



An Adaptive Model of Bureaucratic Politics

Author(s): Jonathan Bendor and Terry M. Moe

Source: *The American Political Science Review*, Vol. 79, No. 3 (Sep., 1985), pp. 755-774

Published by: American Political Science Association

Stable URL: <http://www.jstor.org/stable/1956842>

Accessed: 25/08/2008 16:21

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=apsa>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.

An Adaptive Model of Bureaucratic Politics

JONATHAN BENDOR
TERRY M. MOE
Stanford University

In this article we outline a new framework for the formal analysis of bureaucratic politics. It departs from standard neoclassical approaches, notably those of Niskanen (1971) and Peltzman (1976), in several important respects. First our approach explicitly models a system of three-way interaction among bureaus, politicians, and interest groups. Second, it allows for institutional features of each type of participant. Third, it is a model of dynamic process. Fourth, participants make choices adaptively rather than optimizing. Fifth, participants are only minimally informed.

The result is a dynamic model of adaptive behavior, very much in the spirit of Simon's (1947) behavioral tradition, that offers a new perspective on political control, bureaucratic power, and the "intelligence of democracy."

For decades, students of public administration have stressed that we can only understand bureaucratic behavior by looking beyond the boundaries of administrative organization (Long, 1949; Simon, Smithburg, & Thompson, 1950). Agencies operate in continual exchange with an institutionally structured environment of politicians and interest groups, and it is insight into this system of interactions that holds the key to explanation. Popular theories of bureaucratic politics, ranging in emphasis from iron triangles (Cater, 1964; Freeman, 1955; Truman, 1951) to regulatory capture (Bernstein, 1955; Huntington, 1952; Stigler, 1971) to interest group liberalism (Loui, 1969), are clearly in agreement on this most general of points, and empirical evidence suggests that it is surely a valid one. This consensus sets bounds on what is considered controversial within the field. We argue now about how much discretion bureaus have, the relative influence of politicians and interest groups, and how these properties vary across different agencies. We do not argue about whether politicians and interest groups are central to an understanding of bureaucratic politics.

Formal models have made some useful contributions to the area, but they have uniformly shied away from modelling the kind of three-way inter-

action routinely described in the empirical literature. Perhaps the two most highly regarded efforts to model bureaucratic politics are Peltzman's (1976) theory of regulatory behavior and Niskanen's (1971) theory of budgets and bureaucratic supply. Consider, very briefly, their basic features.

Peltzman builds upon the earlier work of Stigler (1971) and Posner (1974) in an attempt to place regulatory capture within a more general theoretical context. He formalizes the problem by positing a regulator who makes policy decisions about transferring wealth between members of two interest groups (in effect, business and consumers). The regulator is a bureaucrat-politician who, based on perfect information about the groups, chooses among administrative options in an effort to maximize votes. Interest group members then decide to vote for or against him based on his wealth-transfer decisions and on their abilities to organize for political action. Analysis of these supply and demand factors then leads to an equilibrium solution that provides insight into the conditions under which the regulator may be captured by one of the groups.

Niskanen's work is a pioneering effort to explain budgetary outcomes—and to show why government is "too big"—with reference to special properties of the relationship between bureaucrats and legislators. He argues that bureaucrats value larger budgets (and perhaps "slack"—see Niskanen, 1975), and that they have distinct advantages over legislators in the bargaining game that determines budgetary outcomes: they have a near monopoly over information about the true costs of production, they are perfectly informed about legislative demand for their services, they are able to control the legislature's agenda by submitting

Received: March 20, 1984

Revision received: September 28, 1984

Accepted for publication: January 24, 1985

We are grateful for the many helpful comments we received from Robert Axelrod, Michael Cohen, John Ferejohn, Thomas Hammond, Richard Hausman, Roger Noll, John Scholz, Herbert Simon, Aaron Wildavsky, the members of the Stanford Seminar on Politics and Organizations, and the *APSR* referees.

take-it-or-leave-it budgetary proposals, and they are often aided by high-demand legislative committees. By putting these weapons to use, bureaucrats are able to win the bargaining game, securing larger budgets and (if they like) operating inefficiently.

Peltzman and Niskanen thus contribute interesting new ways of thinking about two of the most important issues of bureaucratic politics: which constituencies win and which lose, and how the level and efficiency of bureaucratic provision compare to the social optimum. Peltzman addresses the first, offering an innovative perspective on the relationship between clientele characteristics and bureaucratic resource allocation. Niskanen addresses the second, shedding new light on the relation between bureaucratic power and governmental outcomes. Yet their analyses are also based on five distinctive simplifications that, in imposing analytical structure on bureaucratic politics, condition the questions they raise and the conclusions they derive. In considering whether these assumptions seem reasonable, we are immediately struck by the most obvious of the five. 1) Both theories clearly ignore fundamental dimensions of political interaction. Peltzman combines politicians and bureaucrats into one actor, implicitly assuming that the dynamics of political control, such as oversight, budgeting, divergence of bureaucratic and political goals, and competition for control among politicians, make no significant difference for bureaucratic behavior. Niskanen gives no systematic attention to clientele groups, assuming their demands are represented by the legislature; and he models the legislature as a passive participant that submits to bureaucratic dominance without putting its own substantial resources to strategic use (see Miller & Moe, 1983). In neither is serious attention given to the distinct, interdependent roles of bureaucrats, politicians, and interest groups.

They also share several other properties which, though less objectionable, are major simplifications that structure their analyses. 2) Both make no real attempt to disaggregate institutional decision makers. Bureaus, legislatures, and interest groups are essentially treated as individual decision units. 3) Both are exercises in comparative statics. They are not dynamic models of interaction, and they tell us nothing about the political process that gives rise to their equilibrium results. 4) Both assume that all decision makers optimize, and thus that they choose on the basis of highly complex calculations beyond the capacity of most people. 5) Both include assumptions of perfect information (although Niskanen does so asymmetrically—the bureau is perfectly informed, the legislature is not).

All of these features are in fact quite common

in neoclassical economic models, and they are surely understandable responses to the great complexity of bureaucratic politics. It is one thing for political scientists to point to multiparty interaction, dynamic processes, institutional context, uncertainty, limitations on cognition, and the like, and quite another to construct models that somehow incorporate these properties. In all modelling efforts, simplification is absolutely essential. The problem is always to figure out which aspects can safely be assumed away and which cannot—recognizing that there are different ways to simplify, and that even the most elegant, deductively powerful models run the danger of being very misleading if founded on simplifications of the wrong kinds.

In this article we outline a model—really, a general framework accommodating a whole family of models—that allows us, via computer simulation techniques, to pursue the same basic issues of constituency influence and governmental size addressed by Peltzman and Niskanen, yet also reflects a very different and more general approach to the formal analysis of bureaucratic politics. It departs from their models along each of the five simplifying dimensions noted above. 1) Most important, it is a model in which the central outcomes of interest—governmental outputs, budgets, and bureaucratic efficiency—are in every sense jointly determined by the interdependent decisions of bureaucratic, legislative, and interest group participants. Attention focuses on the entire system of political relationships and on the integral roles these participants play as part of that larger system. 2) It allows us, if we like, to incorporate institutional features of each type of participant, for example, by recognizing decision processes internal to bureaus or legislatures. In this article we take one of many possible steps in this direction by introducing a majority-rule legislature. 3) The model is dynamic, allowing us to explore the process of interaction and adjustment as it unfolds over time. 4) Decision makers are assumed to adapt in simple ways to their environments, moving in directions that appear to promise them greater utility. They do not optimize, nor do they carry out complex calculations. 5) Participants are only minimally informed about their environments and about each other, basing their adaptive decisions on feedback about the success of their prior decisions.

These elements have a certain natural compatibility. A model of three-way interaction is intrinsically complex, and the assumption of optimization would only magnify the problem by forcing inquiry into an array of complicated game-theoretic considerations. Because formally incorporating a third actor jeopardizes the mathematical tractability of bureaucratic models in the

neoclassical tradition, it is no accident that they address only dyadic interaction.¹ Assumptions of adaptive behavior, on the other hand, are not only far more reasonable empirically, but they are also particularly well suited to the task of modeling dynamic, interactive decision making under uncertainty. From a purely technical standpoint, it is a relatively easy matter to model individual decision makers in terms of utility functions and adaptive rules, impose a structure (reflecting, e.g., institutional context) on their interactions, and then allow a computer to map out the corresponding implications for individual choices and collective outcomes as they occur over time. Large numbers of interacting decision makers and varying institutional contexts can be handled without any real difficulty, given the computer's admirable skill at calculation. This technical feasibility does not, of course, imply that any given model will provide an adequate or insightful explanation of bureaucratic politics. But it is an important advantage which, in combination with the greater realism of its assumptions, sets this approach apart from the conventional neoclassical approach and makes it particularly worthy of investigation.

Thus, our interest in constructing a model that incorporates a basic fact of political life—interaction among bureaucrats, politicians, and interest groups—has led us to follow a strategy that departs from the norm. What we are proposing, in fact, is a dynamic model of adaptive behavior that is very much in the spirit of the behavioral tradition associated with Simon (1947), March and Simon (1958), Cyert and March (1963), Crecine (1969), Cohen, March, and Olsen (1972), Axelrod (1976), Padgett (1980), Cohen (1981, 1984), and Cohen and Axelrod (1984). Works in this tradition, grounded in the limitations on human decision makers, have long pointed to adaptation and dynamic process as central to an understanding of organizational behavior, and they have pioneered in the development of computer-assisted models of organizational behavior. Because our model is in some sense an application of this theoretical perspective to the study of bureaucratic politics, its success (or lack of it) says something not only about the nature of bureaucratic politics, but also about the value of putting Simon's behavioral tradition to new use.

The Structure of the Model

Our general framework can be applied in a variety of ways. Here we focus on bureaucratic

behavior in a classic type of political environment—one characterized by interest group conflict—and we develop a simple model to explain bureaucratic output, budgets, and efficiency. The model's central components are an agency interested in some combination of bigger budgets, more slack, achieving policy goals, and avoiding oversight;² a legislature made up of 101 elected politicians, each of whom is interested only in getting re-elected;³ and two interest groups, one that benefits from the agency's program and one that is hurt.⁴ None of this is anchored in particular agencies or policy areas. For purposes of illustration, however, we will develop the analysis below with reference to an agency engaged in consumer protection regulation. On this interpretation, which is obviously just one of many possibilities, the agency's output is taken to be the level of regulatory enforcement, and its enforcement activities are assumed to be supported by consumers and opposed by business.

In broad outline, the model is designed to reflect the circular flow of influence characteristic of representative government. Citizens pressure legislators through elections, legislators influence the bureau through budgets and oversight, the bureau affects citizens through the costs and benefits generated by regulatory enforcement—and the circle is closed when citizens link their electoral support to legislators' positions on agency-relevant issues. This iterative process works itself out in the following way.⁵ (See Figure 1.)

Every period, the legislature transmits to the bureau a new budget that may differ by as much as 10% from its budget in the prior period. Each

²Although there is no agreement in the literature on exactly what motivates bureaucrats, these elements are prominently mentioned. Note that our model allows the agency to value just one of these, or any weighted combination, giving us a basis for exploring the implications of alternative bureaucratic goals.

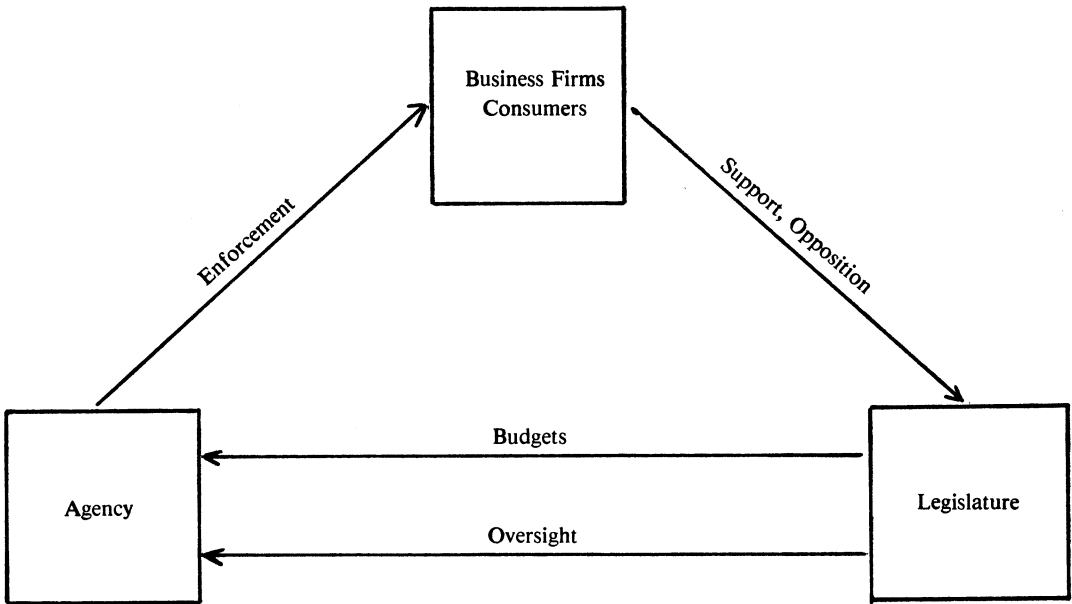
³On the importance of the re-election incentive, see Mayhew (1974) and Fiorina (1974). We chose 101 for the size of the legislature because it is large enough to prevent the institutional dynamics from being obscured by the idiosyncratic behavior of one or a few legislators, and as an odd number, it allows for a clear majority winner. Had we chosen a larger (odd) number, say 501, the model's results would be identical in all essential respects to those presented here.

⁴These groups need not be monolithic organizations, and may even be sets of unorganized citizens. We think of them, abstractly, as sectors that generate votes, money, and other resources in response to political outcomes. Their comparative effectiveness in doing so can be varied by parameter settings of the model.

⁵Readers interested in the finer details may write the authors for a copy of the computer program.

¹For an interesting exception relying on general equilibrium analysis, see Fiorina and Noll (1978).

Figure 1. The Flow of Influence



legislator must take a position somewhere on a scale between -10 and $+10$ percent.⁶ In deciding how to vote, the individual legislator compares his past votes with past changes in the mix of support and opposition he has received from the two interest groups, that is, changes in his electoral security. If he changed his position in a pro-agency direction last period—say, by moving from $+2.5$ to $+3.1$ on the budget scale—and his electoral security subsequently increased, then he will want to take another step in this direction, with the size of the step determined by the size of his utility gain. Similarly, had this earlier move been sanctioned by a drop in electoral security, he would now reverse course and support a smaller, perhaps even a negative, budget increase. The same kind of adaptive logic applies for other possible permutations. Once each legislator has adopted a new position, the entire legislature votes and the median position determines the new budget that is transmitted to the agency.

The legislature also engages in oversight, although it is oversight of a very limited sort. It is

activated in the current period only if the legislature increased the agency's budget in the prior period and subsequently discovers, via feedback from the interest groups, that the agency produced less enforcement with this larger amount of money.⁷ This obvious inefficiency results in "hassling" from the proconsumer legislators (those occupying positions greater than 0 on the budget scale), and the intensity of hassling varies with the political weight of this consumerist contingent (i.e., with the number of legislators to the right of 0, and how far to the right they are on average). If oversight is not triggered in the current period, then the level of hassling decays exponentially.

The legislature's work ends once the new budget and hassling levels are decided. These new quantities are now inputs to the agency that alter its utility, telling it whether the environment is approving or disapproving of its past behavior. As noted, its utility function contains one or more of four potential components—budget, slack (the amount of the budget not devoted to enforcement), hassling, and policy (in the form of an enforcement level the agency prefers to all others,

⁶Initial positions are assigned by random draw from a normal population with mean 0 and standard deviation 4. Hence, the first change in the budget equals the median vote of this random sample. The random character of the first budgetary change does not affect the model's basic results.

⁷This oversight strategy, similar in spirit to the fire-alarm monitoring analyzed by McCubbins and Schwartz (1984) and Weingast (1983), economizes on the information-gathering resources of the legislature.

its ideal point)—and therefore the legislature's impact on its utility will depend in part on precisely what combination of these the agency happens to value. When relevant, budget and slack contribute positively to utility, whereas both hassling and departure from the agency's ideal point contribute negatively. Once the agency has assessed its utility change, it follows an adaptive strategy in setting a new value for its choice variable, efficiency. Thus, if it increased its efficiency last period and was rewarded with an increase in utility, it will increase efficiency again in the current period. And so on for the other logical possibilities.

Efficiency is defined as the fraction of the budget spent on enforcement, and therefore ranges between zero and one. Enforcement is generated as a linear function of the actual amount of money spent on it—that is,

$$\text{enforcement} = c \cdot \text{efficiency} \cdot \text{budget}, \quad (1)$$

where c is a constant representing the units of enforcement purchased by a unit of spending.⁸

Having chosen a new level of efficiency, the agency generates a new level of enforcement which in turn yields new levels of costs for business and benefits for consumers. The business costs imposed by regulation increase at an increasing rate (increasing marginal costs), while benefits to consumers increase at a decreasing rate (decreasing marginal benefits).⁹ Specifically,

⁸One reasonable interpretation of enforcement "units" is person-hours devoted to enforcement. Because there is no unique scale for such units, c can be any positive constant; for convenience, we assume throughout that $c = 1$. Similarly, there is no unique scale for measuring budgets. Although the budgets of real agencies can range from the millions to the billions, for simplicity—and because the scale does not affect the model's basic results—we have set the initial budget equal to 100. This, together with our assumptions about efficiency and the constant c , determine the scale for enforcement.

⁹This is a standard assumption and is certainly a reasonable place to start. Yet costs and benefits may take on far more complex patterns empirically, and they may be interrelated in various ways—as, for example, when regulatory costs bring about business failures and subsequent loss of jobs for consumers. Our curves are not inconsistent with many of these scenarios (e.g., diminishing marginal benefits for consumers captures the notion that the rewards of enforcement are diluted by its negative by-products), but new work along these lines, empirical as well as analytical, is clearly needed if our theories are ultimately to be well designed.

$$\begin{aligned} \text{benefits} &= k1 \cdot \text{enforcement} \\ &- k2 \cdot \text{enforcement}^2 \end{aligned} \quad (2)$$

$$\text{costs} = k3 \cdot \text{enforcement}^2 \quad (3)$$

where $k1$, $k2$, and $k3$ are positive constants. In our initial version of the model, consumers evaluate the benefits of regulation without taking into account the taxes needed to finance the agency. In the final version, consumers base their evaluation of regulation on net benefits: the value of enforcement minus its tax costs.¹⁰ Under both versions the social optimum is the point at which enforcement's marginal costs and marginal benefits are equal, with budgetary costs incorporated into the calculation.

For both business and consumers, the political stakes rise as enforcement rises, and both therefore devote more resources (campaign contributions, votes, etc.) to the electoral process in support of politicians who promote their interests. Consumers increase their contributions in proportion to the benefits they derive from enforcement; firms increase their contributions in proportion to the costs imposed on them:


$$\begin{aligned} \text{consumer's political resources} &= k4 \\ &+ k5 \cdot \text{benefits} \end{aligned} \quad (4)$$

$$\text{firms' political resources} = k6 + k7 \cdot \text{costs}, \quad (5)$$

where $k4$ through $k7$ are positive constants. (See Figure 2.) Thus, the more an interest group is helped by a program, the more it is willing to contribute to politicians supportive of the agency. Similarly, the more a group is hurt by a program, the more it is willing to contribute to politicians who oppose the agency. Empirically, however, the existing system of interest organizations is unbalanced in favor of business, and this imbalance is enhanced by free-rider problems that disproportionately plague the emergence of new consumer groups (Olson, 1965; Stigler, 1971). There is accordingly a closer correspondence in the model between additional costs imposed on business and its political contributions than is true of the analogous benefit-to-contributions translation for con-

¹⁰Because the budgets of regulatory agencies are only a tiny fraction of the federal budget, it is a reasonable first approximation to assume that an agency's beneficiaries are more aware of benefits than tax costs. Moreover, we gain a clearer understanding of the distinctive contributions of the model's various components by dealing first with the simple nontax world and introducing taxes later.

Figure 2. The Pluralist Equilibrium in the Taxless Model



Cost and Benefit Functions (not graphed)

Costs = $.023E^2$
 Benefits = $20E - .04E^2$

- E = Enforcement
- PE = Pluralist Equilibrium
- E^* = Social Optimum
- R_f = Political Resources of Firms
- R_c = Political Resources of Consumers

Resource Functions

Firms $R_f = 10 + .20 (.023E^2)$
 = $10 + .20$ (costs)
 Cons. $R_c = 50 + .05 (20E - .04E^2)$
 = $50 + .05$ (benefits)

sumers. (In equations (4) and (5), constant k_5 is smaller than constant k_7 .)¹¹

¹¹Note that our interpretation of organizational advantage is that firms are better able to respond to changes in enforcement than are consumers. This in itself says nothing about the absolute level of resources when enforcement is zero. We assume, for present purposes, that when there is no enforcement, consumers have more political resources than firms do. This makes sense for the following reasons: consumers are numerous, they vote, voting requires little organization, and when there is no regulation, the issue is likely to be salient to voters. Moreover, the opposite assumption is easily accommodated by our model, as we suggest in note 24 below.

Once business and consumer groups decide upon resource levels, they expend them for or against individual legislators in accordance with the latter's positions on agency budgets. Legislators favoring budgetary increases are judged as pro-agency and therefore as proconsumer; they get resource support from consumers and resource opposition from business in amounts that are a positive function of the size of the budget increase they favor.¹² The same logic

¹²In versions of the model where consumers take taxes into account, consumers will support legislators voting for smaller budgets if the taxes required to finance the

applies for legislators favoring reduction in the agency budget.

Legislators now find that owing to their earlier adaptive shifts in position on the budgetary scale, they face altered patterns of electoral support and opposition—that is, changes in their utility—anchored in the adjusted evaluations of business and consumer groups. This new feedback prompts them to adapt by choosing a new budgetary position for the coming period, following the decision rule outlined above. This completes the circle and sends the process into another iteration.

Before summarizing the results, we would like to pause briefly to underscore a few basic aspects of this model. First, all of its decision makers are self-interested: they seek to realize their own goals, whatever they might be, and they are not directly concerned with the well-being of other participants.¹³ Thus, if the model generates societally attractive outcomes, or if one participant makes decisions beneficial to one or more of the others, these will rarely be intended consequences, but rather the by-products of self-interested behavior. In this respect, we adhere to a central theme of rational choice models of regulation and bureaucracy, and we adhere as well to a classic line of pluralist theory.

Second, our participants are endowed with very little knowledge. Legislators do not know the marginal electoral value of another dollar added to the regulatory agency's budget. Bureaucrats do not know the marginal gain of one more unit of enforcement. What these decision makers do know is rather modest. Legislators know such things as how they voted on the previous appropriations bill and how full their campaign war-chests are. The regulatory agency knows only what percentage of the budget it has devoted to enforcement and how well off it is; it knows nothing directly of the effects of regulation, and it must choose a level of efficiency in ignorance of the true consequences of its actions. In short, as incrementalist analyses of the political process have long contended, these participants adapt as best they can based on what little they know: They muddle through (Braybrooke & Lindblom, 1963; Lindblom, 1959).

Third, our participants are simple adaptive decision-makers. They learn about their environ-

ments only to the extent that they link new information about changes in behavior with new information about changes in utility, and they adapt by following a trial-and-error procedure that prompts them to repeat rewarded actions and to avoid sanctioned ones. They do not generalize or otherwise develop belief structures about their environments, nor do they try to guess the future. The specific model we use here, then, assumes perhaps the simplest possible adaptive scheme—a reasonable place to start, given that we are moving into uncharted territory. More complex adaptive strategies can always be introduced once simple models of this sort are better understood.¹⁴

Fourth, the distribution of influence is typically pluralistic: no one is in charge. The structure of modern democratic government virtually guarantees that citizens, politicians, and bureaucrats will interact through a network of relationships in which no single participant can truly dominate or be entirely autonomous, and in which each must adapt to the decisions of others.

Finally, the model emphasizes the role of indirect influence, of the importance—more, the necessity—in American politics of getting what you want via other decision makers. The bureau likes bigger budgets, but cannot directly influence the legislature. Its behavior directly affects the interest groups, but these groups cannot influence the agency directly; they must try to induce the legislators to move the bureaucracy in the desired direction. The legislators' electoral fortunes are directly affected by the groups, but in order to better their chances the legislators must work through the agency. These indirect paths of influ-

¹⁴The adaptive strategy our actors follow, a simple trial-and-error procedure operations researchers call "hillclimbing," is a weak strategy: it gets trapped on local maxima, makes inferential errors, and is often quite slow. And its implied cognitive processes underestimate the sophistication of real decision makers. But positing trial-and-error yields a significant modelling benefit: although a weak strategy, it is highly general (Rich, 1983). Unlike most optimization techniques, hillclimbing can be used in an extraordinarily wide variety of task environments—singlepeaked or multi-peaked, linear or nonlinear, deterministic or stochastic—and the modeller can avoid specifying detailed beliefs and heuristics of the decision makers. In contrast, when psychologists model a problem-solver facing a particular task, they find that the heuristics are extremely context-specific. Chess heuristics do not closely resemble theorem-proving heuristics (Newell & Simon, 1972). Since political scientists are more interested in systemic behavior than in individual behavior, we are willing to trade some accuracy at the microlevel for generality at the macrolevel, and in the class of adaptive strategies, hillclimbing is one of the most general.

bureau outweigh the benefits of regulation. Typically, however, consumers help legislators voting for larger budgets and oppose those voting for reductions.

¹³This is not invariably true. In some runs the agency has policy goals that do not derive from preferences for budget or slack. These policy goals could derive from more general, nonegoistic ideologies. But the general sense of the model is of self-interested behavior.

ence and the uncertainties they create are the hallmarks of democratic government.

Simulation Results

A computer model of this sort necessarily has many parameter settings, and this can obscure its central tendencies. To highlight major patterns, we will first present simple versions and then add complications one at a time.

Version 1: Budget-Maximizing Agency

Clearly, an agency with only one goal is simpler to analyze than an agency with several. Consistent with most formal theories of bureaucracy, our baseline version is a regulatory agency that cares only about budgets. The outcomes generated by this system—the equilibrium levels of enforcement, efficiency, and budget—are presented in Table 1. We suspect they surprise most readers. Formal models of bureaus, from Niskanen on, have typically concluded that budget maximization leads either to gross inefficiency or gross oversupply of output. Indeed, this is approaching the status of conventional wisdom. But our model suggests quite a different effect of budget maximization: that it can make an agency responsive to the power configuration in its environment.

The system's logic of adjustment is easy to discern. The agency quickly discovers that increases

in its efficiency are rewarded by the legislature—whose members, unbeknown to the agency, are disproportionately rewarded by consumer groups for supporting budget increases; this occurs because, at the initial enforcement level of 80 (an arbitrary starting point), consumer groups out-contribute business groups, and many legislators are motivated to shift to the right on the budget scale.¹⁵ The agency adapts by jacking up its efficiency. The combination of higher efficiency and bigger budgets provided by the legislature rapidly increases enforcement. This increase, however, eventually mobilizes business, and the consumers' early success proves to be self-limiting. Once business outspends consumers (which, given equations (4) and (5), occurs after enforcement exceeds 184), legislators tend to be rewarded for cutting rather than increasing the agency's budget. The legislature therefore reverses course—yet it will not go far below 184, for then consumers begin to outspend business again, and a legislator who continues to vote against the agency will tend to see his opponent's war chest rise.

These oscillatory adjustments damp out over time as the system settles down to its equilibrium enforcement level of 184: the level at which the

¹⁵We experimented with different initial enforcement levels. Equilibrium outcomes remained the same.

Table 1. Simulation Results

Version	1	2	3	4	5	6A	6B
Parameters							
Budgets	Yes	Yes	Yes	Yes	Yes	No	No
Slack	No	Yes	No	No	Yes	No	No
Policy	No	No	No	No	No	Bus	Con
Oversight	No	No	No	No	Yes	No	No
Adaptation	Fast	Fast	Fast	Slow	Fast	Fast	Fast
First Move	Incr	Incr	Decr	Decr	Incr	Incr	Incr
Nontax model							
Enforcement	184	Deg	Deg	184	184	CR	184
Budget	249	Deg	Deg	360	371	CR	196
Efficiency	.74	Deg	Deg	.51	.50	CR	.94
Tax model							
Enforcement	178	87	56	171	170	101	174
Budget	210	1509	1485	321	382	1473	197
Efficiency	.85	.06	.04	.53	.45	.07	.88

Key. Deg, Degenerate (enforcement and efficiency \approx 0)
 CR, Compromise
 Con, Proconsumer
 Bus, Probusiness
 Incr, Increase
 Decr, Decrease

political resources of business and consumer groups are equal (as described by the intersection of their resource functions—see Figure 2). We call this the pluralist equilibrium, because it obviously formalizes the old pluralist idea of a balance of group forces. As the system approaches this equilibrium level, the other key variables—efficiency, the budget—hit steady-state values consistent with it, and legislators become tightly distributed around the median position of 0 on the budget scale.

In a pluralist political environment, then, the budget-maximizing agency of Version 1 cannot impose output levels that exceed the amount determined by the environment's power balance. Here, unlike in the usual formal models of bureaus, any oversupply of output is due not to the agency's lust for revenues, but to a power mismatch in the bureau's environment.¹⁶ Because legislators are responding to electoral rewards and sanctions from interest groups rather than to the blandishments of the bureau, the legislature simply refuses to provide the agency with appropriations that, combined with its high efficiency, would yield so much enforcement that the electoral chances of incumbents would be dimmed. The agency would surely like to obtain much larger budgets—but legislators have no incentive to go along, and they hold the purse strings. Empirically, this result makes perfect sense.

In part, the intuitive appeal of the excessive output hypothesis derives from a specific class of policy sectors, where the high demand interest group is well organized. This includes porkbarrel projects (Shepsle & Weingast, 1981) and, more generally, policies that create compact winners and diffuse losers. But, as Wilson (1980) has argued, many policy sectors lack this property, and in these the excessive output hypothesis is less plausible. Many regulatory agencies either have compact losers as well as compact gainers (the NLRB, the ICC) or, still worse from the stand-

point of the excessive output hypothesis, diffuse gainers and compact losers (the EPA, the Consumer Product Safety Commission). Our model incorporates Wilson's insight that the relative compactness of interest sectors influences bureaucratic outcomes. It does so by means of two parameters representing the fraction of the benefit or loss deriving from enforcement that is subsequently translated into political resources. The more compact the group, the higher the fraction. Because empirically business is more compact than consumers, the business parameter exceeds the consumer parameter in all versions of our model, and the pluralist equilibrium is therefore farther to the left—that is, the two resource curves intersect at a smaller level of enforcement—than would be the case if the groups were equally able to translate costs and benefits into political resources. Marginal changes in parameters from these settings have predictable effects. As consumers become more compact, the pluralist equilibrium shifts to the right, and as business becomes more compact, the pluralist equilibrium shifts to the left. Because outcomes are determined by the pluralist equilibrium in this version of our model, changes in relative compactness produce—through a chain of interactive adjustments—changes in regulatory performance, which formalizes Wilson's assertion that bureaucratic outcomes reflect group compactness and diffuseness.

Version 2: The Effect of Slack

Presumably bureaus prefer bigger budgets because of what they can buy, such as greater prestige, higher salaries, and more output or perquisites. Version 1 left those higher-order goals unspecified. We now specify one such goal: slack, defined here as appropriations devoted to any use other than output. The agency in Version 2 seeks slack in addition to larger budgets. Table 1 reveals the effect on enforcement, efficiency, and the budget. They differ dramatically from Version 1's pluralist equilibrium. In Version 2, efficiency plunges to zero. Ironically, though *all* the legislators quickly start voting in accord with consumers' desires, the consumers do not benefit at all, because enforcement is fast approaching zero despite the constantly increasing budget. This is puzzling: everyone is happy except consumers, even though they outspend firms in *every* period. Unlike Version 1, this result does not accord with intuitions based on simple power reasoning. What is going on?

The explanation is simple. A slack-seeking

¹⁶This difference between our model's results and Niskanen's reflects different assumptions about the distribution of resources rather than differences in decision-makers' goals. In his model, the deck is stacked in favor of the bureau: it has crucial informational advantages and in effect imposes outcomes on the legislature (Miller & Moe, 1983). In our model, because key resources are not monopolized by the bureau, outcomes are jointly determined. In a system with a pluralistic distribution of power, budget-maximization by a bureau does not necessarily lead to superoptimal budgets, just as in a competitive market, profit-maximization by a firm does not necessarily lead to superoptimal profits. (Indeed, in a perfectly competitive market, firms earn zero profits.) Such systems often do not exhibit simple links between individual motives and collective outcomes.

bureau quickly discovers that it can improve its well-being by lowering its efficiency. Because it decreases productivity faster than the legislature increases its budget, output falls. As enforcement falls, consumers continue to outcontribute firms. But though the legislature is responding to the wishes of consumers, it has only one instrument, the budget, and budgets are singularly ineffective instruments in such circumstances. A pro-consumer legislature is caught in a bind. On one hand, because the more powerful interest group desires more output, it makes some sense to give the agency more money. In this line of reasoning, budgets are resources needed to produce enforcement. But if the agency is a budget-seeker, both intrinsically *and* because it has multiple uses for appropriations, then revenues are also incentives. Thus increasing the budget not only gives the bureau more resources to do its job, but also rewards it for past behavior. Here the reward is inappropriate, since the agency's inefficiency caused the problem in the first place. But if budgets are the only instrument, what else can the legislature do? It does not make sense, at least not in the short run, to deal with a problem of diminished output by denying the agency the resources necessary to produce more output. Thus the bind, produced by a combination of the bureau's monopoly position and the legislature's impoverished repertoire, generates a pathological cycle, a degenerate case of huge budgets and negligible output.

The dilemma the legislature faces in trying to use the budget as both a production mechanism and an incentive mechanism has been recognized in the empirical literature (Quirk, 1981), and it emerges quite naturally from our model. By contrast, the conventional notion associated with Migué and Bélanger (1974) and Niskanen (1975) is that slack seeking leads to smaller, more nearly optimal budgets and outputs. In their models, a preference for slack prompts the agency to maximize the difference between the budget and production costs rather than simply the size of the budget itself, and this leads it to prefer (and receive, given its alleged power) a budget-output combination that is closer to the social optimum than the combination preferred by a budget-maximizing agency. In our model, slack seeking makes the situation unambiguously worse. Budgets explode rather than shrink. Outputs (enforcement) do shrink, but they go to zero as the agency chooses to become totally inefficient. And the agency is consistently rewarded for this behavior by a proconsumer legislature that throws money at it in order to gain the electoral support of consumer groups. Other things being equal, society is far better off with an agency solely concerned with maximizing its budget.

Version 3: Superstitious Learning

Noting that the agency in Version 2 garnered much bigger budgets than that in Version 1 did, the reader may now doubt our explanation of Version 1's behavior. If reducing productivity can procure such enormous budgets, why did the budget maximizer of Version 1 increase its efficiency? Is that adaptive behavior?

Adaptive, yes; optimal, no. The fundamental attribute of an adaptive decision-maker is pragmatism: if an action worked once, try it again. Subtle causal inferences about *why* an action paid off are avoided. This strategy opens the door to inferential errors. An action may have worked because of characteristics of the environment rather than of the decision. This is so in Version 1. Because initial enforcement is less than the pluralist equilibrium, the legislature will vote for larger budgets *regardless* of how the agency behaves. This is what learning theorists call a benign environment: anything the agency tries is rewarded. The combination of a benign environment and adaptive decision making renders the bureau's first move important, for it is the first move that, given the inevitable reward, leads the agency down one path (greater efficiency) or the other (less efficiency).¹⁷

The explanation for Version 1's behavior is now apparent. The agency's first move is a random choice: lacking experience to guide it, the bureau is equally willing to try increasing or decreasing efficiency. Version 1 is unlucky and increases efficiency. It is rewarded for doing so. It repeats the choice, and again is rewarded. Convinced it has discovered the road to bureaucratic paradise, it steadily boosts its productivity. The agency is the victim of *superstitious learning*: misunderstanding the causal structure of its environment, it has erroneously attributed its success to its own conduct.¹⁸

Confirmation of superstitious learning is provided by Version 3. Luckier than Version 1, its first move is to decrease efficiency. It too is rewarded by a larger budget. Equally convinced

¹⁷Ours is a model of deterministic adaptation: if increasing efficiency is followed by more utility, then the agency will increase efficiency in the next period with certainty. In a model of probabilistic adaptation, success would increase the odds that the agency would become more efficient. Probabilistic adaptation is empirically more realistic, and also less likely to result in superstitious learning because, even after an initial success, a decision maker has a chance of trying another alternative.

¹⁸It is well known that adaptive behavior in benign environments produces superstitious learning (Lave & March, 1975).

of its brilliance, it continues to reduce its productivity. The consequences are virtually identical to Version 2's explosive results: huge budgets, complete inefficiency, negligible output (Table 1). The control pathology of Version 2 is repeated. The proconsumer legislature teaches the bureau that inefficiency pays.

Therefore, the new conventional wisdom is not entirely wrong. If the legislature has only one instrument (the budget) and constituents ignore taxes, an adaptive budget-seeking bureau is as likely to become inefficient as it is productive. Indeed, only the extreme myopia of deterministic adaptation made the agency of Version 1 wind up at the pluralist equilibrium: the agency in Version 2 is far better off.

Version 4: The Rate of Adaptation

The superstitious learning of Version 3 introduced the idea that the power characteristics of the bureau's environment are mediated by institutional structure (the legislature's instruments) and by the agency's adaptive strategy. Unlike a simple pressure model, in our model the power of interest groups, though important, is not conclusive. Version 4 explores the effects of a second property of adaptation: the speed of adjustment.

A budget-maximizing agency is by definition indifferent to policy and therefore has no intrinsic desire to lower enforcement. Yet in Version 3, though consumers outspent firms, the agency drove enforcement to zero. Because enforcement = efficiency · budget, the agency was reducing efficiency faster than the legislature was increasing the budget. This observation suggests that the agency's rate of adjustment is crucial. Version 4, keeping everything else constant, introduces an agency that adapts more slowly. Table 1 records the effects of this change on the equilibrium enforcement, budget, and efficiency. Because Congress pumps in funds faster than the bureau becomes unproductive, enforcement reaches and temporarily passes the pluralist equilibrium. When this occurs, the legislature reverses its practice of granting it larger budgets, and the agency for the first time gets negative feedback suggesting that still lower efficiency is unwise. The eventual result is the pluralist equilibrium.

There are two important lessons here. First, the pluralist equilibrium of Version 1 is more than an accidental by-product of superstitious learning. The inertia introduced in Version 4 ultimately destroys the agency's benign environment and requires that it make choices in the face of negative feedback. That it settles down at the pluralist equilibrium, rather than somewhere else, indicates that there is something truly general about this particular level of enforcement. The second lesson

is that inertia performs a systemic function of real importance in this version of the model. Because bureaucratic inertia is ritually condemned in both popular and scholarly accounts, this is an intriguing result that underlines the value of further inquiry into its positive aspects.

Version 5: Oversight

Clearly the consumers are handicapped by the paucity of instruments available to the legislature. The sole response in Versions 2 and 3 to declining enforcement is to throw more money at the bureau, but this only teaches the agency to become still less productive. In the real world, however, such a vicious spiral could not continue indefinitely, for society would eventually be devoting all of its resources to regulation, and such a system is not viable. Systems that survive, therefore, must have mechanisms for checking these explosive forces. One such mechanism is taxes. We will turn to this shortly. Here we focus on a second mechanism, legislative oversight.

Specifically, Version 5 of the model endows the legislature with an oversight capacity to complement its budgetary authority, and assumes in addition that the bureau dislikes the hassling that oversight entails. Note that the latter condition is necessary if oversight is to have any impact on bureaucratic behavior. A bureau insensitive to hassling will simply ignore the legislature's complaints. Although this may certainly happen empirically (in fact, legislatures probably feel it happens all too often), it means that oversight would be superfluous to the model.

In keeping with the model's overall theme that decision makers rely on simple heuristics, the oversight mechanism is exceedingly crude. It is activated whenever two events occurred in the previous period: first, the legislature gave the agency a bigger budget; second, enforcement declined. Although we assume that legislators do not have a good grasp of the agency's internal affairs, they do hear about the enforcement level from the interest groups, and of course they know what the budget is. These two pieces of information suffice to yield an inference of mismanagement whenever the agency produces less output with more money. In such instances, the legislature hassles the agency with time-consuming hearings, bad publicity, threats to bureaucratic careers, and similar measures.

This mechanism is crude, but it turns out to be surprisingly potent. As we saw in Version 2, an agency that likes slack and budgets creates the degenerate solution: an explosive budget with efficiency and enforcement moving to zero. Version 5 takes the same context and introduces oversight. Simpleminded and myopic though it is, epi-

sodic legislative hassling has a striking impact on the process of bureaucratic adjustment. The agency is no longer in a benign environment: as it moves toward the degenerate solution, the combination of bigger budgets and drops in efficiency produce mounting sanctions in the form of hassling, until finally the costs of greater inefficiency outweigh the benefits. The bureau then reverses course. After a series of adjustments and counter-adjustments, it winds up at the pluralist equilibrium. Thus, legislative oversight interrupts the process of superstitious learning and saves the system from the degenerate solution, producing the same equilibrium level of enforcement as Version 1's budget-maximizing model—although, owing to its circuitous time path, at a higher budget and lower level of efficiency.

Oversight does not produce socially optimal behavior, nor does it guarantee legislators the best of all possible electoral worlds.¹⁹ It does, however, perform a crucial systemic function that puts a rein on the agency and allows regulation to "work." This occurs even though the legislature knows virtually nothing about bureaucratic costs, performance, or efficiency, and even though it reacts only on those occasions when the evidence of inefficient performance is painfully obvious.

The impact of oversight also helps to illustrate the special status of the pluralist equilibrium. We have presented five runs of our model now; two have produced the degenerate solution, three the pluralist equilibrium. In fact, it is a general property of our model (in the absence of taxes and agency policy preferences) that these are the only two solutions. The pluralist equilibrium, therefore, appears to be the natural equilibrium of a stable system, for it is the only solution compatible with the survival of the system as a whole.

Version 6: Agency Policy Preferences

In their motivational assumptions for bureaus, the standard formal models have clearly been influenced by the analogy of profit- or revenue-maximizing firms. Hence the emphasis on budgets and slack. But although it is ordinarily true enough that firms are not directly motivated by output, this is a less plausible assumption for bureaus, for there is good evidence that bureau-

cratic officials are often motivated by policy preferences (Aberbach & Rockman, 1976). The battles between liberal career civil servants and conservative political appointees in the Nixon and Reagan administrations are hardly battles over budgets and slack. They are battles over policy, over different notions of the "ideal" levels and directions of bureaucratic performance.

To isolate the effect of policy preferences, Versions 6A and 6B introduce agencies that care only about policy. We assume throughout that the agency's preferences are those of its dominant coalition, and that the composition of the dominant coalition is strongly influenced by presidential appointments. Although a more comprehensive model would surely include the president in an explicit fashion, this at least gives us an indirect opportunity to investigate whether presidents can shape regulatory performance by, in effect, changing some of the parameters of the system. In a regulatory context where firms oppose consumers, it is reasonable to categorize the agencies as either probusiness (perhaps reflecting appointments by a Republican president) or proconsumer (perhaps due to Democratic appointments). Abstractly, we consider an agency to be probusiness if its ideal enforcement level is less than the social optimum, as proconsumer if its bliss point exceeds the social optimum. To minimize ambiguity, we proceed by comparing a probusiness agency whose ideal level of enforcement is 100 (smaller than both the social optimum of 150 and the pluralist equilibrium of 184) with a proconsumer agency whose ideal point is 300 (larger than both).

Not surprisingly, a probusiness bias leads to less enforcement than the pluralist equilibrium. Seeking to reach its optimal enforcement of 100, the bureau tends to diminish its efficiency. However, the agency's success in reducing enforcement mobilizes consumers, which ensures a proconsumer legislature. Thus ensues an odd conflict between the two branches: the legislature pushes money on a reluctant bureau, which combats the added revenues by further decreasing its efficiency. The legislature responds by giving still more money, and the pattern continues. The result is that enforcement cycles in a compromise region bounded by the agency's ideal policy position and the pluralist equilibrium. As enforcement cycles, budgets eventually explode, and efficiency plunges virtually to zero.

Thus, a characteristic feature of a system where the pluralist equilibrium exceeds the agency's bliss point is conflict between the legislative and executive branches. (Our simple model does a fair job of simulating, for example, the short-term recent conflict between Congress and the EPA.) What happens if the agency is proconsumer? At first

¹⁹As one example of its suboptimality, note that the oversight strategy permits declining productivity if output is rising. Because this can happen only when the budget is increasing, this toleration represents systemic slack: when the budgetary climate is good, for either macroeconomic or political reasons, hierarchical control loosens. There is a strong similarity here to the dynamics of slack accumulation in the behavioral theory of the firm (Cyert & March, 1963).

blush, it appears that a proconsumer bias should produce a result symmetric to the probusiness bias: enforcement would equal a weighted average of the pluralist equilibrium and the agency's bliss point, with ongoing struggle between the legislature and the agency. Yet this is wrong. It reflects an extension of simple power reasoning that is inappropriate given the institutional structure of the model. As Version 6B suggests, the agency's decision variable, efficiency, cannot drive enforcement permanently higher than the pluralist equilibrium. The agency's policy preference has no significant effect. Jacking up productivity briefly augments enforcement, but this disproportionately mobilizes business, shifting the legislature to a budget-cutting posture. Adjustments and readjustments eventually restore the pluralist equilibrium.

To understand the agency's impotence more clearly, imagine that the agency simply sets efficiency equal to 1 and keeps it there. It would then be doing its utmost to reach its ideal point, but its utmost is inadequate. The legislature controls the budget, and legislators will be induced by the parity of interest-group campaign contributions to cut the budget until enforcement equals the pluralist equilibrium.

The asymmetrical impacts of proconsumer and probusiness policy preferences result from the distinctive manner in which the legislative and bureaucratic control mechanisms interact. The legislature can force a budget on a probusiness agency, but it cannot force the agency to spend all (or any) of it on enforcement, for the choice of efficiency is up to the agency alone. Thus, there is a constant struggle giving rise to the compromise region. On the other hand, the legislature can respond to the expansionist designs of a proconsumer agency by cutting its budget, and the agency can only fight back by increasing its efficiency. Once the agency reaches perfect efficiency, budget cuts are fully translated into declines in enforcement, and the agency can do nothing about it. Thus, the pluralist equilibrium.

It is a bit strong to say that a proconsumer bent has *no* influence: agencies that try to drive enforcement beyond the pluralist equilibrium become highly efficient. This result has a curious implication for antiregulation interests: if these interests faced a choice between a proconsumer agency and a budget seeker disciplined by oversight, they should prefer the zealous bureau. The explanation for this counterintuitive preference is that, owing to the power configurations of the bureau's environment, both agencies will wind up producing the same amount of enforcement, but the proconsumer agency will cost less. Thus a realistic assessment of power dynamics should make antiregulation forces look more kindly upon

career bureaucrats who singlemindedly support their opponents.²⁰

Overview: Basic Types of Solutions of the Model without Taxes

Although we have thus far analyzed only seven versions, the reader has been introduced to all the basic types of solutions attained by our much-more-extensive experimental manipulation of the system. Despite the combinatorial complexity of alternative goals, adaptation rates, congressional instruments, and initial values, all the versions reach one of three outcomes: the pluralist equilibrium, the degenerate solution, or the compromise region.²¹

Analyzing the versions within each type strengthens the inferences made on the basis of pairwise comparisons. Consider first the degenerate versions. With only one exception, all runs producing the degenerate solution were characterized by three properties: the agencies like slack, they adapt quickly, and Congress lacks oversight. When the three properties are present, the agency's policy preferences are irrelevant, as is budget-seeking.²² Normatively, the combination of a slack-seeking bureau and a legislature that

²⁰In real bureaucracies, policy goals usually comingle with less lofty objectives. What happens when the agency seeks other goals—budgets, slack, and avoiding oversight—in addition to policy? The three familiar patterns are recreated. Adding slack to the agency's motivation recreates the exploding pattern of Version 2: huge budgets, abysmal inefficiency, little enforcement. Adding budgets is less damaging. The budget seeking probusiness agency merely cycles closer to its ideal point, and the proconsumer agency is basically unaffected. Thus, slackseeking and budgetseeking differ more than is usually believed. Because revenues are allocated either to slack or output, a preference for more perquisites directly impairs efficiency. Budgetseeking lacks this directly corrosive effect. Finally, applying oversight to a slackseeking proconsumer bureau restores the pluralist equilibrium; overseeing a slackseeking probusiness agency reproduces cycling in the compromise region. (The authors will provide a complete table of results upon request.)

²¹Although one cannot prove that a simulation model will remain confined to a given pattern for all parametric combinations, extensive experimentation makes us confident that realistic parametric settings result in one of the three solutions.

²²We should note that these statements pertain to "on-off" versions of policy preference, that is, certain parameters are set equal to 1 (on) or 0 (off). One can increase the relative importance of any goal by increasing a parameter's value from 1 to, say, 3. We experimented with these possibilities, and found that no new types of solutions appeared. Sometimes, however, a version would switch from one type to another.

does not monitor is unfortunate. Empirically, we believe the combination to be rather rare.

All of the versions attaining a compromise solution share one property: a probusiness policy preference. No other attribute much mattered. If the agency lacked a business orientation, no combination of the other parameters would drive it into the compromise region; if it had it, there were few combinations that would move it out of the region.

Four patterns characterize the pluralist equilibrium category. First, most proconsumer agencies reached the pluralist equilibrium. The only ones failing to do so liked slack and were not overseen. Second, of the bureaus indifferent to policy, all those disciplined by the crude oversight procedure reached the pluralist equilibrium. Third, the rate of adaptation matters. Surprisingly, fast learning is not always a blessing, nor slow learning always a curse.²³ Several degenerate versions are transformed into pluralist equilibrium solutions when they adapt slowly rather than quickly. A slow adjustment rate means that Congress is increasing the budget faster than the bureau is decreasing productivity. The net result is increased enforcement, indicating that the legislature can push on a string—if it pushes fast enough and is willing to pay the price.

Fourth, although all the pluralist equilibrium versions produced the same amount of regulation, they did so with varying degrees of efficiency. Of those that were most efficient, not one both sought slack and did so with impunity. There were some bureaus that liked slack, but in all these cases the legislature monitored it. Finally, many more proconsumer than probusiness agencies are very efficient pluralist equilibrium versions.

The Final Step: Introducing Taxes

It has been convenient thus far to assume that firms and consumers ignore taxes, treating governmental policy as free goods or bads generated by the bureau. There are empirical grounds for using this as a starting point. The budgets of regulatory agencies are typically small compared to either the national budget or the social costs and benefits of regulatory policy. In such contexts a model without taxes is a reasonable approximation. There are also important analytical grounds. This approach has allowed us to develop a theoretical foundation that highlights the basic forces

of the system—forces that, it turns out, are constrained and masked by the role that taxation ultimately plays. By waiting until the final step to introduce taxation, we can better understand the more general model and the distinctive contributions of its parts.

Taxes might be added to the model in different ways. To keep matters simple, we assume that all budgets are financed by taxes and that all taxes are paid by consumers. Little is lost by assuming that firms escape taxation, since they already oppose regulation anyway. The interesting effects are on consumers.


Formally, these effects derive from a change in how consumers calculate the benefits of enforcement: we now assume that the value consumers ascribe to regulation equals the benefits of regulation minus its tax costs, where taxes just cover the agency's budget. What they receive they must now pay for. As a result, how much consumers value any level of enforcement depends on the budget required to produce it, and therefore on agency efficiency. As efficiency increases, the required budget decreases, taxes decrease, and consumers gain. More generally, the consumer benefit function shifts every period as the agency chooses new efficiency levels. It has become an endogenous component of the model.

This has two far-reaching consequences. First, because the consumer resource function is anchored in the consumer value function, it no longer describes a fixed relation between enforcement levels and resource contributions. It, too, is dynamic. As agency efficiency increases, the resource function shifts upward, reflecting greater consumer value and support of agency efforts at each enforcement level. Second, because the consumer resource function is now dynamic, there is no longer a unique point at which the consumer and business resource functions intersect. That is, there is no longer a single pluralist equilibrium. Instead, as illustrated in Figure 3, shifts in the consumer resource function define an infinite number of pluralist equilibria, each representing a potential balance of power for the system. Thus, although the meaning of the pluralist equilibrium remains unchanged, it is no longer an exogenously determined point to which the system gravitates. The pluralist equilibrium itself has become endogenous.

What can we say about political behavior in this new world of taxpaying consumers? Following the same logical steps as before, computer simulation indicates that this is one of those fortunate cases in which adding complexity creates simpler, more unified patterns of behavior. The salient transformation is that the degenerate solution and the compromise region of our earlier model now vanish—with rare exceptions, all system outcomes

²³The seemingly paradoxical benefits of slow learning have been noted before: see, for example, Lave and March (1975).

Figure 3. Pluralist Equilibria in the Tax Model



Cost and Benefit Functions (not graphed)

Costs = $.023E^2$
 Benefits = $20E - .04E^2 - \text{budget}$

- E = Enforcement
- F = Efficiency
- E^* = Social Optimum
- R_f = Political Resources of Firms
- R_c = Political Resources of Consumers

Resource Functions

Firms $R_f = 10 + .20 (.023E^2)$
 $= 10 + .20$ (costs)
 Con. $R_c = 50 + .05 (20E - .04E^2 - \text{budget})$
 $= 50 + .05$ (benefits)

are pluralist equilibria.²⁴ What we have is a very stable balance-of-power system.²⁵

²⁴Of course if the slopes and intercepts of the resource functions were significantly altered, the system would not necessarily reach a pluralist equilibrium. With resources continuing to be linear functions of convex costs and concave benefits, there are two cases to consider. First, if the firms' resources exceed the consumers' at all enforcement levels, then the system would obviously not attain a pluralist equilibrium. Instead, no enforcement would be produced. Second, if firms out-mobilize consumers at very low and very high enforcement levels, but the consumers out-mobilize firms at intermediate quantities, there are two locally stable out-

comes: the interior solution of the pluralist equilibrium and the corner solution of zero enforcement. In this case the more diffuse group has to overcome a mobilization threshold in order to sustain the program.

²⁵The exceptions are odd cases of persistent instability. One is a proconsumer agency that values slack, which fluctuates indefinitely. Experimentation with this case shows that if policy is weighted a bit more or slack a bit less in agency utility, equilibrium is reached. Thus, here and elsewhere, instability is the exception, not the rule. Moreover, in contrast to the nontax model's

To see why the degenerate solution disappears, consider an agency that seeks both slack and budgets. In the taxless model, this bureau decreased its efficiency relentlessly. Because it reduced enforcement, consumers were more mobilized than firms, inducing legislators to pump up the bureau's budget. Rising revenues taught the bureau that sloth pays. It therefore became still less efficient, output fell, legislators pumped in more money, and the cycle continued, leading to exploding budgets and zero efficiency. But in the tax model, consumers recognize the cost of the bureau's profligate ways: their support for the agency diminishes as efficiency declines, thus giving legislators less incentive to add budgetary fuel to the bureaucratic fire. Eventually, the alienation of support becomes so large that the bureau can no longer gain from reducing efficiency, and it reverses course, ultimately settling down—with the rest of the system—to a stable balance-of-power equilibrium.

More generally, the forces that earlier produced the degenerate solutions of zero efficiency, zero enforcement, and infinite budgets are constrained by the incentive-effects of taxation, which induce consumers and therefore legislators to put on the brakes. The result is hardly an absolute blessing for consumers, since enforcement and efficiency are low and budgets are much higher than necessary. But the system is in stable equilibrium, positive levels of enforcement are achieved, consumers do realize net benefits on the exchange, and all of this is far preferable to the degenerate solution.

Similar logic explains why the compromise region vanishes in a world with taxation. Recall that the compromise region arose because of a policy difference between a probusiness agency and the induced preference of the legislature, which is striving (in effect) toward the pluralist equilibrium. The agency, trying to reach its probusiness bliss point, reduces efficiency; the legislature responds by throwing money at it, which produces a see-saw contest in which enforcement cycles within a region bounded by the agency's bliss point and the pluralist equilibrium, accompanied by plummeting efficiency and exploding budgets. When consumers internalize taxes, however, the legislature finds itself at a disadvantage in its struggle with the agency. The legislature's crucial weapon is now treated as a cost by consumers, who will not long support throwing

money at the agency. Budgetary increases diminish the consumer resource function, shifting the pluralist equilibrium to the left. The compromise region therefore contracts—but, because the agency's bliss point is fixed, this contraction is entirely due to a movement of the pluralist equilibrium toward the agency ideal. The net effect of these new forces is to destroy the compromise region, along with the rising budgets and declining efficiency that upheld it. It is replaced by a stable equilibrium very near the agency's bliss point. For a probusiness agency motivated purely by policy, efficiency is very low (.07) and enforcement is virtually right at the agency's ideal (101). But all the relevant values are quite stable, and balance of power prevails in the environment.

As in the degenerate case, then, the forces creating the compromise region continue to operate: the struggle between the agency and the legislature is underpinned by an explosive potential for infinite budgets and zero efficiency. But this potential is constrained by the disciplining role of taxation, which saves the system from budgetary disaster—and, in the process, creates a context in which business firms are better able to get what they want from government.

Aside from the elimination of the degenerate solution and the compromise region, the tax model is virtually identical to our earlier model in its implications for the system's outcomes. Virtually all outcomes from the tax model are now pluralist equilibria, whose relative values clearly reflect the familiar operation of the same basic set of forces, that is, slack, oversight and so forth have the same effect as before.

Discussion

In this article we have proposed a general framework for pursuing a theory of bureaucratic politics. By design, it incorporates a range of components that more conventional models of bureaucratic politics have simplified away, components that are fundamental to the dynamic interaction among bureaucrats, politicians, and interest groups. In surface respects it cannot avoid being highly complex, and it may appear for that reason to fall into the trap that models of the neo-classical variety have tried to avoid. Yet, as our analysis has shown, its essential property is not complexity but simplicity—and it imposes coherence on bureaucratic politics by drawing upon and integrating the complexity inherent in the substantive context, not by assuming important aspects are irrelevant. Among its most general implications are the following.

1) Under a wide range of conditions, the regulatory system tends toward the pluralist equilib-

degenerate solution and compromise region, the unstable cases in the tax model are not explosive: all variables cycle within reasonable bounds. Even its instabilities are stable by comparison.

rium.²⁶ The logic of adjustment is very much a reflection of traditional pluralist thought. As Truman (1951) argued long ago, mobilization begets countermobilization, and changes in the underlying balance of power among social groups are translated by political institutions into corresponding changes in policy outcomes.²⁷ Due to this process of negative feedback, success in this kind of policy sector is self-limiting.²⁸ As we have

²⁶Is the pluralist equilibrium the same solution that completely rational decision makers would reach? This is an important question, both normatively and empirically, but it has no easy answer. Part of the problem is that there are many ways to specify an optimization model. Analytical choices must be made regarding, for instance, the information available to each actor, how they formulate beliefs including expectations about the future, the nature of their decision rules, for example, maximize expected utility or minimax loss, and how the actors themselves are conceptualized, for example, are the interest groups unitary rational actors, coalitions, or sectors of decentralized decision makers? Obviously, different choices along these dimensions could easily produce different systemic outcomes, so we cannot say in general what a shift from adaptation to optimization might imply. We can, however, offer two specific conjectures. First, if the situation were modelled as a cooperative game—that is, actors can costlessly make binding agreements—with the interest groups taken as rational actors, then we conjecture that if an equilibrium exists, it is to produce the socially optimal amount of enforcement (Coase's theorem). Because in general the pluralist equilibrium is not equal to the social optimum (see text, below), the adaptive model does not converge to the equilibrium attained by these rational agents. Second, if the situation were modelled as a non-cooperative game among perfectly informed legislators, with the interest groups and the agency taken as passive reaction functions rather than as strategic agents, then we conjecture that the pluralist equilibrium is the Nash equilibrium of the game. (We would like to thank an *APSR* referee for suggesting this conjecture.) More generally, however, computing the equilibria reached by optimizers in a multiperiod world of the sort represented here would be very difficult—indeed, a fullscale project in its own right.

²⁷Here, as in virtually all dynamic models, the parameters are held constant in order to investigate the system's long-run tendencies. Because empirically these coefficients may change, one may never observe a regulatory system in equilibrium. Instead, it may always be moving toward equilibrium.

²⁸We must point out that we have assumed that the regulatory policy system is not so important that losing is debilitating. In such a world, the positive feedback of the Matthew effect—"To him who hath shall be given"—creates an unstable system of cumulative advantages (Dahl, 1971). The mobilization of previously unorganized interests such as minorities and environmentalists indicates that the American system is described more by negative than by positive feedback (Landau, 1973).

shown, moreover, this result is entirely consistent with—indeed, to be expected from—rationally adaptive behavior, and it holds even though bureaucratic and legislative participants play distinctive, self-interested roles in creating policy from social inputs. The common notion that rationality and the translation effects of institutions are somehow inconsistent with pluralism, then, is overdrawn. Although more elaborate versions of our model (incorporating, for example, legislative committees) may point to conditions that imply alternative expectations, there is every reason at this stage to stress the contribution that pluralist ideas can make to our understanding of bureaucratic politics, as well as their compatibility with more "modern" lines of analysis.

2) This equilibrium is reached in a groping fashion by decision makers who only dimly understand the effects of their own actions. Thus, it is highly consistent with the descriptive literature on policymaking, which emphasizes the complexity of the policymaking environment, the frequent missteps, the unintended consequences, as well as the "disjointed incrementalism" of policymaking in a polyarchical system (Dahl, 1971; Lindblom, 1965). Indeed, in many ways the model is a formalization of the "intelligence of democracy." With Lindblom, we attribute the intelligence (what there is) of policy outcomes to properties of the system rather than to the brilliance of individuals. Our actors are not brilliant. Like the decisionmakers in the "Science of Muddling Through" (Lindblom, 1958), our actors ignore side effects, consider only a limited range of alternatives, and adapt as best they can to a complex environment. Through the feedback and mutual adjustment occasioned by their interaction, they tend collectively to produce coherent, sensible outcomes that the individuals themselves neither intend nor have the power to bring about.²⁹

3) The system is not socially optimal—but the direction of suboptimality differs from what the critics of pluralism have claimed. Since Olson's (1965) work first appeared, it has been common to emphasize the advantages of compact groups over diffuse groups and to explain the suboptimality of policy outputs via this disparity. In our model, this logic would imply a definite bias in favor of business and against consumers, and thus an equilibrium level of enforcement lower than the social optimum. Yet this argument is not in general true, and for the specific model of this

²⁹Indeed, it is intriguing to note that this system of myopic and adaptive decisionmakers may be converging to the same outcome that would be attained by perfectly informed and perfectly rational legislators engaged in a noncooperative game, as conjectured in note 26.

article it is quite false. The reason turns on the difference between the pluralist equilibrium and the social optimum. Enforcement generates costs and benefits for society, and its level is socially optimal when net benefits are maximized, that is, when the marginal benefits and marginal costs of enforcement are equal. The pluralist equilibrium occurs, however, at the enforcement level for which the political resources of consumers and firms are equal. There are accordingly two crucial dimensions of difference between the two solution concepts. First, the social optimum is based on the cost and benefit curves themselves, whereas the pluralist equilibrium is based on the political resource curves. Second, the social optimum is based on marginal curves, the pluralist equilibrium on total curves. The combination of these factors implies that the pluralist equilibrium may be above or below the social optimum.³⁰ Since the Pluralist Equilibrium is the "natural" policy outcome of our system, this means that regulatory enforcement may often be above the social optimum, even though consumers are at a distinct organizational disadvantage relative to business. In the model presented here, this is precisely what happens: business is assumed to be four times more effective at translating values into resources, but the pluralist equilibrium usually lies above the social optimum anyway.

4) The power of bureaus to get what they want has been exaggerated. Although bureaus can sometimes move to a paradise of exploding budgets and slack, they cannot do so when checked by even a primitive form of legislative oversight, nor when their adjustment is slowed by inertia, nor when their constituents are aware of the burden of taxes. In the most general case they wind up at the pluralist equilibrium. There, budgets and slack may vary depending upon the time path of the process, but the agency has clearly been "led" by the system's logic of adjustment and has not engineered its preferred bureaucratic outcome. When the agency has policy preferences, its influence is more apparent—but it is asymmetrical, and it only works to produce smaller levels of enforcement than critics lead us to expect. Specifically, a probusiness agency can move the system toward its own ideal point, increasing its budget and slack in the bargain. A proconsumer (and therefore expansionist) agency,

on the other hand, is helpless to prevent movement to the pluralist equilibrium.

At a more aggregate level, these conclusions provide a different perspective on the growth of government. In the formal modelling literature on this matter, we have seen the blame for big government assigned to different institutions. First, the bureaucracy did it (Niskanen, 1971). Next, Congress did it (Fiorina, 1977). Now, we argue that the routine functioning of a pluralist system is responsible for the size of government. Outcomes are collective by-products of the decentralized choices of several institutions, not the unique responsibility of any single one. And there is no uniform bias toward bigness—government may often be too small rather than too large.

5) The legislature and, indirectly, consumer groups face a fundamental dilemma in the use of control mechanisms. By cutting budgets, the legislature can effectively prevent an agency from producing "too much" enforcement. But it does not work the other way round: a probusiness agency producing "too little" can absorb a larger budget by increasing inefficiency. Moreover, because the budget is an incentive in itself as well as a means of generating output, a policy-indifferent agency may view increased budgets as rewards for inefficiency and low production. If the agency's constituents are not cost-conscious, this view of budgets will prompt the bureau to move toward the degenerate solution—just what consumers seek to avoid. Fortunately, even if constituents ignore taxes, the legislature has a second mechanism—oversight—which plays a crucial role in saving the system from explosive budgets and inefficiency. The wonder of oversight is that it requires so little of the legislature and nonetheless works so well in reversing systemic pathologies.

6) Bureaucratic inertia can be good for the system. Although this seems odd, given the bad reputation inertia has acquired, it actually makes sense: for inertia constrains the flexibility of bureaus in their pursuit of ever-larger budgets and slack, inhibiting their adaptation.

7) The implications of the model pose distinct problems for liberal supporters and conservative opponents of regulation. For liberals, the dilemma arises because the peculiar asymmetry of the budget mechanism operates to their disadvantage. A conservative majority can achieve lower enforcement quite effectively by cutting the budget, but a liberal majority is far less able to engineer higher enforcement by increasing the budget, and its efforts can easily backfire by rewarding agencies for poor performance. Conservative opponents face a different dilemma: a probusiness agency is able to triumph over a proconsumer legislature, but only at the cost of gross inefficiency. They are thus torn between the com-

³⁰To clarify further, that the social optimum derives from marginal cost-benefit curves whereas the pluralist equilibrium derives from total resource curves tends to make the latter solution exceed the former—yet the greater the relative advantage of business in mobilizing resources, the farther to the left the two resource curves will intersect, and the smaller the pluralist equilibrium will be. Either force may predominate.

peting desires for less regulation and more efficient government.

Conclusion

This model is just a beginning, for we purposely ignore institutional and strategic aspects of policy-making that must eventually be part of a satisfactory theory. The task for the future is to move toward more elaborate models that take these additional aspects systematically into account. In our view, one of the real advantages of the general framework we have developed here is that it easily accommodates efforts to move in a variety of interesting directions, thus encouraging the proliferation of a family of models whose similarities and differences may well provide important new insights into bureaucratic politics. The following are but a few of many elaborations that seem to us both feasible and promising.

Legislative Committees. Legislatures can be assumed to vote only on alternatives generated by committees, which are therefore in a position to use their agenda control to shape outcomes for the legislature and, less directly, for the system as a whole.

Electoral Districts. Legislators can be assumed to be elected from geographically distinct districts, with the distribution of group costs and benefits and the ensuing balance of group resources varying from district to district.

The President. The role of the president can be explored not only by means of the appointment mechanism but also by introducing an OMB with budgetary and monitoring powers and perhaps by incorporating the president's veto power.

Interest Group Strategy. Groups can be allowed to adopt various strategies in their efforts to influence outcomes, for example, contributing only to supporters (not against opponents), basing all contributions on the distance of the legislators from the median, or rewarding opponents if they move in the right direction.

Bureaucratic Innovation. Agencies can be allowed to devote a portion of their budgets to the generation of new programs. These innovations would then alter the production of enforcement or its consequences for business and consumers.³¹

Modes of Adaptation. Decision makers can be assumed to adapt in more sophisticated ways to their environments. Among other things, this may involve probabilistic adaptation or the development and modification of belief structures.

Whatever elaborations may be introduced, they will be unified by a common framework, a macro-

theory of politics based on the microfoundations of bounded rationality. This theory, particularly its formal representations, stands to provide new insights into the nature of bureaucratic politics by placing Simon's "administrative man" (1947) squarely in the context of the larger political system. It also stands to illuminate longstanding issues concerning the intelligence of democracy. For decades political scientists have wondered whether voters are too ill-informed, or politicians insufficiently rational, for the overall health of polyarchies. The classical response of pluralist thinkers to these concerns is that, although we cannot depend upon individual decision makers, we can rely upon properly designed systems that "pit ambition against ambition" and correct for the myopia and ignorance of its constituent elements. Our framework provides a formal means for systematically investigating these issues.

References

- Aberbach, J., & Rockman, B. Clashing beliefs within the executive branch: The Nixon administration bureaucracy. *American Political Science Review*, 1976, 70, 456-468.
- Axelrod, R. *Structure of decision: The cognitive maps of political elites*. Princeton, N.J.: Princeton University Press, 1976.
- Bernstein, M. *Regulating business by independent commission*. Princeton, N.J.: Princeton University Press, 1955.
- Braybrooke, D., & Lindblom, C. *A strategy of decision*. New York: Free Press, 1963.
- Cater, D. *Power in Washington*. New York: Random House, 1964.
- Cohen, M. The power of parallel thinking. *Journal of Economic Behavior and Organization*, 1981, 2, 285-306.
- Cohen, M. Conflict and complexity: Goal diversity and organizational search effectiveness. *American Political Science Review*, 1984, 78, 435-451.
- Cohen, M., & Axelrod, R. Coping with complexity: The adaptive value of changing utilities. *American Economic Review*, 1984, 74, 30-42.
- Cohen, M., March, J., & Olsen, J. A garbage can model of organizational choice. *Administrative Science Quarterly*, 1972, 17, 1-25.
- Crecine, J. P. *Governmental problem-solving*. Chicago: Rand McNally, 1969.
- Cyert, R., & March, J. G. *A behavioral theory of the firm*. Englewood Cliffs, N.J.: Prentice-Hall, 1963.
- Dahl, R. *Polyarchy*. New Haven, Conn.: Yale University Press, 1971.
- Fiorina, M. *Representatives, roll calls, and constituencies*. Lexington, Mass.: D.C. Heath, 1974.
- Fiorina, M., & Noll, R. Voters, bureaucrats, and legislators: A rational choice perspective on the growth of bureaucracy. *Journal of Public Economics*, 1978, 9, 239-253.
- Freeman, J. L. *The political process: Executive bureau-legislative committee relations*. New York: Random House, 1955.

³¹For work that moves in this direction, see Levinthal and March (1982).

- Huntington, S. The Marasmus of the ICC: The commission, the railroads, and the public interest. *Yale Law Journal*, 1952, 614, 467-509.
- Landau, M. Federalism, redundancy, and system reliability. *Publius*, 1973, 3, 173-196.
- Lave, C., & March, J. *Introduction to models in the social sciences*. New York: Harper & Row, 1975.
- Levinthal, D., & March, J. A model of adaptive search. *Journal of Economic Behavior and Organization*, 1982, 2, 307-333.
- Lindblom, C. The science of "muddling through." *Public Administration Review*, 1959, 10, 79-88.
- Lindblom, C. *The intelligence of democracy*. New York: Free Press, 1965.
- Long, N. Power and administration. *Public Administration Review*, 1949, 9, 257-264.
- Lowi, T. *The end of liberalism*. New York: Norton, 1969.
- March, J. G., & Simon, H. *Organizations*. New York: Wiley, 1958.
- Mayhew, D. *Congress: The electoral connection*. New Haven, Conn.: Yale University Press, 1974.
- McCubbins, M., & Schwartz, T. Congressional oversight overlooked: Police patrols versus fire alarms. *American Journal of Political Science*, 1984, 28, 165-179.
- Migue, J., & Belanger, G. Toward a general theory of managerial discretion. *Public Choice*, 1974, 17, 27-43.
- Miller, G., & Moe, T. M. Bureaucrats, legislators and the size of government. *American Political Science Review*, 1983, 77, 297-322.
- Newell, A., & Simon, H. *Human problem solving*. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
- Niskanen, W. *Representative bureaucracy*. New York: Aldine-Atherton, 1971.
- Niskanen, W. Bureaucrats and politicians. *Journal of Law and Economics*, 1975, 18, 617-644.
- Olson, M. *The logic of collective action*. Cambridge, Mass.: Harvard University Press, 1965.
- Padgett, J. Bounded rationality in budgetary research. *American Political Science Review*, 1980, 74, 354-372.
- Peltzman, S. Toward a more general theory of regulation. *Journal of Law and Economics*, 1976, 19, 211-240.
- Posner, R. Theories of economic regulation. *Bell Journal of Economics and Management Science*, 1974, 52, 335-358.
- Quirk, P. *Industry influence in federal regulatory agencies*. Princeton, N.J.: Princeton University Press, 1981.
- Rich, E. *Artificial intelligence*. New York: McGraw-Hill, 1983.
- Shepsle, K., & Weingast, B. Political preferences for the pork barrel: A generalization. *American Journal of Political Science*, 1981, 25, 96-111.
- Simon, H. *Administrative behavior*. New York: Macmillan, 1947.
- Simon, H., Smithburg, D., & Thompson, V. *Public administration*. New York: Knopf, 1950.
- Stigler, G. The theory of economic regulation. *Bell Journal of Economics*, 1971, 2, 3-21.
- Weingast, B. A principal-agent perspective on congressional-bureaucratic relations. Presented at the Fifth Carnegie Conference on Political Economy, Carnegie-Mellon University, June, 1983.
- Wilson, J. *The politics of regulation*. New York: Basic Books, 1980.