

Chapter 11

Policy Punctuations: Cascades, Sieves, and Institutional Friction

We closed Chapter 10 with a picture. The picture is a way of looking at all budget changes in OMB subfunctions since WWII. We prepared a simple frequency distribution of budget changes by ordering the changes from the biggest cutback to largest increase, and counted the number of changes for each interval. We showed how this frequency distribution could be seen as a combination of time series (each subfunction traced across the post-war period, added up with all the other subfunctions). It can also be seen as a combination of each budget change for each subfunction for a single year, added up.

What affects budgets must also affect the policy process more generally. In this chapter, we lay the groundwork for comparing policy distributions by developing some concepts that will help us generalize this budget picture to the entire political system. First we review some basics of general policy change models. Then we turn to a discussion of the three major known mechanisms in social science for generating the leptokurtic distributions of Figure 10.13 — policy cascades, decisional sieves, and institutional friction. We show how each of these mechanisms can lead to leptokurtic change distributions. In the next chapter, we compare a variety of policy distributions in a manner that allows the combination of the static analyses common in the analysis of political institutions with the dynamics of policy choice.

Policy Change

We want to understand the processes that are responsible for patterns of change in policies—basically the solutions directed at perceived problems. We are less interested in explaining particular events and policy responses to these events. The prediction of events and

responses to them is probably a fool's errand, at least at this stage of development in the policy sciences. But this does not leave us helpless, and here we develop a set of new perspectives that can address this complexity without the baggage of demanding point predictions—that is, the prediction of particular events and policy responses to them.

Patterns of change can be described as time series, in which we look for major shifts in the variables of interest, or as change distributions. A *time series* simply follows a variable of interest (like a category of expenditure) through time. We have looked at many of these series in the previous chapters of this book. A *change distribution* takes the period-to-period change (year-to-year change for government expenditures, for example) and studies them without respect to when they occurred. If we look at a picture of a time series, we are inevitably drawn to the study of particular changes, and we tend to think about what caused the specific changes. If we look at a picture of a change distribution, we study the pattern rather than the particulars of changes over time—how many large, strong punctuations occur and how many simple adjustments occur.

We define a policy change as a period-to period alteration in the commitment of a government to an objective. This may involve a budgetary commitment, or a commitment of the legitimacy of government through a statutory enactment, or the commitment of organizational capacity in implementing a policy. But, as we have repeatedly emphasized in this book, it can also involve an investment of attention by policymakers to an objective. Attention is a scarce good, and its allocation toward a policy objective is an important indicator of policy commitment.

In studies of the policy process, the incremental model of policy change serves as the standard. Whatever was being done last year is continued this year, with minor modifications.

Of course many major policy initiatives do not fit the incremental standard. They change the course of policy fundamentally. But the incremental model is a serviceable starting-point for any study of policy change.

A major problem is judging just how large an increment must be to count as a major policy change. Any particular policy change in the eye of one beholder can be trivial; in the eye of another, huge. In the absence of an agreed- upon standard, the question of policy change versus continued policy stability is unanswerable. As a consequence, we adopt a distributional perspective here; any change is judged relative to the overall pattern of policy stability and change within an issue area or policy arena.

Two distributions of policy commitment are important. The first is the *policy outcome distribution*. An outcome distribution is a frequency distribution of the levels of effort expended in a variety of policy arenas. For example, it could be the budgetary amount allocated to a policy arena on an annual basis, adjusted for inflation to assess real commitment. The outcome distribution is a static representation of the state of policy commitments at a given time. As a consequence, it fails to distinguish the initiation of past commitments from present changes. Far too many studies of comparative public policy rely in essence on the outcome distribution, postulating all sorts of causal mechanisms to account for a static outcome distribution in one nation versus another. A few studies, notably that by Peter Hall (1992), focus on policy changes in response to similar inputs from the environment; this approach is far more successful.

More essential for our project is the *policy outcome change distribution*. An outcome change distribution is a frequency distribution of change patterns of policy outcomes. It is a distribution of period-to-period changes in a variable of interest—for example, a tabulation of the sizes of year-to-year changes in budget allocations to a particular policy area. Policy changes

may be tabulated by first differences that are raw period-to-period changes, or by annual percentage changes. The latter is especially useful when aggregating change processes across policy arenas. Focusing on first differences in such circumstances will lead to over-weighting the large-budget arenas.

In the study of policy outcome distributions, the Normal, or familiar bell-curve, distribution becomes a strong standard for comparison. It is an important comparison for policy studies because it is the result of an incremental policy process. An incremental policy process is one that is approximately at equilibrium. Interests vie for advantage within a generally-understood framework that precludes the overwhelming victory of one set of interests over another (Redford 1969). The decision-making process in such a policy arrangement is characterized by limited comparisons (generally proposed changes are judged in comparison to the status-quo) and mutual adjustments among interests. The resulting outcomes of the policy process are increments (or decrements) when compared to the existing policy.

Considerable debate, much of it based on misunderstandings, has characterized the notion of incrementalism. Much of the debate has centered on 'how large' a policy change would need to be to qualify as 'non-incremental'. That argument can be avoided by using a distributional perspective. A major advantage of the distributional perspective is that it leads automatically to a model for incremental policy-making that can be compared to actual policy outcomes. If many small factors impinge on the course of public policy, the result will be an incremental pattern of change often characterized by policy scholars as 'policy drift'.

This is the typical situation when Congress assigns policy responsibility to the executive bureaucracy and its own oversight committees. Interest groups are attracted to the new center of policymaking authority. As the law is implemented within the general confines of the statutory

requirement, interest groups reach an equilibrium (Redford 1969), even if some interests are benefited at the expense of others in the initial statute. As a consequence, subsystem politics is incremental politics.

Analyzing Incremental Change

Equation (1) is the traditional equation for incrementalism . Political forces are balanced and the broader public is inattentive. Policy today, P_t , is just policy last year, P_{t-1} , plus an adjustment component, ϵ_t .

The adjustment component, ϵ_t , is important. It makes the assumption that the adjustment to last year's policy is made up of lots of more or less independent factors. Some of these factors that may influence increases in the policy, others decreases, but many of them just cancel each other out.¹

$$P_t = P_{t-1} + \epsilon_t \quad (11.1a)$$

$$(P_t - P_{t-1}) = \epsilon_t \quad (11.1b)$$

Equation 1a describes a random walk over time. Statistically, incrementalism in which policy participants within a sector are at equilibrium will trace a *random walk* through time. A random walk is a time series in which the next value is completely random.² The first differences of a random walk (that is, the value of the variable at time 2 minus the value at time 1), considered together, form a Normal distribution.

¹ We add the assumption that $\epsilon_t \sim \text{IID}(0, \sigma^2)$. This just means that the adjustment, or 'error', terms is drawn from similar distributions with 0 mean (there are no systematic factors affecting policy change) and similar variance. If many factors influence each error term, then the central limit theorem proves that, in the limit, the distribution would be Normal.

² In the sense of being drawn from a probability distribution of mean 0 and finite variance. The classic illustration is the case of counting the number of heads in a number of tosses of a fair coin. Will the next coin be heads? Taking the number of heads that occurred previously each time we toss a coin is a random walk through time.

Equation 11.1b shows that the year-to-year change in budgets is affected only by the random factor. Incremental policymaking generates a Normal distribution of policy change (Padgett 1980). Last year's policy is modified by several independent factors, many of which off-set one another. If we subtract last year's change from this year's change, we get an expected value of zero. There is no expected direction of change.

An important modification of the incrementalist model we term 'incrementalism with upward drift'. Government budgets in most developed democracies have moved upward since the Second World War, a pay-off from successful economic management. In such an environment, incrementalist politics of mutual partisan adjustment and successive limited comparisons (Lindblom 1960) is played out within a growing pie. Moreover, many programs are funded by formulas that include the size of the target population or some sort poverty level floor. These aspects of program funding can result in growth of budgets as the economy grows.

This suggests that the year-to-year upward drift is proportional to last-year's budget. In economics the starting-point for a similar problem, the growth of firms over time, is similar to incrementalism in politics. The Gibrat thesis asserts that the growth of firms over time is proportional to the size of the firm. Incrementalism with upward drift would imply a similar thesis with regard to government budgets.

If the stronger form of upward drift, the Gibrat thesis, were applied to government budgets, it would imply that the annual percentage change (rather than the annual dollar change) would be constant—up to the random component. So in this formulation, the annual percentage change in budgetary commitments would follow a random walk, and the annual proportional (or percentage) difference would follow a Normal distribution:

$$P_t = (1 + k)P_{t-1} + \varepsilon_t \quad (11.2a)$$

$$\begin{aligned} P_t &= P_{t-1} + kP_{t-1} + \varepsilon_t \\ (P_t - P_{t-1})/P_{t-1} &= k\varepsilon_t \end{aligned} \quad (11.2b)$$

where k is the proportional growth increment.

It is critical to understand that a straightforward incremental policy process will invariably lead to an outcome change distribution that is Normal. And vice-versa: any Normal distribution of policy outcome changes must have been generated by an incremental policy process. Any time we observe any non-Normal distribution of policy change, we must conclude that incrementalism cannot alone be responsible for policy change. That is why distributional analyses are so critical to policy studies.

As we shall see, policy change distributions are almost never characterized by a Normal distribution. A simple model of incremental decision-making can be quickly ruled out. If we study lots of policy change distributions, we find a particular pattern of non-Normality. The distributions are leptokurtic. This means they have fat tails and tall, slender central peaks in comparison to the Normal. In comparison to the Normal, there is simultaneously too much change (the fat tails represent large changes in a policy change distribution) and too little change (the strong central peaks indicate little change from the previous policy commitment). What is absent in such distributions is moderate change—the classic 'weak shoulders' of leptokurtic distributions.

The immediate result is a rejection of the incremental model of decision-making as a complete model of policy change. Yet the strong central peaks indicate that most of the time policy-making is indeed incremental. It is the extreme tails that disrupt the pattern.

Now we see theoretically what was evident to the eye in Chapter 10: leptokurtic policy change distributions rule out incrementalism as a basis for the study of policy change by putting

it in full context. Ignore the tails, and you get an incremental-style politics. Look at the tails, and you will immediately find the signature non-incremental pattern of episodic and disjoint change.

Cascades, Sieves, and Friction

Focusing on whole distributions of changes in policy commitments leads us to examine the causal processes that can generate such distributions. Simple additive processes such as those captured in the incrementalist models are precluded, since they generate Normal distributions (West and Deering 1995). Simple interventionist models in which an external event generates uniform responses from all policy process distributions are unlikely, as we demonstrated in Chapters 9 and 10. The major budget punctuations are not associated in a simple fashion with external events—rather, internal dynamics, such as the arguments made by presidents and party leaders—can amplify responses by attending to particular aspects of the flow of events. Moreover, as we noted in Chapter 8, each policy arena seems to need a separate causal explanation because of the complexities inherent in political choice. A distributional approach pushes us to a level of analysis that transcends these particulars, but of course at the loss of important detail.

In the case of the fat-tailed distributions characteristic of policy changes, three basic causal processes are implicated: *cascades, sieves, and friction*.³ We may understand these processes as ways of explaining the extreme values that occur in the tails of policy distributions (Sornette 2000). Cascades are processes in which positive feedback dominates—each change begets another even larger change. These large ‘correlated’ changes add cases to the tails of a frequency distribution of policy change like that depicted in Figure 10.13. Institutional friction

³Thanks to John Padgett and John Brehm for highlighting this distinction in the Organizational Theory and State Building Workshop at the University of Chicago.

occurs when institutions retard change, but result in a large-scale ‘jump’ when the built-in friction is overcome. The large changes from such policy leaps similarly add cases to the tails of the policy change distribution. Sieve processes come about when decision-makers apply ever-greater constraints to a decision-making process. Simple constraints rule out more options than more severe ones, causing ever-greater changes from the status quo and resulting in leptokurtic policy change distributions.

These three processes may all operate simultaneously, or they may not. It may turn out, for example, that cascades are the key to understanding how friction is overcome in policy systems. However they are clearly not sufficient—like fads, the pressures generated may just fade away. Decisional sieves may be decision-making reflections of cascades—but they may not. We just don’t know, but the isolation of these three stochastic mechanisms has the decided advantage of specifying complexity. Rather than claiming that everything is contingent, contextual, and complex, we can begin to try to reason through just what mechanisms are responsible for the observed complexity in any given case of policy change.

Here we will discuss briefly the roles of cascades and sieves in policy choices, then develop in considerable detail the friction model. It is critical because it has the potential of unifying static institutional analyses with the dynamics of public policy studies.

Cascades

Cascades explain the extreme values in a policy change distribution by a process in which one action begets other actions, independently of the nature of the external event stimulating the initial action. Cascades result from positive feedback effects (Baumgartner and Jones 2002). Positive feedback effects may be contrasted with more typical negative feedbacks, which occur when a process causes a dampening effect on itself (often indirectly). For example, in a well-

behaved market, scarcity of a good causes the price to rise. This encourages producers to make more of the good available, which exercises downward pressure on prices. As prices fall, producers earn less per unit, and shift some productive capacity somewhere else. Positive feedback occurs in markets when a price rise causes a further price rise. In financial markets, bubbles can occur in this circumstance. Crashes occur when a price drop causes further price drops—the self-correcting feature of negative feedback has been replaced by positive feedback. Positive feedback effects are very much in evidence in politics, as well—bandwagons in elections, for example, and the destruction of policy subsystems in the face of social mobilization (Baumgartner and Jones 1993).

In social science, a major source of cascades is the process of monitoring and mimicking. In many situations, people may observe carefully not the real world directly, but how others around them are responding to the real world. Then they take action based not on real world indicators but on the responses of others. This process results in chain-reaction dynamics, in which major changes can occur rapidly. Social cascades have been used to study elections and policy change in political science (Bartels 1988; Baumgartner and Jones 2002), fads in sociology (Bikhchandani, Hirshleifer and Welch 1992), peculiarities in pricing in economics (Becker 1991), residential segregation (Shelling 1972) bubbles and crashes in economics (Shiller 2001, Sornette 2003), revolutions and collective action (Granovetter 1975, Chong 1991, Kuran 1989) and a variety of other topics (Kirman 1993), all with considerable success.

These cascades cause punctuated behavior because (in a pure situation) either there is no change in behavior or everybody changes. In many situations, there can be different groups engaged in monitoring and mimicking, but they are not tightly connected to one another. For example, Democrats in Congress organize separately from Republicans, and may be sensitive to

different informational cues. This deconcentrated organizational structure limits the size of the cascade—one of the arguments that in effect the constitutional framers made on behalf of the proposed system of divided and separated powers.

While some have modeled cascades as fully rational informational shortcuts in decision-making (Hershleifer 1995), these attempts are unpersuasive (Jones 2001). In the case of financial markets, mimicking can result in severe and unpredictable crashes, resulting in a risk structure that should deter decision-makers capable of rational expectations from engaging in mimicking. Most people studying these processes base them on bounded rationality, in that people caught up in cascades are goal-oriented but make systematic mistakes (compared to a rational expectations model).

Imitation and herding are the bases of a class of models of financial markets that are designed to explain bubbles and crashes. Computer simulations of these processes have demonstrated how imitation by investors can lead to market crashes. Interestingly, the heart of the problem is in communication. In these models, if investors do not communicate with one another, then mimicking cannot occur, and crashes will not occur. More correctly, any crash will be a proportionate function of an external event, so that only major calamities could cause crashes.

But if they communicate, then crashes can occur, and the more widespread the communication, the more likely the crash. In such circumstances, we may observe a crash that seems not to be connected with external events. Some seemingly trivial happening can cause a massive sell-off, because the system has moved into a very fragile state. Summarizing this line of research, Didier Sornette (2003: 159) writes that the risk of a crash “increases dramatically when the interactions between traders becomes strong enough that the network of interactions

between traders self-organizes into a hierarchy containing a few large, spontaneously formed groups acting collectively”.

Actual financial data have also been studied using Markov switching models. These models simply treat financial time series, such as stock market returns, as being generated by one of two ‘regimes’—a normal regime, when things work according to standard market theory, and a ‘bubble’ regime, when positive feedbacks are in effect (Schaller and Van Norden 1997; Sornette 2003: 167-168). Similarly, political scientists have noted the ‘critical periods’ in which attention toward a policy matter causes a shift from the typical negative feedback processes characteristic of policy subsystems to the intense period of positive feedback associated with major reforms (Baumgartner and Jones 1993, 2002). Conceiving of a kind of Markov on-off switch for positive feedback systems has the potential of adding substantially to the theory and analysis of complex systems.

Sieves

A second way to explain extremes occurs when policy-makers try to solve problems presented by external challenges, and must sequentially search alternatives rather than examine all comprehensively (Padgett 1980). In complex and uncertain circumstances, decision-makers will winnow the realistic alternatives until they find one that will satisfy the conditions imposed by the external challenge. First they decide in which direction to move from the initial status quo point. Should a budget be cut or increased? Should we move toward more or less regulation of an industry? Then they cycle through the available options until one option satisfies the constraints of the situation. If the alternatives are ordered, as they are when deciding how much to increase or cut a budget, then as constraints get more difficult to satisfy, a decision-maker will have to move further away from the status quo. Simple constraints lead to simple adjustments

with few alternatives in the vicinity of the status quo rejected. But difficult constraints imply large movements from the status quo

In Padgett's serial choice model, constraints are generated from two considerations: whether program increases (decreases) are justified 'on the merits' and whether the increase (decrease) is justified under 'the current fiscal climate' (1980: 364). The second consideration in effect brings into a decision-making situation external and contextual considerations. So the serial choice model can be seen as a particular mechanism of disproportionate information processing. *Both* the conditions of merit and fiscal/political climate must be satisfied.

If the probability of deciding that the program deserves a specific budgetary increase is p_i , and the probability of deciding that the same alternative is acceptable under the current political climate is k , then the probability of accepting the particular alternative is $P(\text{accept}) = p_i \times k$. The 'i' subscript indicates the program; there is no subscript for k because the political climate is presumed to be constant across programs. This may well not be the case; Republicans in Congress may support military spending while insisting on domestic program decreases. But it is a good start.

If the decision-maker judges the merits of a program to be high, then the probability of an increase based on the merits would be positive. But if the fiscal climate were judged to be unfavorable, then multiplying $p_i \times k$ would yield zero (because k , the probability that the fiscal environment is favorable, is zero). No matter how meritorious the program, there will be no budget increase until the fiscal climate improves. With an improvement in the fiscal climate, for very meritorious programs, whose budgetary increase was zero in the past, suddenly receive increases.

While incremental budgeting leads to a Normal distribution of policy changes, sequential search leads to strongly punctuated outcome changes. Padgett derives the particular budget change distributions that theoretically should emerge from his analysis, and finds that the serial choice model leads to a double exponential distribution. The ‘double’ comes from the fact that budgets can be cut or raised, so there are positive and negative values for the curve. The two curves meet at the 0 point (that is, zero percent change in the budget) The exponential distribution has fatter tails and a more slender peak in comparison to the Normal—it is leptokurtic.

That is what would hold for a single program. But there are lots of programs, each of which has different characteristic parameters—different means and variances, for example. If we combine programs, as we did in Figure 10.13, the result is a double Paretian distribution, which is even more fat-tailed and slender peaked—leptokurtic—than the exponential. And, indeed, the budget change distribution in Figure 10.13 is Paretian.

Padgett’s serial choice model is a model of decision-making; the cascade model described above is a policy process model. Lots of decision-makers would interact in the cascade model—as for example, when all Members of Congress got concerned about corporate governance simultaneously in early 2002. Clearly these interactions can be transmitted down into bureaucratic decision-making as part of the political climate—modeling in effect the notion of ‘political feasibility’. No securities regulatory body could conceive of decreasing regulation in the face of the corporate scandal. The ‘political climate’ constraint leads to a probability of zero for decreases in regulations even though regulators may well have thought that desirable (indeed, the head of the SEC, responsible for regulating financial accounting, had spoken of bringing a more ‘accountant friendly’ environment to the agency in 2001).

Friction

The cascade model suggests that decision-makers are basically equal in their inputs on a situation.⁴ That is not true in politics. The president's actions have a higher probability of generating downstream changes than a rookie Member of Congress. Moreover, the locus within the policy process affects the general probability of generating favorable outcomes. It may be fairly easy to get Members of Congress to focus on a problem, and hold a hearing. But it may be very difficult to get a statute passed. We need, in effect, to impose an institutional structure on the random interactions of the cascade and sieve models.

The final way leptokurtic policy change distributions may emerge is through friction. If there is resistance to change—and this resistance can be political, institutional, or cognitive—then there can be a ‘jump’ in a policy commitment when this friction is overcome. Institutional friction is the most familiar process to political scientists, but it has tended to be thought of as static and often critiqued as ‘gridlock’. What is characterized as gridlock may be better characterized as intuitional friction: hard to overcome, but major shifts in policy may occur when it is overcome.

Let us turn to a real-world policy example. On August 28th, 1950, President Truman enthusiastically signed Social Security reforms he had urged for years, reforms that expanded Old Age and Survivors Insurance (OASI) benefits by 77 percent, expanded the covered population dramatically, and decreased the required contributions in the system. The result was a transformation in the program from a small system covering only 6 percent of the elderly in 1940 into a program “firmly established as part of American public policy (Sparrow 1996: 34).”

⁴ This section draws from Bryan D. Jones, Tracy Sulkin, and Heather Larson, 2003. Policy Punctuations in American Political Institutions. *American Political Science Review* 97: 151-72.

The 1950 legislation radically transformed the small program established in major amendments to the Social Security Act in 1939. The 1950 statutory changes caused an explosion in Social Security expenditures. From FY1949 to FY1950, real expenditures grew 3%. From FY1950 to FY1951, they grew 25%, and the next fiscal year grew 37%—the largest two-year percentage increase in the history of the program—even though most payments would come much later as the newly-covered retired. By 1952, expenditures had increased by an astounding 71 percent, and expenditures increased ten percent a year or greater for the next three fiscal years (True 1999).

Between these two landmarks, Congress enacted only two very minor adjustments to the program. This almost complete absence of legislative output was not for lack of trying. Presidents Roosevelt and Truman urged change; major advisory commission reports indicated the desirability of reform; many bills were introduced, and both House and Senate committees held hearings and reported legislation. Sporadic, but vocal and persistent, calls for reform emerged immediately after the 1939 enactments, and continued until 1950.

Moreover, there were good objective reasons for action. Sparrow (1996: 39) calls the failure to enact reform “puzzling”, and points out that “a further expansion in the Social Security system would have guaranteed a large net increase in federal revenues, since most of the government’s obligations would not be incurred until contributors retired. In the meantime, the government would receive desperately needed funds and would ease inflationary pressures by limiting consumer spending.” In other words, a ‘window of opportunity’ existed; the issue occupied the national agenda both before, after, and during the war; great effort was expended in proposals,

bill writing, and hearings; yet Congress nevertheless failed to pass legislation. When Congress finally acted, the result was not incremental adjustment, but major policy reform.

In democracies at least, it is easier to talk about an issue than to get serious action on it. In the United States, executive support or even support of legislative leaders may not ensure passage of a popular act; the system requires concurrent majority support in both houses of the legislature in addition to the President's acquiescence. In the case of Social Security, Republicans posed important counter-arguments to Social Security expansion based on wartime revenue need—a kind of 'lockbox' argument mimicked sixty years later. Airing of arguments takes time; opposition to change can be entrenched; even extraordinarily good ideas can be thwarted for a long time.

We can take the Social Security story one step further. Figure 11.1 diagrams the percentage change in real Budget Authority and major Social Security Amendments since the Second World War. These are the social security laws that were rated as among the most important statutes enacted by our weighting system, as described in Chapter 8. Major statutory enactments are associated with major changes in budget commitments, even though changes in social security commitments affect future retirees far more intensely than current ones. Even a policy area that is thought of as quite stable and consistent is rent with periods of punctuations (True 1999). All budgetary punctuations are not associated with major statutory reforms—there are in effect more budgetary punctuations than major statutes—but many are. It would seem that the major story in social security is not major reform then incremental adjustments, but major reform, incremental adjustment *and* substantial policy reform.

[Figure 11.1 about here]

The general lesson is that policymaking institutions are ‘sticky’—they do not respond simply or directly to demands or needs. Social, economic, and political changes are not automatically registered in public policies in democratic systems. Moreover, agenda access itself may also exhibit ‘stickiness’, albeit on a less dramatic scale. The agenda space is severely limited, and many issues can compete for the attention of policymakers. In the case of Social Security, during the “missing decade” (Sparrow’s term) many bills were introduced in some years (54 in 1941), while very few were introduced in others (17 in 1943) (Sparrow 1996: 60). Executive and legislative attentiveness shifted from intense to absent during the long period of policy inactivity as other matters clamored for governmental attention.

Political institutions impose costs on policy action in direct proportion to how far a policy proposal has proceeded in the lawmaking process. It is easier to get an issue discussed than it is to get serious attention for a specific line of policy action; it is easier to get a hearing on a bill than to schedule it for a vote; it is easier to get one house of Congress to pass a bill than to get it enacted. These institutional costs act in a peculiar way. They keep the course of public policy steady and unvarying in the face of lots of changes; that is, they do not allow for continuous adjustment to the environment. Early decision theorists termed this pattern ‘incremental’. But these costs also cause major policy changes when dynamics are favorable—a ‘window of opportunity’ opens, in the language of Kingdon (1994).

Policymaking and Earthquakes

We may think of policy processing within an institutional framework somewhat analogous to the geophysicists’ plate tectonics. The earth’s crust is divided into major segments, or plates, that slide against one another. Plate tectonics explains continental drift; it also accounts for earthquakes.

If we order observed earthquakes from the smallest to the largest, we will observe many, many, very small earthquakes (the incremental), and a number of very large ones (the punctuations), but very few moderate quakes proportionately. The earth's surface plates slide against one another, driven by powerful forces in the earth's core (the 'inputs'). Even in the face of powerful dynamics, however, much of the time plate friction holds them together, allowing small magnitude quakes to release some of the pressure. But as time progresses, and the pressure builds, at some point a potentially disastrous 'jump' must occur, bypassing a more moderate response to the pressure and resulting in the large magnitude quake. The 'inputs' of plate pressure are not directly translated into 'outputs' of earthquakes.

The Social Security story parallels this description of plate tectonics. During the 1940s focused demands resulted in minor adjustments to the 1939 basic statutory structure—minor incremental adjustments.⁵ When reform came in 1950, it came in monumental fashion. No moderate adjustments occurred anytime during the period.

Geophysicists cannot observe the friction of plates directly. Instead, they measure the outputs from the process (the earthquakes), and study their frequency distribution. We are in a similar position, but with one important advantage. It is relatively easy to order political institutions according to the extent to which they impose decision costs on policymaking activity. To the extent that a political institution adds decision costs to collective action, the outputs from that institution will exhibit periods of stability ("gridlock") interspersed with periods of rapid change. The higher the decision costs that must be overcome to achieve a collective goal, the more punctuated the outputs are likely to be.

⁵ Seamen affiliated with the War Shipping Administration were added to the system in 1943 and survivors of veterans killed in the war were added as beneficiaries in 1946 (Sparrow 1996: 36)

Costs in Political Institutions

We need to be somewhat more precise about the idea that decision-making costs lead to institutional friction. The payoff will be a very general framework for studying political, economic, and social change where interactions among actors are structured by institutions. An *institution* may be defined as a set of individuals acting according to common rules resulting in collective outcomes. Institutional rules are not neutral, in the sense that different rules often lead to different outcomes (Jackson 1990: 2). These aggregations of individuals interacting according to rules react to information from the environment and come to a collective response (even if the collective response is simply the sum of individual actions, as it is for markets and elections and roll-call voting in legislatures).

Decision-making systems impose four kinds of costs in making decisions in response to a changing environment: decision costs, transaction costs, information costs, and cognitive costs. *Decision costs* are costs that actors trying to come to agreement incur. They include bargaining costs and institutionally-imposed costs, such as those built into a separation of powers governing arrangement (Buchanan and Tullock 1962; Bish 1973). *Transaction costs* are costs that parties incur after they come to agreement (North 1990). In market transactions, these involve such items as the cost of insuring compliance to contractual agreements, and other payments to third parties to complete the transaction. It ought to be clear that in advanced democracies decision costs in the policymaking process heavily outweigh transaction costs. Bringing relevant parties to agreement in a system of separated powers (decision costs) generally outweigh the costs of holding hearings, enacting statutes, or changing budgetary allocations once agreement has been reached (transaction costs). In any case, we combine these costs in our analysis, terming them together 'decision costs'.

Information costs are search costs—costs of obtaining information relevant to making a decision. These are costs that exist when a person (or an organization) *wants* to make a decision. *Cognitive costs* are costs associated with the limited processing capacity of any social institution comprised of human beings. These are costs that occur because people *don't know* they need to make a decision. If one is not attending to a key component of the environment, then he or she cannot decide whether or not to incur search or information costs.

Institutional costs in politics may approximate the manner in which friction operates in physical models. When friction is introduced into idealized physical models, nonlinear systems result (Bak 1997). Such open systems result in an output pattern that is episodic and punctuated, with extraordinary difficulty in making point predictions. Earthquakes are an example. Predicting a particular earthquake is not currently possible, but the patterning of earthquakes follows a lawlike pattern—the power function of the Gutenberg-Richter law (Rundle, Turcotte, and Klein 1996; Schroeder 1991, West and Deering 1995; Bak 1997; Sornette 2000).

Information Processing in Institutions

The manner in which a policymaking system responds to information is critical in assessing policy change. Probably the major problem with the initial incrementalist model of policy change is that it did not incorporate the flow of information from outside the system. No matter what the external challenge, the system responded incrementally. That is quite unrealistic, and lead to models that are easily rejected. If we can understand how decision-making systems respond to information in the absence of any institutionally-imposed costs, then that idealized model can serve as a basis of comparison for systems that impose such costs.

A hypothetical fully efficient decision-making institution that imposed no costs would respond seamlessly to the world around it. That is, it would incorporate all relevant aspects of the information it encountered, and would ‘use up’ all the information in its decision-making process. The outputs of such a system would perfectly reflect the information flows coming from its environment (Simon 1996). If there were big changes in the environment, the system would respond with big changes. Similarly, small changes would generate only small changes. The major example of such a cost-free system is the classical model of a competitive economy.

In such a pure system,

$$R = \beta S \quad (11.3)$$

R = response = ΔO = change in output

S = information (signal)

β = benefits derived from the information flow (< 1).

The system reacts directly to the input flow by changing its output. What happens in real institutions in which decision-making costs are imposed? If costs are assumed to act linearly on the system, then

$$R = \beta S - C \quad (11.4)$$

Our hypothetical system continues to respond directly to the input flow. Now, however, it will not act until it recovers the costs that must be invested in reacting to the flow of information.

Equation (11.4) captures what political scientists know as Eastonian systems theory (Easton 1953).

The information coming into the system is not enough alone to force any change. We need to take into consideration decision-making costs as well. Where costs are low, signals of low power get reflected into public policy. Where costs are high, only the strongest signals are translated into public policy.

This is somewhat too simple, however. In politics costs are imposed only when actors take the trouble to use the system to block action. For minimal changes, actors that would normally be opposed might not take the trouble. For major changes, they can mobilize and make use of the system to try to block changes. This leads to a model in which costs go up as the strength of the signal increases. While we cannot know exactly the form of the equation translating inputs into outputs, we do know that it is of the general form of Equation 11.5. The signal and institutional costs interact with each other to magnify the effects of the signal.

$$R = \beta S \bullet C \quad (11.5)$$

Distributions

In complex systems, such as policymaking systems, it is very difficult to predict with any precision the particular output at a particular time. Moreover, it is generally extremely difficult to measure the information flow independently of what actors attend and respond to. That means matching the response to the signal is really not possible outside of detailed case analysis such as the investigation of the disintegration of the space shuttle Columbia in 2002. A major lesson of that work is the tendency of people to ignore or downplay important facts that don't fit in with the preferred way of viewing the world.

One response might be to think of news coverage as an input variable, and to try to relate policy outputs to 'the news'. This is not an appropriate strategy, however, because media scholars have shown that in many cases the news responds to what policymakers are doing, and are subject to manipulation by them (Bennett 1990).

As a consequence, we have focused on full distributions of outputs. So we return to the issue of what distributions with and without institutional costs would look like. Figure 11.2 depicts a frictionless, cost-free policy system, a system with institutional costs, and an interactive system.

The frictionless system is highly sensitive to incoming information. For a hypothetical one-unit change in relevant information, the system responds with a proportional level of outputs. The system with institutional costs ignores signals of low intensity, then responds proportionally to the strength of the signal after some threshold in signal strength is reached.

[Figure 11.2 about here]

All other things equal, a decision-making system with interactive costs as depicted in Figure 11.2 would lead to a leptokurtic policy change distribution. It should be easy to see why. Think about budgets. If we start at the point of the last budget allocation, and signals in the policymaking environment indicate the need to spend a little bit more, then the frictionless system does exactly that. In the system with costs, there is no change when the signal is weak, however. There is no change around the point of zero change for low signal strength—that is, for low power signals, the system returns a value of zero. Then, as signal strength increases, budget punctuations will occur because of the absence of marginal adjustments to information flows. Finally, the interactive system depicted in the graph (it is idealized, because we don't know the exact form of the relationship) suggests even more leptokurtosis in output change distributions.

In idealized information processing systems, the output (decision) result is entirely a function of the information received. This is a standard against which other processes may be compared. In ideal systems, if the information flow were distributed in a Normal or Gaussian fashion, then the output flow would be Normal, as the Normal distribution is invariant under a linear transformation. If the input is Normal, and the system just takes the information and produces policy outputs proportionally, then the outputs will be Normal.

There are good reasons to expect the Normal as an input distribution. Decision-makers are bombarded with information from many sources, and they may seek out other sources to inform themselves. How do they prioritize the information from these many sources, and how do they combine them in making decisions? The best way would be to weight the information streams by importance and add them to make an index. If participants in an institution receive information from independent diverse streams and weight and sum these diverse streams in an index, then the resulting distribution would be Normal by the Central Limit Theorem, at least in the limit.⁶ That is, across lots of decision-makers and lots of decisions, the distribution of information (coded, for example, as more or less favorable to a particular alternative) would be Normal.

Let us now turn to less-than-perfect human systems. The model of implicit indicators, described in Chapter 4, is our approach to the addition of ‘cognitive costs’ in models of policy change. If individual decision-makers rely on a limited set of indicators to monitor their environments, and update them or include newly-salient aspects of the environment in the decision-making calculus episodically, the result will be a flow of ‘news’ (that is, the information flow that the decision-maker attends to) that is not Normal (Jones 2001: Chapter 7).

Basically the cognitive architecture of the decision-maker imposes a selective bias on the flow of information. Of course decision-makers in politics will not cling forever to bad information, but they undoubtedly believe it far beyond its utility. When the information is

⁶ The input flow, however, may not be Normal because whatever processes generate that flow do not approximate the Central Limit Theorem. In that case, we lose a convenient standard, but the basic argument should nevertheless hold: the higher the institutional costs on collective action, the more punctuated the outcome pattern. Three CLT assumptions can lead to difficulties. They are the independence assumption, the assumption that a decision-maker does not rely too heavily on a single source, and the finite variance assumption (very extreme values cannot occur).

exposed as faulty, the decision-maker must shift dramatically to a new understanding of the situation. In effect, the decision-maker locks choice into a set of facts based in the past, and must update in a punctuated manner in the face of change that cannot be ignored. The ‘news’ is leptokurtic.

If the news is leptokurtic, outputs in completely efficient institutions will be leptokurtic. Since collective outputs from markets and elections are simple aggregates of individual decisions, outputs will follow news flows. So even in cost-free systems, if we add cognitive costs, outputs will be leptokurtic.

Complexity in Information Processing

The preceding discussion reflects the complexity of human decision-making systems. We have here tried to simplify by analyzing institutional costs within a simple framework. The key question is how people interacting in political institutions process and respond to signals from the environment. If institutions add friction to informational inputs, then outputs will not be directly related to inputs. But how will inputs and outputs differ in policymaking systems? We posit that whatever the input flow, the output flow will be both more stable (ignoring many important signals) and more punctuated (reacting strongly to some signals).

Lots of work in political science points toward an information-processing approach with political institutions playing major roles. Institutional analyses show that a ‘policy core’ exists that is not responsive to changes in preferences (for example, through replacement of legislators in elections), but when preferences change enough to move the pivotal legislator’s preferences outside the core, then major punctuations in policy can occur (Hammond and Miller 1987; Krehbiel 1998). Policy process scholars have argued that policy agendas change when

attentiveness and mobilization are directed at particular aspects of a complex environment, raising the probability of major policy innovations based on new ideas. Again, stability (when attention is not directed at the issue) and punctuation (when it is) occur in a single process (Baumgartner and Jones 1993). Similarly, in elections, first past the post voting systems and partisan identifications by voters operate together to add great stability to election patterns that are nevertheless occasionally disrupted by realignments.

Institutional decision costs will add to the kurtosis of output distributions. Difficulty in changing the status quo results in incremental-style decision-making rather than reform. This shows up in output distributions as the tall central peak associated with leptokurtic distributions. When change occurs, it requires substantial mobilization to overcome the stasis associated with the workings of political institutions and the tendency of humans to adopt rules of action that are difficult to change. As a consequence, when change occurs, it tends to be relatively extreme. This results in the characteristic ‘heavy tails’ and ‘weak shoulders’ of leptokurtic distributions.

Conclusions

The American policy process follows a distinct pattern. Attention is directed at a perceived problem. The matter is raised to the formal agenda. Policy action may or may not result. If it does, serious budgetary commitment to the matter may or may not occur.

Each stage of the process requires the expenditure of considerable resources to overcome the dual frictions of the attention spans of actors and the cost structure of political institutions. Political activists work to move their favored policy into the limelight, but the attention of policymakers may be directed elsewhere. When attention shifts, bills can be introduced and hearings can be scheduled, matters with low institutional costs. But the next stages, passage of

laws and major budgetary commitments, require much more effort—because of the design of American governing institutions.

The interaction of political demand and institutional costs is written in traces of policymaking of activities across time. This trace is, as a consequence of this interaction, punctuated with periods of intense activity and periods of more quiescence. If we aggregate the time traces of policy change, their signature frequency distributions will be leptokurtic.

Three models of change have been developed to understand this signature pattern: policy cascades, decisional sieves, and institutional friction. Each alone can account for leptokurtosis in policymaking distributions. When combined, we have an order of complexity that defies simple analysis.

Staying focused on the distributions, however, allows patterns to emerge. If we order these distributions from those from institutional arrangements that impose the least costs on collective action to those that impose the most costs, the kurtosis of the associated distributions should move from low (close to the 3.0 associated with the Normal) to high. In the next chapter, we do exactly this.

Figure 11.1: Percentage Change in Real Social Security Budget Authority and Major Social Security Enactments




Figure 11.2: Information Processing Policy Systems with Institutional Costs

