given external stimulus. Treating them as separate events, in contrast, admits that adaptation may not constitute a stimulus response, nor need learning necessarily be reflected in some form of adaptation. When they do occur in tandem, moreover, they need not appear in 'compatible' forms; complex forms of learning may precipitate very simple forms of adaptation, and so on. Thus, once the guidance mechanism is admitted, formulation can escape the stimulus–response straitjacket that denies institutions any room for purposive maneuvering beyond that which is provoked by their policy environments.

The Learning Component

At the center of the efforts to revise the traditional process model of policy-making is the notion that actors throughout the policy process manifest the ability to learn from experience. The dynamic character of the process comes from the fact that the changes attributable to this learning are ongoing and not confined by any sequential ordering of policy-making events. Yet, the learning that does take place is largely confined to individuals; organizations are thought to 'learn' only in the passive sense of being a repository for the learning done by its members. Once the organization is viewed as an institution, however, the character of the learning that takes place changes.

By the time an organization has persevered long enough to assume the status of an institution, its norms and traditions are conveyed through custom and ideology (Child and Kieser, 1981); its long-term memory is reflected in symbols, myths and a conceptual framework for interpreting experience (Feldman and March, 1981); and it encodes inferences from its history into routines that are independent of the actors who execute them (Levitt and March, 1988). From this perspective, an institution is not limited to setting the context of its members' learning, it defines what will count as error and determines the appropriate forms of detection and correction.

Not all institutions learn in the same way. Some narrow the scope of permissible error to include only disruptions in the existing pattern of activity, while others admit flaws that challenge their basic norms and principles. In the former instance, there is little need to call on long-term memory of earlier patterns of activity since the existing pattern serves as the benchmark for assessing error. Further, the negative feedback conveyed by the disruption contains all of the necessary diagnostic information for error-avoiding adjustments. This is conditioned rather than contingent learning. No elaborate understanding of cause and effect is necessary, nor is there any need for anticipating where the adjustments will eventually lead.

In the latter case, however, memory plays an important role; the steps from accumulated experience to the principles that guide judgments about performance need to be retraced. Things learned earlier may need to be 'unlearned' as a part of developing a new set of norms and priorities. Argyris and Schon (1978) call this 'double-loop' learning, since it involves a kind of feedback on the feedback itself. Hedberg (1981), along the same lines, refers to this as 'turnaround' learning, as it extends beyond simple turnover in causal parameters or response repertoires to the institution's conception of what constitutes effective performance.

The variety of possible forms of learning can be summarized once we parcel the learning process into an acquisition phase – the trial – and an assessment phase, wherein results are compared to some benchmark for

performance. The trials may be simple, engaging an immediate, separate reaction for each piece of information; or they may be compound, where the feedback of multiple trials is allowed to accumulate over time before any reaction occurs. The assessment phase may involve comparisons between the immediate past and current performance – widely known as trial and error – or between current performance and expectations. Learning based on trial and error is accomplished through simple enumeration, that is, the direction of the process is not set ahead of time. Learning based on trial and expectation, in contrast, takes place within a bounded framework that anticipates the locus of corrections.

Cross-tabulating these two phases, acquisition through simple or multiple trials and error-based or expectation-based assessment, yields four distinct forms of learning. These four can be arranged into a hierarchy of increasing demands on institutional memory, on interpretive capability and a broadening scope of admissible information. 'Simple trial and error' is highly focused on a single channel of information, it involves the iterative application of previously conditioned responses. At the next level, 'multiple trials and cumulative error' adds accumulated information and the retention of experience as an aid to assessment; error is placed in a retrospective context of earlier efforts at correction. For our third level, 'trial and expectation', there is not only an accumulation of experience but the addition of a refined sense of contingencies, perhaps in the form of causal schema, that permits the institution to move away from error avoidance to the pursuit of more strategic objectives. Finally, there is 'multiple trials and cumulative expectations' which permits an assessment of the institution's objectives based on both retrospective and prospective information. This is not to say that information can somehow be processed in a way that approaches the synoptic ideal, but only that learning may be both forward- and backward-looking. No claims of comprehensiveness need to be made.

Each form of learning can be viewed as an expansion on the forms that precede it in the hierarchy. The form of learning characteristic of an institution, however, is a necessary but not sufficient condition for adaptation. The fact that a given institution has the capacity to 'unlearn' its norms or to adjust its expectations will not dictate the form that its external responses will take. In effect, the form of learning sets the opportunity for responses of different kinds and scope; the broader the scope of learning, the wider the range of possible adaptation. We can now turn to the adaptation function of the guidance mechanism and consider how the two functions, learning and adaptation, work in tandem.

The Adaptation Component

In a discussion of adapting the provision of government services to changing circumstances, Hood (1986) identifies four possible responses that

formulators might make. Assuming of course, that changes can be detected and a capacity for adaptation exists, the first type of response is to engage in piecemeal adjustment of the policy currently in place. This might involve varying amounts of resource inputs or clarifying certain functional details. In any event, adjustments are reflexive and marginal in the classic sense of disjointed incrementalism, and protect the option to reverse course should problems develop.

A second response also operates on the existing policy but does so by changing its mix of components, introducing new combinations of established elements to change performance without altering the basic design. Policy succession in the sense of Hogwood and Peters (1983) has this complexion. A third response simply involves borrowing proven designs from other institutions or time periods, assuming roughly comparable trials and basic similarities in assessments. Rather than adding new flesh to the bones of an earlier initiative, the aim is to replace these bones with new ones by transplanting from policy that seems to be performing well in an analogous situation. This pattern of imitation and diffusion is likely to be affected by the relationship among institutions in a given context, and the mix of competitiveness and cooperation. The final response is to fashion a new, presumably better, formulation to replace the old. This would be similar in form to the cases of programmatic innovation discussed by Polsby (1984).

These four follow a natural order in terms of their relative requirements for, what Hood calls, 'engineering capacity' and 'down time', the interruption of performance associated with formulation activity. If we translate these two terms into information and comprehension requirements, we have the hierarchy of analytical strategies proposed by Braybrooke and Lindblom (1963). The basic insight that these four approaches differ first and foremost in their information demands has weathered the 20-year span between the Hood and Braybrooke works. The matter still to be settled, however, is at what point an attempt to meet these demands becomes worthwhile. At what point should the institution move from piecemeal revisions to consider changes in basic policy design?

For some, the answer lies in the character of the environment (Landau, 1973; Cobb and Elder, 1981). Some environments demand innovativeness as the price for survival, others will require or permit only small, cautious maneuvers. If we are willing to deny autonomy or purposive action on the part of institutions, or focus on the constraints affecting individual actors, then this makes sense. For others, the form of adaptation likely to occur is closely linked to the structure of institutions (Lindblom, 1977; Wildavsky, 1979). A decentralized, variegated and diverse institution like the market will respond in a piecemeal fashion while a centralized, hierarchical, homogeneous, institution, say, a bureaucratic apparatus, will engage in less tentative strategies.

The point is not that markets cannot foster innovation, surely they

can and do, or that the typical bureaucratic apparatus normally eschews incremental decision-making, but rather that the structural form of an institution places binding constraints on the adaptive behavior of individual actors. When combined with the environmental determinism of the view immediately preceding, this position relegates the institution to an intervening role while, at the same time, offering prescriptions for structural forms deemed more compatible to the environments and actors of interest. Decentralized structures foster adaptation, centralized ones inhibit it. The idea that institutions might possess guidance mechanisms with varying levels of learning and adaptation and, further, that they might employ them deliberately is foreign to this way of thinking.

Basic Types of Guidance Mechanisms

In an intriguing analysis of how society manages risk, Wildavsky (1988) advances two basic strategies that characterize current policy responses: a strategy of resilience, involving a decentralized, rapidly moving trial and error process; and one of anticipation, where a slow-moving central authority preempts trials in order to avoid error altogether. Not surprisingly, he argues persuasively for the former as the more desirable way to adapt to uncertainties about safety. The interesting aspect of this work for our purposes is that his conclusion is built on several claims that seem incompatible with our guidance-mechanism notions or, for that matter, with the idea of institutional autonomy.

For Wildavsky, resilience involves learning and anticipation does not; accordingly, resilience promotes adaptation to environmental change and anticipation does not. We agree with the claim that learning is generally superior to non-learning – that is one reason we feature learning as a key guidance function. We disagree, however, that learning occurs in one predominant mode (the trial and error form of conditioned learning) and that learning and adaptation are synonymous. Learning is not only more ubiquitous as an organizational phenomenon but occurs in greater variety; non-learning from this perspective is a straw man – without some form of learning going on, an organization, by most accounts, is doomed. The narrow definition of learning as trial and error, however, is consistent with the twin assumptions that learning must be inferred from adaptation and that piecemeal forms of adaptation are more efficient. A guidance mechanism presupposes just the opposite. To the extent that learning is necessary but not sufficient for adaptation, and that environmental selection is a weak rather than determinant force, the guidance-mechanism notion offers a richer conception of the institutional dynamics underlying policy formulation.

If we combine the four levels of learning with the four levels of adaptation, there are 16 possible guidance mechanisms. Since learning is taken as a necessary condition for adaptation, however, only 10 of these are

plausible; the lowest level of learning will typically permit only the lowest level of adaptation, while the next higher level of learning will permit both the lowest and the next higher level of adaptation, and so on. The variety of mechanisms can be constrained further by assuming that the four mechanisms, lying along the diagonal of the 4×4 matrix of learning and adaptation shown as Figure 2, subsume the off-diagonal combinations in their operation. Again, this is consistent with learning as a necessary condition. As illustrated in Figure 2, a reactive mechanism pairs simple trial and error with piecemeal adjustment. At the other extreme, an anticipatory mechanism pairs multiple trials and cumulative expectation with policy innovation.

Two other mechanisms occupy the intermediate positions; both are labeled 'apprehensive' to suggest a posture that is neither purely reactive nor entirely anticipatory. One is retrospective, in the sense of relying on historical experience, while the other is prospective and can accommodate expectations tied to some latent assumptions about causal relationships. The apprehensive-retrospective mechanism can generate either policy succession or piecemeal adjustment as responses. The apprehensive-prospective mechanism, given its capacity to interpret contingencies, extends these forms of adaptation to include imitation and the use of analogous designs instead of historically sanctioned ones. Given space limitations, the narrative remaining concentrates more on the reactive and anticipatory mechanisms, than on either of these two less extreme ones.

A simple analogy to reactive and anticipatory guidance can be found in

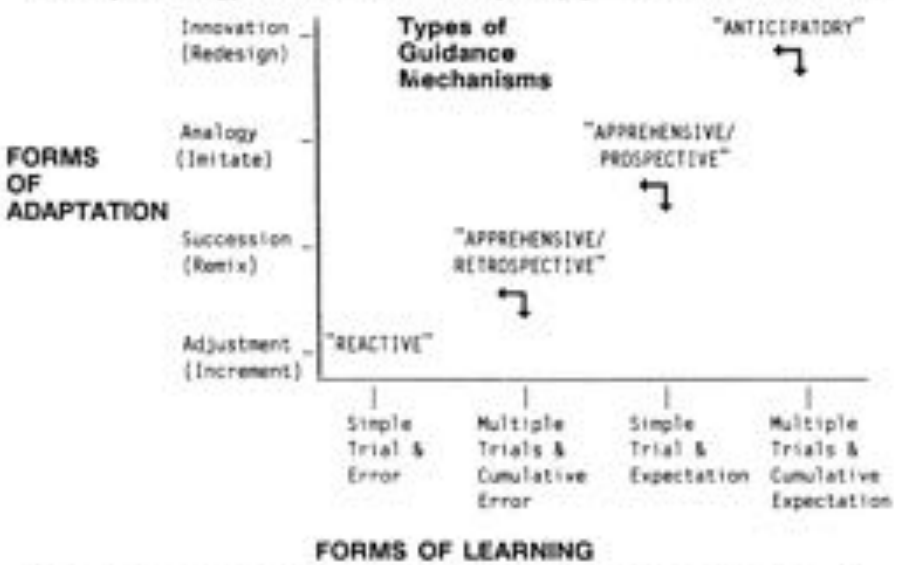

Figure 2. Basic Types of Guidance Mechanisms Arrayed By Their Distinctive Forms of Adaptation and Learning

preventive maintenance and error-based quality control; both depend on models of failure, but the former anticipates errors by replacing components before they fail while the latter waits until a threshold defining non-random failures is crossed before taking any action. An anticipatory mechanism relies heavily on prior information to identify exposure to risk and then acts to interrupt the failure-generating process before failure can occur again; in health care, this mode takes the form of primary prevention, disrupting the path of disease vectors before infection recurs. A reactive mechanism, on the other hand, depends upon detection and the setting of an appropriate threshold for action; this is secondary prevention which focuses on keeping conditions from getting worse once performance has been perceived as impaired. A reactive mechanism responds not so much to some structured memory of the past or to patterns of error unfolding over time – as does the anticipatory mode – but to each error individually, as it occurs in the immediate present.

In a reactive mode, detection must be sufficiently precise to separate one's own errors from the errors of others; further, it may be 'passive' in the sense of waiting for performance messages to be brought to the formulator, or 'active', by involving inspection and aggressive monitoring. Passive error detection is perhaps the more familiar, including examples such as Hirschman's 'exit' and 'voice' mechanisms (1970) and pluralist notions of the 'squeaky wheel' serving as an economical, although biased, device for signaling problems. One convention expressing the stakes in setting the appropriate threshold for action is the tradeoff between Type I and Type II response errors. In setting the threshold for defining error, one must balance the expected consequences of missing a failure against those of acting on false information when there was no real problem.

In effect, the prospects of too little action with potentially disastrous consequences are weighed against those of too much action and the steady dissipation of resources. As the cost of 'false alarms' – based on the misperception of failure – drops relative to the benefits of undertaking remedial action even on slim evidence, the reactive mode may begin to resemble the anticipatory one. As this happens, the emphasis will appear to shift from minimizing the total costs associated with response errors to minimizing the risks of service disruption and other untoward policy outcomes. Nonetheless, the reactive mode will continue to base its corrective responses exclusively on *ex post* information despite any greater willingness to absorb the costs of false positives in its detection decisions.

As a mechanism for managing error, the reactive mode is likely to steer policy formulation toward responses that feature the capacity to make ongoing adjustments based on a continuous but conditional effort to detect mistakes. Keeping the scale of the corrective responses small is presumed to keep the scale of prospective mistakes small and thus potentially subject to remedy. Nonetheless, having to wait for errors of a noticeable size to occur means that the timing of responses will always lag behind the con-

ditions that precipitated the error in the first place. Further, responses must follow an awkward path of over- and under-correction until they eventually converge on a level of performance below the threshold-defining error. Ironically, as convergence proceeds, less and less information will be available about performance; since the mechanism responds primarily to error, smaller errors mean less information (Weinberg and Weinberg, 1979). As a guidance mechanism, then, the responsive mode has a built-in tolerance not only for detection error but for performance errors that accumulate around the margins of its corrective responses.

The amount of information available for piecemeal, corrective adjustments will depend almost entirely on the guidance mechanism's tolerance for error. A high threshold for error will transmit little information except in the worst of circumstances. Of course, by then it may be too late. Underestimating mistakes is likely to inhibit revisions of any sort; overestimating mistakes, on the other hand, is likely to have a chilling effect on change, with any one potentially leading to additional mistakes. In either event, formulation effectively takes on a conservative cast and becomes bound to the reversibility of any corrective response should error be detected.

Guidance for Policy Formulation

For a better understanding of the mediating role of guidance mechanisms, there are two ways to proceed. One can proceed inductively by examining the information requirements behind a particular policy instrument (or program alternative) and then consider the kinds of guidance mechanisms that might accommodate them; inferences can then be drawn to the normative theory of policy-making most compatible with that mechanism. Alternatively, one can begin with a given theory of policy-making, deduce its guidance assumptions, tie them to a particular mechanism, consider the type of information it is likely to favor and, finally, assess the instruments compatible with information of this type. Both inductive and deductive methods should converge on the same set of conclusions about the likely match-ups between theories of formulation and instrument choices.

For dealing with the question of knowledge transfer, we zero-in on only one segment in the chain of inference mentioned above. Attention focuses on the information functions of different guidance mechanisms – both learning and adaptation – and their treatment of knowledge of different kinds coming from identifiable sources. The question of whether formulation should be guided by feedback from other stages, by results from policy research, or by commitments to certain operating principles, can then be addressed with greater awareness of the normative as well as the analytical implications of possible answers.

As a simple illustration of the inductive approach, consider the case of

a tax structure that both satisfies political constraints on its variability and provides an acceptable balance between equity and efficiency considerations. The performance of such a structure as a policy instrument, however, depends directly upon what government can observe: if the attributes targeted for taxation are only indirectly observable, there will be a tradeoff between the costs of obtaining information and the distortions introduced by relying on indirect measures (Atkinson and Stiglitz, 1980). This tradeoff lies at the heart of the concept of a guidance mechanism that steers policy improvement. If obtaining the necessary information is deemed too costly or the risks of failure perceived as too great, then the class of these instruments is likely to be ruled out, a priori. Notice, as well, that the information demands effectively translate into a judgment about what administrative burdens might feasibly be imposed on the organizations charged with implementation. While there may be an analytical, and perhaps normative, link between the design of instruments and their implementation, it is one mediated by guidance considerations.

For the deductive approach, we can begin with the normative theory of incremental policy-making. Borrowing from Goodin's summary of normative incrementalism (1982), there are two claims bearing directly on the issue of guidance and error management. First, the real effects of policy interventions cannot be anticipated prior to actually experiencing them. And second, even if the effects could be anticipated, our evaluative response to them could not be, prior to our actually experiencing them. This not only implies that policy instruments can only be refined and developed as errors occur, but that the goals these instruments serve are themselves likely to change in response to the same errors. Reactive guidance based in trial and error then extends beyond policies to our schedule of preferences about policies. In effect, the threshold for defining error that forms an integral part of reactive mechanisms becomes an endogenous random variable: that is, as errors appear, the threshold defining them may be altered in some unforeseen but adaptive way.

One normative function of this adjustment process is to reduce opposition (see Braybrooke, 1985). Internal *ex ante* opposition can be minimized through the tentative character of the ends being served by formulation, while external *ex post* opposition can be accommodated through either redesign or compensation of some kind. From an anticipatory perspective, the key question is not whether opposition can be accommodated by adjusting means and ends but whether we can anticipate it fading once a policy is firmly in place. A recent discussion of desegregation policies by Hochschild (1984) points up the contrast between reactive and anticipatory modes of guidance. To the extent these policies can be judged a success, they have not only diverged from a piecemeal approach to formulation but have relied on anticipatory guidance, assuming that opposition would diminish over time so long as the bulk of necessary changes were introduced all at once.

Along similar lines, Goodin invokes the notion of retrospective rationality drawn from the literature on paternalism (1982: 41). Here, the anticipatory test is whether one can assume that people's preferences will change once the policy has had its intended effect, turning initial opposition into eventual support. Time perspectives clearly differ between the anticipatory and reactive modes and, consequently, so do the kinds of errors they are likely to find important.

Finally, we turn to the issue of knowledge transfer and pay special attention to the feedback and policy research centered on a single stage in the policy process – implementation. The implementation stage of the process model, in which the pronouncements of government are put into effect, has taken on increasing importance in understanding policy formulation (Hjern and Porter, 1981; Sabatier, 1987). It has been argued that a concern for successful implementation should guide all formulation activity. While heavy reliance on strategic information about implementation may appear a practical approach to the difficult problem of making more effective policy, it can also make formulation more conservative and potentially less responsive to other kinds of inputs (Linder and Peters, 1987). The institution's guidance mechanism intervenes between information drawn from the later stages of policy-making and any effect it might have on subsequent policy action.

After a thorough review of the policy research on multi-actor policy implementation, O'Toole (1986) concludes that, perhaps in contrast to direct feedback there are few well-developed recommendations for policy formulation to be found in this work, and those that do appear are invariably contradictory. To some extent, researchers may be reluctant to fashion recommendations because of either ambiguity over the policy significance of their observations or the perception of barriers to utilization. Those who do make recommendations may find their efforts undercut by critics within their own research community. It is difficult to inspire much confidence when there is no consensus on so basic an issue as what constitutes implementation success and failure.

Recommendations may reflect certain basic disagreements over effective strategy. For example, should policy stress management control or flexibility, or should the focus be on changing policy instruments or tightening their administration? This is reminiscent of Simon's (1947) discussion of the 'proverbs of administration', and their contradictory nature. Alternatively, disagreements may be grounded in different research perspectives. By now, it is clear that observations on implementation looking from the top-down are markedly different than those made from the bottom upward. Still, the basic linkage between implementation and formulation is likely to assume a similar form in either case.

For our purposes, the question is not so much how many recommendations are being made and put to use, or even how good they are, but how the linkage between implementation and formulation takes place –

when and if it does. The issue of whether implementation studies have anything worthwhile to offer is put aside in favor of asking how the formulation process goes about consuming its more palatable offerings. The search for a plausible intervening mechanism distinguishes our approach from that common to the knowledge utilization literature. The functions of analytical information and patterns of substantive use by policy-makers featured in the utilization literature are generally related to the issue of impact. In contrast, our concern is with structure and process; what kinds of systemic mechanisms are operating behind the scenes, structuring and conditioning that impact?

Appraising the quality of the findings emerging from implementation research becomes less important once information no longer serves as the independent variable whose impact is of primary interest. Instead, these findings become one source of performance data to be processed by a guidance mechanism whose character will affect any subsequent policy response. Information gathered from internal monitoring and detection devices serves as an alternative source, likely to be mixed in some fashion with data from other sources. Any given guidance mechanism is likely to admit only a select menu of sources, however, and to process the resulting mix of information in a distinctive way. As a result, we can expect the same set of data on the system's performance of a given policy task to elicit a wide variety of remedial responses, if any are elicited, depending on the particular mechanism in place.

Whether implementation information can serve as a valuable guide for policy formulation depends, as we have argued, on which guidance mechanism is in place. Findings from the 'top-down' approach to implementation tend to be cast in prospective terms, focusing on the features of a policy that might be changed to make it less likely to fail. Of the two guidance mechanisms discussed above, only the anticipatory would respond at all to information of this sort; it is largely irrelevant to a reactive mode of error control. On the other hand, findings from the 'bottom-up' perspective often emphasize the capacity to adjust means and ends to accommodate and thereby reduce error on a serial basis. The view is not prospective so much as retrospective. Here, the aid to formulation might come not as design prescriptions but in the form of an expanded sensory capacity to augment the learning capacity of reactive guidance mechanisms.

Conclusion

In the abstract, there is a rough parallel between recent efforts to develop a theory of policy analysis and the institutional theory of policy formulation that we propose here. Both adopt a meta-analytic perspective on the links between normative theories and analytical structures; in effect,

the choice of a normative theory sets limits on the kinds of analytical structures that are plausible and, in some instances, permissible. Both are moderately relativistic in holding that no one analytical structure is inherently superior and yet each delimits the kinds of policy initiatives to be taken seriously. For Bobrow and Dryzek (1987), for example, the relevant normative theory is a theory of knowledge that leads to the adoption of one analytical approach to policy analysis over another. The analytical approach adopted, in turn, will largely dictate the kinds of policy initiatives under consideration. The meta-analytical problem is to choose the appropriate approach in light of its ties to a distinct normative theory and a given problem context. For Anderson (1987) the analytical structures in question correspond to modes of analysis and the normative theory to particular philosophical traditions. Here, the meta-analytical problem is to assure balanced representation among the most appropriate modes given the requirements of ideal deliberation.

The meta-analytical features of guidance mechanisms as analytical structures are similar. First, their form varies across different normative theories of policy-making and thus offers reliable clues to a given theory's biases toward information and feedback. Second, the mechanism in place will serve a gatekeeper role that selectively favors some policy instruments and initiatives over others. From an empirical perspective, this can tell us much about information requirements and tradeoffs in evaluating prospective policy instruments, about institutional strategies for error detection and control, and finally about the impact of normative theories of policy-making on the workings of the formulation process.

Under the adverse environmental conditions associated with interdependence and unpredictability, the selection of a guidance mechanism will set the overall contours of the formulation process and place requirements on the performance of instruments that may be independent of objectives and other *ex ante* constraints. Design under adversity, then, is very much a product of the guidance mechanism implicit in one's theory of policy formulation. Rather than assuming that the design of an instrument will influence the course of implementation or that one's views on implementation should shape the design of instruments, basic assumptions about guidance that underlie policy-making effectively delimit both formulation and implementation.

In particular, we are concerned with the effects of formulating policies as purely reactive to events in a complex and turbulent environment, as opposed to policies that attempt to anticipate events or involve sufficiently flexible designs so that changes in the environment do not prevent the attainment of accepted policy goals. While trial and error learning through incremental adjustments has long been accepted by most American analysts as the most 'rational' manner of making policy, such learning and information-processing may not be adequate for more complex environments. The conventional wisdom dictates just the opposite; as environ-

ments grow more complex, policy formulation should become more conservative. In the absence of a well-articulated framework for policy, however, trial and error may be no guide at all. If we simply muddle through without a strong logical and theoretical basis for a given policy design, we will have no means of understanding the source of many errors that may occur (Goodin, 1982).

Similarly, other experience-based systems of information generation, such as those tied to various diagnoses of implementation failure, may suffer from many of the same logical weaknesses. From this perspective, the question of whether implementation research can, or perhaps should, serve as a principal source of guidance for policy formulation takes on a more general, analytical meaning. The guidance provided by implementation research is best judged in the context of other sources of guidance, that is, as one of a variety of distinct information sources whose influence on formulation will be mediated by the way formulators have implicitly chosen to manage and control their errors. Approaches to error correction will be bound up with the way errors are detected and appraised; together these constitute, in systems terms, a 'guidance mechanism' for course assessment and correction.

Policy formulation can be treated as a dynamic process affected by information coming from subsequent stages in the policy-making process. The assumed complexity and turbulence of the policy environment will place the requisite emphasis on the management of perceived errors in policy performance. The way implementation will affect later rounds of policy reformulation will be seen to depend largely on the guidance mechanism in place. In general, the question of whether or not implementation prospects and experiences should be instrumental in fashioning policy revisions will be answered differently depending upon one's normative presumptions about policy-making and the guidance notions these imply. The idea of a guidance mechanism, while only a theoretical fiction, provides a useful set of concepts for organizing the chaos introduced by admitting institutional autonomy at the same time as one embraces a more dynamic conception of the policy process.

References

Anderson, C. (1987) 'Political Philosophy, Practical Reason, and Policy Analysis', in F. Fischer and J. Forester (eds) *Confronting Values in Policy Analysis*. Beverly Hills, CA: Sage Publications.

Argyris, C. and D. Schon (1978) *Organizational Learning*. Reading, MA: Addison-Wesley.

Atkinson, A. and J. Stiglitz (1980) *Lectures on Public Economics*. New York, NY: McGraw-Hill.

Bobrow, D. and J. Dryzek (1987) *Policy Analysis By Design*. Pittsburgh, PA: University of Pittsburgh Press.

Braybrooke, D. (1985) 'Scale, Combination, Opposition – A Rethinking of Incrementalism', *Ethics*, 95: 920–33.

Braybrooke, D. and C. Lindblom (1963) *A Strategy of Decision*. New York, NY: The Free Press.
Child, J. and A. Kieser (1981) 'The Development of Organizations Over Time', in P. Nystrom and W. Starbuck (eds) *Handbook of Organizational Design*. New York, NY: Oxford University Press.
Cobb, R. and C. Elder (1981) 'Communication and Public Policy', in D. Nimmo and K. Saunders (eds) *Handbook of Political Communication*. Beverly Hills, CA: Sage Publications.
Cohen, M., J. March and J. Olsen (1972) 'A Garbage Can Model of Organizational Choice', *Administrative Science Quarterly* 17: 1–25.
Davis, O., M. Dempster and A. Wildavsky (1966) 'A Theory of the Budgetary Process', *American Political Science Review* 60: 529–47.
Davis, O., M. Dempster and A. Wildavsky (1974) 'Towards a Predictive Theory of Governmental Expenditure', *British Journal of Political Science* 4: 419–52.
DiMaggio, P. (1988) 'Interest and Agency in Institutional Theory', in L. Zucker (ed) *Institutional Patterns and Organizations*. Cambridge, MA: Ballinger.
Dror, Y. (1985) *Policy Making Under Adversity*. New Brunswick, NJ: Transaction.
Dunn, W. (1982) *Public Policy Analysis: An Introduction*. Englewood Cliffs, NJ: Prentice-Hall.
Dunsire, A. (1986) 'A Cybernetic View of Guidance, Control and Evaluation in the Public Sector', in F. Kaufmann, G. Majone and V. Ostrom (eds) *Guidance and Control and Evaluation in the Public Sector*. New York, NY: Walter de Gruyter.
Emery, F. and E. Trist (1965) 'The Causal Texture of Organization Environments', *Administrative Science Quarterly* 18: 21–32.
Feldman, M. and J. March (1981) 'Information in Organizations as Signal and Symbol', *Administrative Science Quarterly* 26: 171–86.
Goodin, R. (1982) *Political Theory and Public Policy*. Chicago: University of Chicago Press.
Hays, R. Allen (1985) 'Perceptions of Success or Failure in Program Implementation: The "Feedback Loop" in Public Policy Decisions', *Policy Studies Review* 5: 51–67.
Heclo, H. (1978) 'Issue Networks and the Executive Establishment', in A. King (ed) *The New American Political System*. Washington, DC: American Enterprise Institute.
Hedberg, B. (1981) 'How Organizations Learn and Unlearn', in P. Nystrom and W. Starbuck (eds) *Handbook of Organizational Design*. New York, NY: Oxford University Press.
Hirschman, A. (1970) *Exit, Voice and Loyalty*. Cambridge, MA: Harvard University Press.
Hjern, B. and D. Porter (1981) 'Implementation Structures: A New Unit of Organization Analysis', *Organization Studies* 2/3: 211–27.
Hochschild, J. (1984) *The New American Dilemma: Liberal Democracy and School Desegregation*. Chicago: University of Chicago Press.
Hogwood, B. and L. Gunn (1984) *Policy Analysis for the Real World*. New York, NY: Oxford University Press.
Hogwood, B. and B. G. Peters (1983) *Policy Dynamics*. New York, NY: St Martin's Press.
Hood, C. (1986) *Administrative Analysis*. Sussex: Wheatsheaf.
Jones, C. (1984) *Introduction to the Study of Public Policy*, 3rd edn. Belmont, CA: Wadsworth.
Kingdon, J. (1984) *Agendas, Alternatives, and Public Policy*. Boston: Little, Brown.
Krasner, S. (1984) 'Approaches to the State: Alternative Conceptions and Historical Dynamics', *Comparative Politics* 16: 223–46.
Landau, M. (1973) 'On the Concept of a Self-Correcting Organization', *Public Administration Review* 33: 533–42.
Levitt, B. and James G. March (1988) 'Organizational Learning', *Annual Review of Sociology* 14: 319–40.
Lindblom, C. (1977) *Politics and Markets*. New York, NY: Basic Books.
Lindblom, C. and D. Cohen (1979) *Usable Knowledge: Social Science and Social Problem Solving*. New Haven, CT: Yale University Press.

Linder, S. and B. G. Peters (1987) 'A Design Perspective on Policy Implementation: The Fallacies of Misplaced Prescription', *Policy Studies Review* 6: 459–75.
Lowi, T. (1979) *The End of Liberalism*, 2nd edn. New York, NY: Free Press.
McFadden, D. (1975) 'Revealed Preferences of a Government Bureaucracy', *Bell Journal of Economics* 6: 401–16.
March, J. (1987) 'Ambiguity and Accounting: The Elusive Link Between Information and Decision Making', *Accounting, Organization and Society* 12: 153–68.
March, J. and J. Olsen (1984) 'The New Institutionalism: Organizational Factors in Political Life', *American Political Science Review* 78: 734–48.
March, J. and H. Simon (1958) *Organizations*. New York, NY: John Wiley.
May, J. and A. Wildavsky (eds) (1978) *Policy Cycles*. Beverly Hills, CA: Sage Publications.
Nakamura, R. and F. Smallwood (1980) *The Politics of Policy Implementation*. New York, NY: St. Martin's Press.
North, D. (1986) 'The New Institutional Economics', *Journal of Institutional and Theoretical Economics* 142: 230–45.
O'Toole, L. (1986) 'Policy Recommendations for Multi-Actor Implementation: An Assessment of the Field', *Journal of Public Policy* 6: 181–210.
Peters, B. G. (1986) *American Public Policy: Promise and Performance*, 2nd edn. Chatham, NJ: Chatham House.
Polsby, N. (1984) *Political Innovation in America*. New Haven, CT: Yale University Press.
Rich, R. (1981) *Social Science Information and Public Policy Making*. San Francisco, CA: Jossey-Bass.
Sabatier, P. (1987) 'Knowledge, Policy-Oriented Learning and Policy Change', *Knowledge: Creation, Diffusion, Utilization* 8: 649–92.
Schroder, H., M. Driver and S. Streufer (1967) *Human Information Processing*. New York, NY: Holt, Rinehart and Winston.
Scott, W. R. (1987) 'The Adolescence of Institutional Theory', *Administrative Science Quarterly* 32: 493–511.
Selznick, P. (1957) *Leadership in Administration*. New York, NY: Harper & Row.
Simon, H. (1947) *Administrative Behavior*. New York, NY: Macmillan.
Steinbruner, J. (1974) *The Cybernetic Theory of Decision*. Princeton, NJ: Princeton University Press.
Sundquist, J. (1968) *Politics and Policy*. Washington, DC: Brookings Institution.
Ungson, G., D. Braunstein and P. Hall (1981) 'Managerial Information Processing: A Research Review', *Administrative Science Quarterly* 26: 116–34.
Weinberg, G. and D. Weinberg (1979) *On the Design of Stable Systems*. New York, NY: Wiley Interscience.
Weiss, C. (1977) 'Research for Policy's Sake: The Enlightenment Function of Social Research', *Policy Analysis* 3: 531–45.
Weiss, C. and M. Bucavalas (1980) 'Truth Tests and Utility Tests: Decision-Makers' Frames of Reference for Social Science Research', *American Sociological Review* 45: 302–13.
Wildavsky, A. (1979) *Speaking Truth To Power*. Boston, MA: Little, Brown.
Wildavsky, A. (1987) 'Choosing Preferences by Constructing Institutions: A Cultural Theory of Preference Formation', *American Political Science Review* 81: 3–21.
Wildavsky, A. (1988) *Searching For Safety*. New Brunswick, NJ: Transaction Press.
Zucker, L. (1987) 'Institutional Theories of Organization', *Annual Review of Sociology* 13: 443–64.

STEPHEN H. LINDER is Associate Professor in Management and Policy Sciences at the University of Texas School of Public Health. He received the PhD in Political Science from the University of Iowa and has held faculty positions at UCLA and Tulane. His principal interest lies in policy

analysis, currently focusing on policy design and implementation issues. His most recent work has appeared in *Policy Studies Review*, *Milbank Quarterly*, and *Journal of Policy Analysis and Management*. ADDRESS: School of Public Health, PO Box 20186, The University of Texas, Houston, TX 77225, USA.

B. GUY PETERS is the Maurice Falk Professor of American Government at the University of Pittsburgh. He received the PhD in Political Science from Michigan State University and has taught at the University of Delaware, Emory and Tulane Universities. His most recent book (with Brian Hogwood) is entitled *The Pathology of Public Policy*. He is currently the co-editor of *Governance: An International Journal of Policy and Administration*, and his research interests include the design of public policy and comparative public administration. ADDRESS: Department of Political Science, 4T20 Forbes Quadrangle, University of Pittsburgh, Pittsburgh, PA 15260, USA.

Paper submitted 4 August 1988; accepted for publication 28 February 1989.