A DYNAMIC DECISION THEORETIC APPROACH TO MODELING POLICE PERFORMANCE

Thomas F. Cooley
Department of Economics

Richard A. Berk
Department of Sociology

Group for Research on Social Policy

University of California -- Santa Barbara

November 1978

This paper was read at the 1978 meetings of the Operations Research Society of America, and support from the NILECJ is gratefully acknowledged.

This work was performed under Grant No. 78NI-AX-0003, Performance Measurement and the CJS: A System-Level Perspective--Phase I, from the National Institute of Law Enforcement and Criminal Justice to the Georgia Institute of Technology. Points of view or opinions stated in this document are those of the authors and do not necessarily represent the official position or policies of the United States Department of Justice.

92504

I. INTRODUCTION

The performance of criminal justice systems is often approached in rather global terms through some set of loosely defined objectives. Such efforts have been thoroughly criticized by Deutsch (1977) among others as lacking a clear and comprehensive conceptual basis. Hypothetical constructs are poorly formulated, and there are few concrete mechanisms linking the inputs of the criminal justice system to the outputs of the criminal justice system. Clearance rates, for example, are assumed to respond to police budgets, but even if empirical relationships are found, their meaning is often ambiguous. It should come as no surprise then, that evaluations employing such perspectives are rarely compelling.

A second difficulty with most criminal justice evaluations is that even when a clear, formal model is posed, policy assessments are typically undertaken in static terms. We learn about the relationship between inputs and outputs at the margin, but little of the sequential nature of criminal justice functioning. We may discover through static comparisons how an increase in police personnel can affect the "supply" of crime, for instance, but little about longitudinal processes by which police officers may affect crime rates. While we would not deny the usefulness of static models, we also believe that much can be learned from dynamic approaches.

Of course, it is one thing to argue for formal dynamic approaches in principle, and quite another to put such preaching into practice. To begin, one must develop dynamic models of the criminal justice system which, though abstract, do not seriously distort the empirical phenomena in question. Then, one must operationalize a large number of hypothetical variables so that their epistemic correlations are well above zero. This is, of course, a non-trivial problem on which even the best dynamic model may founder. But, given the set of operationalized variables, one still needs to return to the empirical world to obtain adaquate measures. Here a range of practical constraints are likely to intervene, perhaps undermining the entire enterprise. Finally, assuming that adaquate data may be obtained, the system's parameters remain to be estimated. Often this introduces yet another group of difficulties.

While we would be among the last to claim that solutions for these and other problems are readily available, we will attempt to demonstrate in this paper that the application of dynamic models to evaluations of criminal justice performance is more than fantasy. We will focus on a subsystem within the criminal justice system, the police, propose a broad theoretical framework, operationalize many of its key concepts, and then suggest how it may be applied in practice. In the process, we hope to not only document the variety of likely pitfalls, but also demonstrate the promise of dynamic approaches to assessments of criminal justice performance.

The viewpoint adopted in the present paper is that part of the problem with attempts to measure the performance of the criminal justice system is the very macroscopic level at which they are undertaken. The criminal justice system is a multicomponent social organization. When it is viewed as a whole, an implicit assumption is made that the goals one ascribes to the system, reflect the goals of individual decision makers within that system. There is consistency in the objectives of all those who transform inputs into performance. This is clearly a heroic assumption. If one views the system at a more microscopic level by choosing a particular component and studying it in some detail then

there is a better chance that the objectives of individual decision makers will be consistent with the objectives of policy makers within the subsystem. In addition, one can begin to discuss the objectives in a more concrete fashion, one can identify factors which affect system performance (and these are likely to be more than just inputs), one can begin to conceptualize what a model of the system will look like, and finally one can begin to develop a measurement strategy. It is the latter component which is the ultimate goal of our analysis, but it is our view that measurement in the absence of a formal conceptualization of the system we are dealing with and the use to which that measurement will be put, is likely to be less useful than a measurement that is oriented to the particular needs of a particular problem.

II. A DECISION THEORETIC FRAMEWORK FOR THE ANALYSIS OF POLICE BEHAVIOR

The subsystem we have chosen to consider in this paper is the behavior of on-duty police officers. This is clearly one of the most micro level subsystems of the criminal justice system, and so it seems an appropriate level at which to begin.

Our first task is to describe the nature of decision making in such a system and the generally accepted framework for decision theory consists of the following:

1. A set of agents (decision makers, controllers, actors).
2. A preference ordering for each agent (utility function, performance criterion, pay-off function).
3. A set of permissible decisions for each agent (sanctioned actions).
4. Agents make some observations regarding the state of the world in which they are operating. The relation between the actual state of the world and the observations is called the information structure.
5. Each agent has a model of the system which is a formal description of the way in which the state of the world and the decisions of the agent affect the outcomes of the system.

The behavior of agents is characterized by a set of decision rules or control rules which describe the decisions to be taken when certain observations are made. We assume that preferences and the set of permissible decisions are exogenous to the system being studied. There are a number of possible ways one could further specify the decision problem at this point. The three most prominant ones are:

1. There is only one agent with one preference function.
2. There are many agents with the same preference function.

3. There are many agents with different preference functions.

The third category is usually called a game and it is clearly not appropriate for the situation we are considering. In addition, mathematical descriptions of such problems do not lend themselves easily to measurement strategies and empirical implementation. The second category is further divisible into a) Team Theory, and b) Hierarchical systems.

While there is a substantial literature on team theory (Marschak and Radner 1972; Marschak 1955) it is the most abstract and formal treatment of the problem of decision making in large systems, and does not lend itself easily to empirical applications. One reason may be that optimal decision rules for teams are considerably more complex (analytically and computationally) than they are for systems with centralized decision making (one agent). It is also not clear to what extent one would want to endow an individual on-duty police officer with decentralized decision making despite the informational and computational efficiencies which might result (although it is also quite apparent that the very nature of their environment requires that they often function this way).

Clearly, however, there is a definite hierarchy to the system we are describing, and it would appear that a hierarchical multi-agent system is a good place to begin. Such a system would consist of a number of interacting subsystems each under the control of one agent. The subsystems would be arranged in a hierarchy of levels. Communication would be lateral as well as hierarchical. Unfortunately, while such systems have been discussed and classified (Simon 1962), the formal mathematical treatment of such systems is still relatively undeveloped.[1]

Finally then we are left with category 1, which, while not completely satisfactory as a description of the decision environment we are addressing, is certainly more promising than most of the others. The one agent/one preference function has well formulated mathematical theory in both continuous time and discrete time, which should lead us to a much clearer understanding of objectives, policies or decision rules, and finally to a measurement strategy that is appropriate for articulating optimal decision rules. The requisite mathematical tools can be found within optimal control theory on which there is by now a vast literature in mathematics, engineering and economics. In subsequent sections then we will envision the problem of controlling the performance of on-duty police officers as one of formulating rules or control signals which may be communicated to them from a central decision maker on the basis of information about their "states" at different times. The rules will be designed to maximize police performance in terms of centrally determined objectives given the environment in which they operate. These issues are discussed in detail in the next section.

III. OPTIMAL CONTROL OF THE POLICE

One of the basic premises of the decision problem is that agents possess a model of the system which formally describes the way the state of the world and the decisions of agents affect system outcomes. More formally, we assume that the functions of an on-duty police officer can be described by a system of discrete time, linear difference equations of the form:

$$y_t = Ay_{t-1} + Bz_t + Cx_t + u_t, \qquad (1)$$

where y_t is a vector of outputs, y_{t-1} is the same vector of outputs in

the previously observed period, z_t is a vector of variables which are exogenous to the system (that is they are not determined by the actions or outputs of the police) but which describe the environment in which they operate, x_t represents a vector of inputs or policy variables which are subject to control (also called control variables or instruments), A,B,C are known matrices of coefficients, and u_t is a vector of unobserved stochastic shocks to the system.[2] By outputs we mean the set of measureable activities which represent the functions of police officers. For example, we might include among the y_t arrests, bookings, security checks, mutual aid assists and arrivals at given locations. The variables z_t describe the "environment" in which police operate. If the unit of analysis is a city precinct, z_t might include the ethnic and demographic makeup of the precinct, the presence of parks, mix of public transportation and other exogenous factors which are correlated with outputs. The variables x_t include policy variables which affect police performance such as the number of officers, patrol cars, routings, and rules governing apprehension and arrests. The u_t represents the random unpredictable occurences which affect the outputs of the system, and which are unknown.

The essence of the optimal control problem is to choose an appropriate set of policies to drive the outputs of the system. The model, as we have expressed it does not explain how policy variables are determined. If we specify a rule or equation for the policy variables, they will become endogenous to the system.

There are two kinds of policies that can be considered. The first is a set of policies which are specified at the beginning of a planning period and are not altered by future events, these are called control rules. The second kind of policy is one in which future policies depend on future events according to some rule. These are called feedback policies, described in the form of an equation

$$x_t = Gy_{t-1} + g \tag{2}$$

where x_t is a vector of policy variables, G is a matrix of coefficients and g is a vector of constants. Both kinds of policies will be relevant for the problem at hand. Clearly, some of the police functions are best governed by rules which are essentially invariant to the occurances on my given day. There are some, however, which are best determined as the result of the functioning of the system at a point in time.

We now need some means of judging the outcomes of a given policy or set of policies. What we want is some way to assess whether the outputs of a policy (the y_t elements or a subset of the elements of y_t) are desirable. An objective function is a scalar function which measures the desirability of the outputs or their characteristics (such as the means, variances or covariances). When the variables are stochastic the objective function will also be stochastic and some parameter in the distribution of this stochastic function must be selected as the criterion (e.g. the mean). One of the most mathematically convenient and tractable forms of the objective function is the expectation of a quadratic function of the stochastic outputs of our model.

Let us assume a quadratic objective function, and then we can state the optimal control problem. Specifically, the problem is to minimize

$$E(C) = E \sum_{t=1}^{T} (y_t - y_t^*)'W(y_t - y_t^*) \tag{3}$$

subject to the constraints of equation (1) (the model of the system), where

the y_t^* (t=1,...T) are exogenously determined performance targets for the outputs y_t, the expectation E is conditional on the initial state of the system (y_0), and where W is a symmetric positive semi-definite matrix which weights the different objectives and establishes their relative importance in the criterion function. Thus, the optimal control problem is to choose values of the x's (policies) which minimize the deviations of the output variables y_t from their target or desired levels y_t subject to the constraint of the equations of the model which describe how the policies determine the outputs. The exogenously determined elements of W reflect the relative importance of the different targets. The number of arrests, for instance, may be twice as important as the number of apprehensions.

The optimal control problem as we have stated it has two components, one deterministic and one stochastic. The deterministic part of the control problem is obtained by setting the random disturbances in our model of police behavior equal to zero, their assumed mean value. We can write the deterministic model as

$$\bar{y}_t = A\bar{y}_{t-1} + B\bar{z}_t + C\bar{x}_t \tag{4}$$

where the bars indicate that we are dealing with the deterministic part of the model. Now let $y_t^s = y_t - \bar{y}_t$ and $x_t^s = x_t - \bar{x}_t$, then the stochastic part of the control problem uses the stochastic model

$$y_t^s = Ay_{t-1}^s + Bz_t + Cx_t^s + u_t \tag{5}$$

which is derived by subtraction of (4) from (1). Many researchers have considered only the deterministic control problem even in situations where the system is known to be stochastic (see Pindyck, 1972) arguing that the added costs and complexity of considering the stochastic control problem may not be warranted given the insensitivity of the model to random shocks. This seems inappropriate for the police, however, because the relevant systems are likely to be extremely sensitive to the stochastic components of the model.

If the control problem is treated as deterministic then the problem is reduced to one of minimizing

$$C = \sum_{t=1}^{T} (\bar{y}_t - y_t^*)' W_t (\bar{y}_t - y_t^*) \qquad (6)$$

with respect to $\bar{x}_t$ subject to the constraint of the model in deterministic form (4). The shocks to the system, however, make it necessary to modify the deterministic control solution x_t by x_t^s in order to compensate for the deviations of y_t from $\bar{y}_t$. To reflect this, the objective function is decomposed as

$$E(C) = \sum_{t=1}^{T} (\bar{y}_t - y_t^*)' W_t (\bar{y}_t - y_t^*) + E \sum_{t=1}^{T} y_t^{s'} W_t y_t^s = C_1 + E(C_2) \qquad (7)$$

The two parts of the control problem then are to minimize C_1 with respect to $\bar{x}_t$ subject to (4) and to minimize $E(C_2)$ with respect to x_t^s subject to (5) The optimal policy x_t then is defined as the sum of x_t^s and $\bar{x}_t$.

There is not sufficient space in this paper to go into a detailed discussion of the various methods of solution of the optimal control problem. The most common method of solution is to apply dynamic programming (see for example Bryson and Ho, 1969; Chow, 1975; and Aoki, 1967) although the problem can also be solved by the method of lagrangian

multipliers or a discrete version of the Pontryagin minimum principle.

What we have done in this section is to lay out the conceptual basis for finding decision rules to govern the performance of on-duty police-officers. We have articulated a method for finding rules (policies) which are optimal, given an objective function which represents desired levels of performance of the system and weights them according to their relative importance. An essential characteristic of this conceptual framework is that it treats the problem in a dynamic framework. Policies and performance are viewed as a dynamic sequence unfolding in time which is the way the system of policing evolves in real time.

There are many issues which we have not discussed or have treated lightly in laying out the problem, and we must now provide some added dimension to the problem. The next section discusses in more detail the formulation of objective functions. This is followed by a discussion of the optimal control problem which suggests some more complex control procedures and algorithms. We then discuss the sorts of measurement and sampling strategies that are implied by the optimal control framework. The final section presents some conclusions and suggestions for future research.

IV. OBJECTIVE FUNCTIONS

In section III we introduced the notion of an objective function which represents the desirability of the outputs or their characteristics in terms of their weighted deviations from target values. The notion that such targets exist may seem naive given likely disagreement on what the targets are and even greater disagreemnet on their relative importance. Nevertheless, it must be recognized that virtually all bureaucratic systems will in fact function with implicit objective functions and implicit weightings of objectives.

It is also important to recognize an objective function is only an approximation of the preferences of the policy maker just as the model described in equation (1) is only a statistical approximation of the functioning of police officers. Yet, if an objective function is found to be imperfect, the function can be revised and the resulting impact on policy performance can be re-evaluated. The crucial point is that explicit statements impose a discipline on the formation of policy that is likely to be missing otherwise. In the best of circumstances, they may lead to substantial improvements in the performance of the system.

The objective function presented in section II is really the most elementary we could have ocnsidered and it will be useful to expand on that basic framework. One of the immediate extensions that can be made is to recognize that there may be costs associated with policies and that the resulting benefits may not outweigh the costs. For example, if the state of the system is such that more arrests could be made or crimes could be reduced by having a more extensive coverage of a given area, then the implied optimal policy might be to have police travel singley rather than in pairs. There may be a cost to such a policy, however, if the personal risk to the police is increased substantially. Even if the optimal policy would be to simply put more pairs of police on duty, this is not done costlessly since there are increased payroll costs, or costs in terms of coverage of other shifts. The costs of manipulating instruments can be incorporated by changing the specification of the objective function to the form

$$E(C) = E \sum_{t=1}^{T} (y_t - y_t^*)'W_t(y_t - y_t^*) + E \sum_{t=1}^{T} (x_t - x_t^*)'R_t(x_t - x_t^*) \qquad (8)$$

where x_t are the policy instruments, x_t^* are the targets for those instruments, and R_t is a positive definite matrix, the diagonal elements of which give the relative costs of deviating from the nominal or target paths of the policy variable. The relative magnitudes of W and R represent the relative costs of control versus the objectives of control.

The great advantage of the quadratic form of the objective function is that it is computationally tractible and yields linear decision rules when applied to linear models. It does have some defects however. One important shortcoming is that it is symmetric. Consequently, it implies that the costs of overshooting a target are the same as the costs of undershooting a target. This may not be a viable assumption for the problem we are studying. It could well be that the costs of achieving too few apprehensions of suspects (for example) are greater than the costs of achieving too many. Another perhaps less important shortcoming is that it is additive. This implies that the expected cost of achieving objectives is essentially independent between periods. We may wish to penalize positive or negative covariance of outcomes between two periods.

The symmetry problem can be resolved by considering objective functions that are piecewise quadratic. For each output variable and each policy variable the range of possible values can be divided into three regions; in the middle region no cost is assigned, but in the two extreme regions costs are assigned quadratically but asymmetrically. Let y_{it}, i=1,...K represent the K components of the vector y_t and let $\bar{\varepsilon}_i$ and $\underline{\varepsilon}_i$ represent the maximum upper and lower allowable deviations from the target trajectory y^*_{it} respectively. Then, the objective function can be expressed as

$$C = \sum_{i=1}^{K} f_i(y_{it} - y^*_{it}) \tag{9}$$

where $f_i(y_{it} - y^*_{it})$ has the form

$$f_i(y_{it} - y^*_{it}) = \begin{cases} w_{1i} \sum_{t=1}^{T} (y_{it} - y^*_{it} + \underline{\varepsilon}_i)^2 : y_{it} \leq y^*_{it} - \underline{\varepsilon}_i \\ 0 : y^*_{it} - \underline{\varepsilon}_i < y_{it} < y^*_{it} + \bar{\varepsilon}_i \\ w_{2i} \sum_{t=1}^{T} (y_{it} - y^*_{it} - \bar{\varepsilon}_i)^2 : y_{it} \geq y^*_{it} + \bar{\varepsilon}_i \end{cases} \tag{10}$$

In this example w_1 and w_2 are the quadratic penalties for deviations below and above the target range respectively. Similar functional forms can be applied to the cost associated with manipulating policy variables. The use of piecewise quadratic objective functions complicates somewhat the computational problems, but not greatly so.

The additivity and temporal independence of the objective function can also be remedied by including intertemporal covariance terms in the function. This is a more complicated procedure and will lead to more complicated optimization algorithms. The seriousness of this problem will depend on the particular circumstances and it is not clear at this speculative state that it is a subject of much concern for the problem at hand.

There may be functional forms other than the quadratic that are more representative of the true social costs of not obtaining adequate

performance from the policy system. However, the definition of societal goals in any parameterizable form is the most pressing problem at this stage, and the quadratic function does not seem especially unreasonable as a first approximation of the objectives to be pursued.

V. ALTERNATIVE CONTROL PROCEDURES

In the development of the optimal control problem presented in Section III it was assumed the agents had full knowledge of the linear system of equations which describe the functions of the police up to the additive stochastic errors. In this section we discuss briefly some of the possible extensions of the basic control problem which account for uncertainty about the structure of the model.

The basic equation describing the functioning of the police is

$$y_t = Ay_{t-1} + Bz_t + Cx_t + u_t \tag{11}$$

where all variables are defined as in section III. In that section, however, we assumed that the matrices of parameters A, B and C were known. This is clearly not going to be the case for the system we are studying. The most likely assumption is that the coefficient matrices will be estimated on the basis of observed data and so the coefficients themselves will be random variables.

This additional element of uncertainty in the system must be accounted for in the computation of the solution to the optimal control problem. The uncertainty in the problem is expressed in terms of the variances and covariances of the parameters and the optimal control policies and the associated expected minimum expected cost can be derived by the method of dynamic programming as in the original problem. The optimal feedback control policies will turn out to depend on the mean vectors and variance-covariance matrices of the estimated parameters. The optimal policies under these circumstances will differ from the policies with known parameters but it is difficult to determine a priori what the nature of the differences will be. The exact computational procedures for this case need not concern us here (see Chow 1975 for an extensive discussion of the computational details), but it is important to be aware of the necessary modifications.

Explicit recognition of the uncertainty about the structure of the system leads to another interesting possible formulation of the optimal control problem, one in which decision makers improve their knowledge of the structure of the system through the process of trying to control it. The basic premise of our conceptual picture of the performance of the policing system is that the measurable outputs are caused (in statistical sense) by the policies. Knowledge of the statistical link between the policies and the outputs, however, is imperfect because it is based on estimates that are obtained from historical data. Suppose that historical policies have been relatively passive and have not shown much variation, while outputs have varied considerably. Then, the statistical link between the policies and the outputs will be weak. It is easy to see that under these circumstances the pursuit of active and varied policies could lead to improved knowledge of the functional relationship between decision rules and policies. Indeed, whatever the historical pattern of policies and outputs it is possible to improve knowledge of the functioning of the system by introducing learning behavior into the model of decision making.

The optimal control problem can be formulated as an adaptive decision rule problem in which learning takes place on the part of the decision makers. There are a number of alternative ways in which the

problem has been formulated. MacRae, (1972) and Prescott (1972) have developed numerical approximate solutions to the adaptive decision rule problem, Tse and Bar-Shalom (1973) and Tse (1974) have developed solution methods which rely on conceptual approximations to the principle of optimality. One of the interesting characteristics of all of the solution methods is that adaptive policies are often more conservative (less active) than non-adaptive policies. That is, the future gains from active experimentation may not offset the costs incurred by the uncertain impact of experimentation. This result, however, depends on the structure of the objective function and system being controlled. It may well be the case that active experimentation will prove beneficial in the case of controlling the police system.

Finally, it should be noted that we have considered only the case in which the model of the system is linear. This restriction was intentional because the linear model is the basic building block of stochastic decision theory. All of the optimal control procedures discussed can be extended to deal with non-linear models but at some cost in terms of complexity and computational difficulty. It is unknown at this stage, whether the linear framework will be adequate for modelling police behavior.

VI. DATA COLLECTION FOR OPTIMAL CONTROL APPLICATIONS

Even the most sophisticated prescriptive models are of little practical use without empirical information on the phenomena in question. In the absence of data, links between the model and the "real" world cannot be established. For our purposes therefore, collection of appropriate data is critical; if there are no data, estimates of the coefficients in equations 4 and 5 cannot be obtained.

The observation that applications of optimal control theory require data states little more than the obvious without some consideration of the kinds of data one should obtain. While it is difficult to provide a firm game plan in the abstract (for reasons that will soon become clear), there are nevertheless several important issues which always must be broached.

Some General Definitional Issues

Looking back at equation 1, it is apparent that four basic kinds of data are required. The first, represented by y_t (and y_{t-1}) includes indicators of system outputs. For police, measures of direct contact with citizens, the speed of response to calls, reported crimes cleared by arrests are perhaps good examples although there are some non-trivial difficulties to be surmounted. To begin, it is often not clear a priori what is to be treated as an end in itself and what is to be treated as the means to that end. Is a large number of face-to-face contacts with citizens a good thing (or maybe a bad thing) or simply a vehicle for better police community relations? To some degree, the difference between means and ends lies in the eye of the beholder or a bit more formally, how one chooses to bound the system. And system boundaries are at least as much political decisions as technical decisions.

A second complication in obtaining output measures is determining the proper unit of analysis. Sometimes appropriate units may reflect rather arbitrary but hardly mutually exclusive levels of aggregation. One might consider, for instance, citizen complaints about "unnecessary" force directed at particular police officers or aggregated to the precinct level. On the other hand, one must not be blinded by conventional measurement decisions and under closer scrutiny, it may not be apparent what the proper units should be. For example, should one measure of

police performance be arrests or some of the components that go into making arrests? The latter might include the speed with which information is relayed to patrol cars, the thoroughness of that information, the speed with which police arrive at the scene, the techniques used in "hot pursuit" of a suspect, and so on.

A third and related difficulty is determining the time intervals by which data on outputs should be organized: by minute, hour, day, week, month or some other period. The issues here are often quite subtle and depend not only on earlier decisions about the proper units of analysis, but on a theoretical perspective about the system under study. In other words, a committment to the particular system one intends to monitor and control implies a theoretical view on the system's dynamics and in turn the time intervals at which measures should be collected and reported. A system that changes slowly over time may not need frequent soundings since little variance of interest will be obtained. A system that changes rapidly, in contrast, will require more frequent measures or important phenomena will be lost. Too often, however, output measures come in time intervals based on convention or organizational convenience. A detailed consideration of the performance of police on patrol, for instance, may be best viewed over five minute intervals rather than summary figures for an eight hour shift.

A second kind of data one must collect is represented by z_t in equation 1and stands for factors that are, from the perspective of police departments, beyond control. These describe the environment in which patrolling police operate: physical characteristics of neighborhoods (e.g. kinds of housing), statutes and ordinances affecting the behaviour of police and citizens (e.g. stop-and-frisk laws), the kinds of police technology available in principle (e.g. communications equipment), and other exogenous factors. Note that exogenous factors may be fixed (e.g. street layouts) or stochastic (e.g. weather), but in either instance, they are in principle observable and immutable.

As in the case of outputs, the definition of what is really exogenous, what the proper units of measurement should be, and the time intervals at which data should be collected are often ambiguous. For example, the yearly budget (once provided) is basically exogenous and clearly places important constraints on how police departments function. However, from year to year police departments may have significant impact on their appropriations and therefore, with a longer time horizon, budgets may be endogenous. Indeed, in the long run, almost any factor affecting the functioning of police is to some degree endogenous. Thinking a bit more about budgets, one could alternatively (or in addition) consider the monthly allocation of department resources to different components of police departments and, for example, the kinds of difficulties that might arise when unusual needs for particular divisions in particular months are not easily met.

A third kind of data required for optimal control is represented by x_t and measures factors over which police departments have control. These are the inputs or policy variables whose manipulation allow police to alter their performance. Again, there are a host of difficult definitional problems to be surmounted, but examples might include the allocation of personnel to various assignments (e.g. the number of police on patrol on weekends), the recruitment and hiring of personnel (e.g. the number of minority officers), administrative regulations (e.g. guidelines on the use of firearms), and dress codes (e.g. when short sleeve shirts may

be worn). Note that with some other perspective on the system under study, even variables such as the allocation of police personnel may be exogenous. If one were examining the detailed actions of police over a particular eight-hour shift, for instance, the number of officers on patrol is exogenous and therefore a fixed parameter.

Finally, there are events that are stochastic in nature and about which no independent and observable information really exists. These are represented in equation (1) by u_t and result from the difference between outputs projected by the system model and the observed outputs. In other words, the u_t are simply the residuals and are subject to a range of possible assumptions. For our purposes, we rely on the usual assumptions implying that the residuals result from a large number of small perturbations to y_t not captured by the model. We also assume that the expected value of u_t is zero and depending on the application, one must also make assumptions about the variance-covariance matrix of u_t (e.g., the covariances are zero).

To recapitulate briefly, we have been arguing for dynamic perspectives on the criminal justice system and particularly the police. This implies not only models in which time plays a critical role, but measures that permit concrete applications. Equally important, one cannot sensibly undertake measurement without one or more _a priori_ models of the system under study. In the absence of _a priori_ models, it will not be clear what to measure, whether the "what" should be treated as output, input, exogenous constraints, or stochastic perturbations, and how often observations should be made. Finally, the existence of a range of popular conventions speaking implicitly to such issues should not be allowed to obscure their problematic nature.

Difficulties with the Usual Data

It should be painfully apparent that a dynamic approach to criminal justice agencies places heavy demands in the quality of available data. One must obtain a sensible array of disaggregated indicators in inputs, outputs, and exogenous constraints. In addition, it is absolutely essential that indicators of each and every variable be collected over time and frequently enough to capture important, policy-relevant longitudinal variation. Typically, therefore, data of requisite quality will not be readily available, and a data-collection effort must be mounted.

Perhaps nowhere is this more obvious than in efforts to develop practical models for monitoring and controlling the behavior of police officers on patrol. While monthly summary statistics, often by precinct, are commonly available, these are usually inadequate.

Consider the following example. Suppose a local police department institutes a policy to aggressively stop-and-frisk any "suspicious" looking individuals on the streets in commercial districts after midnight. The immediate goal is to decrease burglaries. Also suppose that the only input measure is when this change in policy occurred and the only output measure is the number of reported burglaries, the latter organized into aggregate monthly figures by precinct. One potential problem is that a wide variety of factors may also affect the number of reported burglaries in any given month and that therefore changes in the number of reported burglaries as a function of changes in stop-and-frisk policies will be impossible to isolate. At best these other factors may be treated as "noise," but even under this convenient assumption, it is likely that the ratio of noise to signal will be so high that accurate parameter estimates will be impossible to obtain. In contrast, a time series of daily observations at least provides for the

possibility of disentangling police effects from noise. Put in other terms, aggregate monthly figures are actually the product of ad hoc smoothing, and are difficult to justify especially when statistical procedures exist which may model both the signal and the noise (e.g., Box-Tiao, 1975).

Assuming that despite using aggregate monthly figures a reduction in burglaries is found, one may still be seriously misled by insufficient understanding about the precise mechanisms by which the reduction occurred. One explanation for the reduction may be that prospective burglars are arrested (presumably with burglar tools) and are effectively put out of circulation. Another explanation may be that the stop-and-frisk policy effectively intervenes in time to prevent burglaries from occurring. Suspects are frisked and told to go home. Finally, stop-and-frisk policies may simply increase the visibility of police in commercial districts, and this alone may deter would-be burglars. Note that each explanation has different policy implications. The first implies that stopping and frisking suspicious individuals basically catches them with incriminating evidence from which they may be convicted of a felony. The second suggests that crimes are prevented in part by postponing them and perhaps by changing the calculus by which burglars operate (Letkemann, 1973). The third suggests that prospective burglars are deterred simply by the presence of police, and that stop-and-frisk practices per se may be beside the point. Note that the first two explanations may be interpreted as support for stop-and-frisk policies. The last explanation means that one may obtain the same result by simply patroling commercial areas more frequently after midnight. In any case, if detailed information were available about what police actually did on patrol, such confusions could be eliminated.

Now consider an extension of the stop-and-frisk example in which what police do on patrol must be linked to <u>when</u> they do it. Suppose one finds that after the new stop-and-frisk policy has been in effect for several months, there is a marked increase in the number of police-citizen encounters leading to violence. On one hand, such violence may result directly from heightened face-to-face contact between police and the kinds of "suspicious" individuals who frequent commercial districts after midnight. Stop-and-frisk practices may be less important than simply more numerous contacts between police and particular kinds of citizens. On the other hand, the stop-and-frisk practices themselves may cause the violence. Police may occasionally use the stop-and-frisk policy to rough up "undesirables" and/or citizens may respond to aggressive police practices with aggression of their own. One could begin to separate these two competing explanations by knowing whether the violence occurred immediately after the initial contact between police and citizens or whether the violence occurred after the frisking was initiated. In other words, one could determine whether the "stopping" or the "frisking" was the source of the violence, but to do so, one would have to have detailed information on the <u>sequence</u> of activities preceding violence.

It is important to emphasize that the kinds of aggregate figures routinely available from official records are typically lacking because both the longitudinal and cross-sectional dimensions have been aggregated. For any given time period, aggregating across individuals encourages the "ecological fallacy." Just because there are more arrests for burglary after stop-and-frisk policies are initiated, for instance, does not necessarily mean that it is the officers who are doing the stopping and frisking who are also making the burglary arrests. For a given police officer, aggregating across time may lead to analogous errors. If one finds that the first "burglary in progress" call to which the officer responds each night occurs on the average at 1:45 a.m. and the first burglary arrest of the night occurs on the average at 2:00 a.m. does not mean that the first burglary call necessarily precedes the first arrest (i.e., these may typically occur on different days). To make matters worse, in the usual official data both the cross-sectional and longitudinal dimensions are aggregated.

<u>Collecting Better Data</u>

It should be apparent that the usual data available from official records are inadequate for dynamic models of police patrolling practices. We turn now to some suggestions of how one might do better.

To begin, one must determine what variables <u>in principle</u> need to be included in one's system model. This implies returning to equation (1) and considering more thoroughly how the system is organized. With this accomplished, one may then turn to the practical matter of finding empirical indicators. Below are some <u>examples</u> of the kinds of variables one might include. Note, however, that these will necessarily depend on a kind of prior decision described above.

I. Outputs

A. Enforcement of Laws

1. arrests
2. citations

B. Maintenance of Order

1. crowd dispersal or crowd control
2. dispersing loiterers
3. quieting neighborhood disturbances (e.g., rowdy parties)
4. breaking up family fights
5. directing traffic

C. Crime Prevention

1. checking commercial and residential premises
2. patroling
3. educating the public (e.g., about proper door locks)

D. Social Services

1. convoy or transportation services (e.g., to hospitals)
2. finding missing adults, children and/or pets
3. first aid
4. counselling (e.g., on whether charges may be pressed)
5. providing information (e.g., directions to lost motorists)
6. referrals to other public agencies

E. Community Relations

1. meetings with school children
2. meetings with citizens
3. recreational services (e.g., the Police Athletic League)

F. Assistance to Other Criminal Justice Personnel

1. testifying at trials
2. providing information to prosecutors
3. providing information to detectives

G. Assistance to Other Public Agencies

1. referrals (e.g., to welfare workers)
2. protection (e.g., patroling in schools)
3. convoy services (e.g., for fire trucks)
4. crowd control (e.g., at public meetings)
5. calling in emergencies (e.g., fires)

II. Fixed Exogenous Factors

A. Legal Mandate

1. definitions of illegal activity by citizens
2. definitions of illegal activity by police (e.g., legal search and seizure)

B. Resources

1. budgets
2. technology
3. personnel

C. "Supply" of Crime

1. demographic factors (e.g., age distributions)
2. economic climate.
3. local "tastes" for crime (e.g., norms about marijuana use)

D. Physical Environment

1. location and mix of neighborhood types (e.g., commercial versus residential)
2. layout of streets
3. street lighting
4. location and nature of parks and other open spaces
5. public transportation routes
6. architectural features of structures (e.g., highrises versus single-family homes)
7. location of schools

E. Community Support

1. amount and kind of crime reported
2. willingness of witnesses to testify
3. willingness of witnesses to provide information

4. availability of informers

5. respect for police doing their duty

6. willingness of citizens to report police misbehavior (e.g., on graft)

F. Support from Other Public Agencies

1. cooperation from schools

2. information from hospitals (e.g., about "suspicious" injuries)

3. information from building inspectors

G. Support from Other Criminal Justice Agencies

1. cooperation from prosecutors (e.g., in deciding which cases to prosecute)

2. cooperation from judges (e.g., in setting bail, granting search warrants and sentencing)

3. cooperation from prison officials (e.g., in granting parole)

III. Stochastic Exogenous Factors

A. Where the Need for Police Services Occurs

1. visible from street or not

2. what neighborhood

3. accessibility

4. inside or outside

B. What Individuals Need Police Services

1. which particular citizens

2. which particular public officials (e.g., a school teacher)

C. When the Need for Police Services Occurs

1. time of day

2. day of week

3. month

4. season

D. Climate

1. temperature

2. precipitation

3. wind (which affects fires, for example)

4. humidity

E. Situational Factors

1. number of citizens present

2. whether drinking is involved

3. whether police response is police or citizen initiated (i.e., was there a complaint)

4. whether weapons are involved

5. density of vehicular traffic

6. relationships between citizens present (e.g., spouses)

IV. Endogenous Factors

A. Administrative Policies and Guidelines about Policing

1. definitions of "necessary" force

2. dress codes

3. protection of citizen rights (e.g., informing arrested suspects of their rights)

4. definitions of "good manners"

5. proper record keeping (e.g., filling out arrest forms)

6. testifying properly at hearings and trials

7. physical fitness

8. dispatching procedures

B. Training of Police

1. knowledge of the law

2. use of firearms

3. self-defense and restraining suspects
4. gathering evidence
5. filling out forms
6. human relations
7. first aid
8. crowd control
9. use of nightsticks, gas, mace and other social control devices
10. high-speed driving

C. Recruitment of Personnel
1. health requirements
2. age requirements
3. educational requirements
4. residential requirements
5. affirmative action
6. phychological screening
7. biographical background (e.g., previous convictions for felonies)
8. aptitude and achievement testing

D. Organizational Structure
1. chains of command
2. centralized or decentralized structure
3. grievance procedures for citizens
4. disciplinary procedures for police
5. promotion practices
6. job descriptions and backgrounds required
7. communications networks

E. Allocation of Personnel
1. by kind of task (e.g., patrol versus detective work)
2. by time of day
3. by neighborhood (or other geographical unit)
4. by time of year
5. in response to special events (e.g., parades)

F. Allocation of Hardware
1. patrol cars, paddy wagons, and ambulances
2. communications equipment
3. firearms, handcuffs and other equipment carried by officers
4. technical equipment in crime laboratories
5. data-processing equipment
6. outfitting of offices and station houses
7. clerical equipment (e.g., typewriters)
8. emergency equipment (e.g., riot gear)

It is perhaps worth stressing once again that the variables we have just listed are <u>examples</u> of the kinds of things one might want to consider and depend fundamentally on <u>how the system in question is bounded and at what level of aggregation one is working</u>. Nevertheless, it is easy to demonstrate that at least many common policy interventions with police can be roughly described with the framework suggested. For example, the Kansas City Preventive Patrol Experiment (Kellog et al., 1974) involved primarily a change in the allocation of personnel and transportation equipment coupled with an alteration in the procedures by which police were dispatched. Similarly, the Rochester experiment in which patrol officers were allowed to play a larger role in investigative work can be characterized by somewhat new job descriptions for police officers.

Unfortunately, even with decisions made about what to measure for the application of optimal control theory, one must still address when to measure. Since time figures fundamentally in any dynamic models, the longitudinal arrangement of observations is absolutely critical.

To begin, we shall assume discrete time as the subscripts in our equations imply, although for the kinds of variables involved, one typically has no choice. In addition, we will require that regardless of when observations are actually taken, the data are ultimately arrayed so that the time intervals between observations is constant. If, for example, time is only recorded when a particular kind of event occurs, one would assume that at other times the event did not occur and simply record the absence of the event for those times (as if one were taking measures, and it was not observed). Were this not the manner in which the data are organized, the values of y_t and y_{t-1} would tap different time lags as one stepped through the data and for any given value of y_t, the values of the other time dependent variables with the same subscripts would often correspond to different moments in time. The alternative of leaving a "hole" in the data for moments in which no observations were taken would virtually rule out conventional statistical procedures for dynamic systems of equations.

The requirement of discrete, equal interval units of time suggests that within the bounds of t_0 and t_T, some form of sampling must be undertaken. Even a decision to collect micro-level data on a particular officer on patrol at 2 minute intervals, for example, implies a form of "systematic sampling" (depending on how one starts the sampling sequence) in which events between the 2 minute intervals are missed.

The recognition that one must necessarily sample in time introduces a range of subtle questions for which no simple answers are likely to exist.

Perhaps most important, one must choose the frequency with which observations will be taken. This implies some formal conception of how often the system must be measured to accurately and efficiently capture the phenomena of concern. Sampling infrequently may obscur important dynamic relationships while observations taken too close together may provide a great deal of redundant information. If one wished to consider the sequence of significant events occuring during a typical 8 hours shift, for example, would one sample at 5 minute intervals, 15 minute intervals, 30 minute intervals or what? 30 minute intervals may miss, for instance, the links between a "burglary in progress" call and what police actually did upon arriving at the scene. 5 minute intervals might provide a wealth of virtually worthless information about an uneventful patrol, perhaps late at night.

In principle, one possible solution might rest on sampling different time intervals for different kinds of processes. One might sample at 15 minute intervals, for example, during uneventful periods and more intensively during periods when events were more densely packed. Assuming this were feasible (more on that shortly), one could then analyze the sparce and dense periods separately or insert reasonable values in the observation gaps in the sparce period (e.g. "nothing" happened). This leaves unanswered, of course, which periods will be deemed sparce and which dense and within them, how frequently to sample.

If a decision is made to sample with different intensities at different times, there are several strategies one might initiate. First, one might determine in advance that some periods during patrol are likely to be linked to dense sequences of events (e.g. between 3 PM and 4 PM, when school lets out). Second, one might make the collection of more frequent data collection contingent on some particular event, regardless of when it occurs. For example, one might begin collecting more frequent observations as a patrol car is dispatched to the scene of a crime. Third, one might space frequent and infrequent observations regularly throughout a patrol period. Data could be collected, for example, at fifteen minute intervals with a 15 minute period of 2 minute intervals inserted once an hour. Finally, one might insert these more intense data collection intervals at random throughout an 8 hour shift. A choice among these strategies clearly depends on what one knows a priori about the system in question. If one knows, for instance, which kinds of events are likely to be followed by important and dense sequences of activities, the second approach might be preferred. At the other extreme of total ignorance, one might rely on the last approach of random insertion of more intense data collection periods.

Assuming that one decided about how time will be sampled, one is still left the knotty problem of actually obtaining the sequence of observations across a range of variables. Basically, there are four possible sources of such information: first hand reports, data from observers, "traces" of events or technological "fixes." First hand reports involve information from participants at the scene (e.g. police officers, citizens, suspects, victims, etc.). Observation data implies the collection of information by "outsiders" who are strategically located (e.g. in police cars) to record what occurs when. "Traces" refers to data "left behind" by events of interest. Dispatcher's logs are perhaps a good example. Finally, technological "fixes" speak to monitoring devices which record sequences of events. One might, for example, place small radio transmitters in police cars which send out signals on some regular basis. Then, it would

be possible to monitor the location of patrol cars without requiring police officers to call in.

Each of the four modes of data collection suggest a variety of tradeoffs. Data from traces is perhaps the most easily obtained, but more likely to speak to the needs of police departments than the needs of researchers. Collecting first hand reports would in principle provide a wealth of detail, but would in practice have to be recorded after the fact. It is difficult to imagine police producing a narrative of a high speed chase as it occurred. Consequently, first hand reports would be subject to errors in recall. Technological fixes have the potential advantage of unintrusiveness, but only some kinds of events may be accurately monitored. Even if individual police officers carried small transmitters that allowed one to listen in to ongoing conversations, physical surroundings would be lost (unless police tried to report on them as well). In the best of all possible worlds, of course, one would design the data collection to capitalize on the strength of each data collection strategy and therefore use some mix of approaches.

It should be clear by now that the application of optimal control theory to police patroling practices is not an easy matter. One needs some a priori model of the system in question, a priori decisions of what and when to measure and then some feasible means to collect the necessary data. Given considerable ignorance of precisely what police do on patrol and no compelling systems of model of patroling activity, the practical use of optimal control approaches will likely remain an elusive goal for some time to come. Nevertheless, we have some suggestions about how one might initially proceed.

Data Collection Through Systematic Longitudinal Observation

Decades of experience in sociology and anthropology have demonstrated that a range of important phenomena can be effectively measured through first hand observation. Variously characterized as "field work" or "participant observation", this approach rests on a)careful observations, b)the use of on-the-scene memory assisting devices (e.g. notes), and c) the later construction of detailed descriptions of what was observed (McCall and Simmons, 1969; Junker, 1960; Schatzman and Strauss, 1973). While field work efforts are often characterized as "only" a "case study", in fact, effective field work researchers capitalize on many of the more powerful quasi-experimental designs described by Campbell and Stanley (1963). To take a simple example, in his important study of police on patrol, Reiss 1968, employed (implicitly) a matched, multiple non-equivalent control group design to compare how police treated black and white suspects. In essence, Reiss collected observational data on a number of police-citizen encounters which were as similar as possible except that some suspects were black and some were white. Similar approaches have been advocated by Glaser (1965).

To date, however, field work researchers have not fully capitalized on the longitudinal potential of their techniques. Analyses tend to emphasize cross-sectional comparisons in which events in one setting are contrasted to events in another, despite the fact that frequently the sources of such contrasts are sets of chronological narratives. Here we propose to build directly on ability of on-the-scene observations to capture longitudinal sequences of events.

In essence, the kinds of data necessary for the application of optimal control theory could probably be collected by "simply" attaching an observer to police as they go about their business. One could imagine the

observer riding in police cars or walking a beat with police officers and recording when various events occured. Perhaps the most straightforward procedure would rest on first noting the event and the time (actually the time at which it began and the time at which it ended) and both could later be transcribed in the form of a chronological log of events. Roughly similar approaches have been successfully employed by several students of police behavior (e.g. Reiss, 1971; Skolnick, 1966; Rubinstein, 1973), although all have stopped short of subjecting the data to a quantitative dynamic analysis. Note that one would not be limited only to what occured when, and a variety of other sorts of data could be collected (e.g. the number of citizens observing a police-suspect encounter, the location of the event, the weather, etc.)

At first blush, one might imagine that the presence of observers would substantially alter the ways in which police behave; normal activity patterns would be consciously adjusted to what police believe observers should be allowed to see. While field work critics have indeed raised such questions under the label of "reactivity", considerable experience indicates that after some initial awkwardness, police (and other kinds of subjects) "habituate". Observers are in fact a very small part of the environment in which police operate, and the regular forces affecting the day-to-day behavior of police soon re-assert themselves. Reiss (1968), for example, was easily able to observe a large number of police "brutality" incidents which police would have no doubt preferred to hide from external scrutiny. In short, within a reasonable interval, observers either fade in the background or are inducted into the "club" and treated as insiders. Given all of the other things police have to think about on patrol and large blocks of habitual behavior, observers appear to have little impact.

(See Manning, 1972, for a review of observational issues in studying police)

What might a typical data set generated through observational methods look like? Imagine a log with times arranged down the left margin, perhaps at one minute intervals. Across the top of the log would be variable names including such things as the event, the starting time, the ending time, and the location (variously defined). Later, a range of other variables could be appended such as weather conditions, the day of the week, and most important, policy relevant indicators (e.g. one person or two person patrol cars, weapons carried by police officers, stop-and-frisk guidelines). The crossing of times by variables would result in a matrix of time bound observations in which some variables would vary over time and some would not. Also, some variables would be nominal (e.g. whether the "event" was a motorist running a red light, whether police policies allow the use of mace) and some would be equal interval (e.g. time, temperature). To illustrate, the log entry for 8:45 PM might show (reading horizontally across variables) "used mace on resisting suspect, raining, 64 degrees, 1500 block Main Street, light traffic, crowd present, two officers, Monday, July 21st." If a police department were interested in the deterrent value of mace, such information across a number of incidents would be of great use. Inserted in a dynamic model of police patrol behavior, one could in principle not only describe what typically occurred, but also provide recommendations in an optimal control context on when mace should be used. One might find, for example, that in the presence of crowds (a stochastic exogenous event), mace is less effective on individual suspects because suspects feel compelled to visably demonstrate their bravery. Hence, one might recommend that mace not be used with crowds present.

Note that use of observationally generated logs circumvents many of the difficulties described earlier. In particular, many of the more troubling problems involved in sampling time are largely eliminated. The log provides a regular, discrete metric on which events can be located as they occur. More dense periods would not require any alteration in data collection methods, but would simply appear as periods of more dense entries on the log.

On the other hand, the practical use of observer generated logs is far from trouble-free. For example, one needs a way to code "events" so that they may be organized into broader categories. Clearly one cannot permit each distinct event to be uniquely represented, but at the same time, one would not want to collapse events into categories which significantly distort their meaning. Thus, one would probably choose to keep arrests for auto thefts distinct from arrests for burglaries, but the difference between arrests for auto thefts by makes of cars would likely swamp any serious analytic effort. However, such decisions clearly rest on the kinds of issues raised earlier and the policy variables of interest. As we have emphasized throughout, different systems applied to different substantive concerns will require alterations in one's analytic strategy.

A second complication involves the precision needed in events initially recorded on the log. For example, one would probably want to treat the apprehension of a suspect as an "event", but what about the actions that compose the apprehension? One could imagine recording a sequence something like "spotted suspect, took chase on foot, cornered the suspect, wrestled suspect to the ground, applied handcuffs, led suspect to waiting patrol car." Should each of these be recorded as distinct events or would the overall event "apprehended suspect" suffice? Again, there is no answer without first considering the system under study and the policy variables of interest.

If, for instance, one wished to model the broad impact of patroling on foot versus patroling in cars, perhaps "apprehension" is sufficiently precise. In contrast, if one wanted to model the impact of one person versus two person patrols on the apprehension of resisting suspects, more detail is probably required.

Given these and other likely difficulties, one might wonder if the data collection procedures we are suggesting are really feasible and useful. Our response is cautiously affirmative. We have had considerable experience (Berk and Berk, 1979) developing dynamic models for household activities using data much like suggested here. Indeed, those data, based on <u>self-report</u> diaries, are in many ways weaker than the data we propose collecting for police. Similarly, urban geographers (Lenntorp, 1976) have recently been modeling the movement of people in time and space using chronological reports of people's behavior. (They have also developed a number of interesting descriptive representations of time-space movement.) Finally, our suggestions should not be viewed as anything but a general description of how one might in principle build dynamic models of police behavior. Clearly, a great deal of additional thought is required, especially coupled with small, pilot studies in which first hand experience may be obtained.

VII. CONCLUSIONS

Given all of the difficulties we have outlined, we would argue against any large scale effort to apply optimal control theory to police behavior. On the other hand, the time may be right to attempt a number of small scale studies employing one of two simplification strategies. First, it

may be feasible to focus on some relatively narrow dimensions of police behavior and apply optimal control theory at a very micro-level. For example, it may be possible to build dynamic models of police-citizen encounters lasting little more than 30 minutes. Observational data might be adaquate, and the short time interval would in practice hold a range of confounding factors constant. If one were able to place observers in a small number of patrol cars for several weeks, for instance, sufficient data of the requisite quality might well be generated. Second, one might attempt to work with models at a much higher level of aggregation if the phenomena in question were relatively simple. For example, one could imagine a dynamic model of the "migration" of crime away from heavily patroled neighborhoods, with the idea of developing optimal reallocation policies to deal with this phenomenon. Clearly, however, even this would require a closely controlled data gathering effort over a longer period of time. It is unlikely that existing data sources could provide all of the requisite information needed to derive optimal policies.

Finally, however, we would like to argue that the adoption of control theory as a conceptual framework for analysis and data collection has certain intrinsic merits. If the process of police behavior is viewed as one of discrete time, dynamic decision making, then the foundation has been laid for improving data collection and ultimately, the performance of the system. Only by viewing the system in a way which reflects the logic of decision making in that system can we come to an understanding of the processes we would like to improve.

FOOTNOTES

1. Some special cases of this sort of system have been considered in which individual agents are responsible for specific kinds of decisions, ie. agent i is responsible for controlling the ith policy variable. This is called the assignment problem and has been discussed by McFadden 1969.
2. The model as written, is what is referred to as the reduced form of a system of simultaneous equations. Even though there may be contemporaneous interactions among the elements of y_t, any system with such interactions can be represented in terms of a reduced form model. (See Theil, 1971)

REFERENCES

Aoki, M.

1967 Optimization of Stochastic Systems. New York: Acamemic Press.

Berk, Richard A. and Sarah F. Berk

The Organization of the Household Day. Beverly Hills: Sage Publications, forthcoming.

Box, G.E.P. and G.C. Tiao

1975 "Intervention Analysis with Applications to Economic and Environmental Problems", Journal of The American Statistical Association, 70(1).

Bryson, A.E. and Y.C. Ho

1969 Applied Opitmal Control. Massachusettes:Ginn & Co., Waltham.

Campbell, D.T. and J. Stanley

1963 Experimental and Quasi-Experimental Designs for Research. Chicago:Rand McNally.

Chow, G.

1975 Analysis and Control of Dynamic Economic Systems. New York:Wiley.

Glaser, B.G.

1965 "The Constant Comparative Method of Qualitative Analysis", Social Problems, 12(3).

Junker, B.H.

1960 Fieldwork: An Introduction to the Social Sciences. Chicago: Univ. of Chicago Press.

Kelling, G.L., T. Pate, D. Dieckman and C.E. Brown

1974 "The Kansas City Preventive Patrol Experiment", in S. Halleck et. al. (eds) The Aldine Crime and Justice Annual, Chicago:Aldine.

Lenntorp, B.O.

1976 Paths in Space-Time Environments:A Time-Geographic Study of

Movement Possibilities of Individuals. Sweden:The Royal University of Lund.

Letkemann, P.
1973 Crime as Work. Englewood Cliffs:Prentice-Hall.

MacRae, E.C.
1972 "Linear Decision with Experimentation", Annals of Economic and Social Measurement 1(4):437-448.

Manning, P.
1972 "Observing the Police: Deviants, Respectables and the Law" in J.D. Douglas (ed) Research on Deviance. New York:Randon House.

Marshak, J.,
1955 "Elements for a Theory of Teams", Management Science 1:127-237.

Marshak, J. and Radner, R.
1972 The Economic Theory of Teams, Cowles Foundation Monography 22 New Haven:Yale Univ. Press.

McCall, G.J. and Simmons, J.L.
1969 Issues in Participant Observation. Reading, Mass.:Addison-Wesley.

McFadden, D.
1969 "On the Controllability of Decentralized Macroeconomic Systems: The Assignment Problem", in H.W. Kuhn and G.P. Szegö (eds.) Mathematical Systems Theory and Economics I, Springer-Verlag, Berlin, pp. 221-240.

Pindyck, R.S.
1972 "Optimal Stabilization Policies via Deterministic Control" Annals of Economic and Social Measurements, 1(4):385-390.

Prescott, E.C.
"The Multiperiod Control Problem Under Uncertainty", Econometrica, 40:1043-58.

Reiss, A.J., Jr.
1968 "Police Brutality-Answers to Key Questions", Transaction, 5(4).

Reiss, A.J., Jr.
1971 The Police and the Public. New Haven:Yale University Press.

Rubinstein, J.
1973 City Police. New York:Farrar, Straus and Giroux.

Simon, H.A.
1962 "The Architecture of Complexity", Proc. American Philosophical Society 106:467-82.

Skolnick, J.H.
1966 Justice Without Trial. New York:John Wiley.

Theil, H.
1971 Principles of Econometrics. New York:John Wiley.

Tse, E.
1974 "Adaptive Dual Control Methods", Annals of Economic and Social Measurement. 3(1):65-84.

Tse, E. and Y. Bar-Shalom
1972 "An Actively Adaptive Control for Linear Systems via the Dual Control Approach", IEEE Transactions on Automatic Control. April:109-116.

END