Donut Shops, Speed Traps, and Paperwork: Supervision and the Allocation of Time to Bureaucratic Tasks

John Brehm, Scott Gates, and Brad Gomez

Supervisory Roles and Subordinate Compliance

Principal-agent models of hierarchical control (e.g., Alchian and Demsetz 1972; Holmström 1982) emphasize the coercive elements of supervision as the means of inducing subordinate compliance. In this view, supervisors extract work from subordinates by rewarding good workers, that is, those who comply with the supervisor's preferences, and punishing inadequate performers (Kadushin 1992). The principal's job is to anticipate the rational responses of agents and to design a set of incentives such that the agents find it in their own interest (given the incentive system) to take the best possible set of actions (from the principal's perspective). These models assume that the principal's ability to induce compliance is constrained by asymmetric information regarding the agent's effort (moral hazard) or ability (adverse selection).

As we have previously found (Brehm and Gates 1997), the coercive abilities of supervisors are limited by more than informational asymmetry. We have found that the coercive capacity of most supervisors is much less important as an influence over subordinate compliance than the subordinates' preferences, the selection process, intersubordinate learning, and contacts with clients (Brehm and Gates 1993, 1997). The demonstrated weakness of supervisory coercion leads us to reexamine the nature of the supervisor-subordinate relationship. If we dispense with the coercive capacities of supervisors, does this leave other supervisory functions as potential levers on subordinate compliance?

Organization theorists writing about hierarchical control have long suggested that the supervisor's role in inducing compliance is not limited to coercion. Chester Barnard (1938) suggested that supervision is a combination of potential functions. Supervisors provide a system of communications, promote the securing of essential efforts, and formulate and define purpose (217). The coercive aspect of supervision is only one means of achieving these functions. Supervisors may act as "teachers" in helping subordinates to identify more efficient ways to accomplish what is in both of their interests. In exhorting subordinates to identify with the mission of the organization, a supervisor may take on the role of "preacher" (an elaboration of Kreps [1984] 1990). A supervisor is a "recruiter" when he or she hires subordinates who share his or her preferences for work (Carpenter 1997; Johnson and Libecap 1994). A "participant" supervisor is also a contributor to work. Supervisors also provide "support," providing political cover so that subordinates can work. Finally, supervisors must execute the role of "coordinator," applying sanctions and rewards to encourage greater work on more productive tasks. Obviously, this list of supervisory roles is not complete and some are potentially contradictory. However, we believe these roles point to an expanded view of supervisors, their tasks, and their ability to persuade subordinates.

The principal contribution of this essay is to examine the supervisor as coordinator within a public bureaucracy, a role that is consistent with both principal-agency approaches and organizational theory. In this role, the supervisor must define and allocate tasks across subordinates (Wilson 1989). Tasks define what it is that bureaucrats, and hence bureaucracies, do. The challenge for the supervisor as coordinator is to match the right subordinates with the right tasks. This requires that subordinates be given tasks that they prefer. However, a perfect match is doubtful, since a subordinate is likely to receive a bundle of tasks, some of which are preferred and some not. What is the capacity of supervisors to coordinate subordinate work across a variety of tasks?

The change in our approach from the simple working/shirking dichotomy to a range of tasks offers many advantages. One advantage is that by looking at specific tasks rather than cumulative time spent in "working" we can permit a much more nuanced understanding of what it is that public bureaucrats do. Our previous analysis of working and shirking by police officers required that we collapse together all tasks that lead toward production when there is obviously considerable variation in the form of those tasks. It is of considerable interest to the supervisor to know whether his or her subordinates are devoting the majority of their time to mobile dispatches at the expense of time spent on paperwork. Further, the notion of "tasks" as the unit of bureaucratic effort squares much more cleanly with organizational theory on routines (e.g., Steinbruner 1974) and bureaucratic politics literature (e.g., Wilson 1989). Finally, by switching to an analysis of time spent among different forms of working, we reduce the sensitivity of our results to the potential for effects of the observer upon the performance of the subordinate. That is, while the amount of time spent working probably increases when the subordinate is being observed, the amount of time allocated to any particular form of work is probably less sensitive to observation.

We expand upon the working/shirking dichotomy by examining the compliance of public bureaucrats across multiple tasks, and we choose police officers as the bureaucrats to study. By examining the amount of time a police officer allocates to specific tasks, we have an opportunity to evaluate the supervisor's role in both maintaining subordinate compliance and coordinating work effort. Based upon our previous models of supervision, we believe that subordinate work across tasks will be a function of supervisory coercion, subordinate preferences, and the solidary attachments of subordinates. Although the findings here relate specifically to police work, we believe that they are generalizable to the behavior of bureaucrats in other contexts. In our previous analyses (Brehm and Gates 1997), we found strong similarities across a wide variety of bureaucrats, from police officers to social workers to federal civil servants.

We proceed in three parts. We begin with an explication of propositions derived from two of our previously published models that attempt to explain why subordinate bureaucrats behave the way they do and who influences them. The first of these models, the enhanced principal-agent (EPA) model, recasts agents' efforts toward the production of output by assuming that the utility of work for individual agents varies across tasks. This model assumes that agents' behavior is influenced by a broader set of incentives than simply the flow of residuals. A second model, an imitative model, draws on social psychological theories of organizational compliance in order to develop a model of intersubordinate learning.

Next we present our data set and explicate a statistical model for the analysis of time allocation across tasks. We then present the findings from the estimation of our model of supervision. Finally, we discuss the implications of our results for understanding both the analysis of time allocation and, most importantly, the theory of bureaucratic supervision.

Supervisors as Agency Coordinators

James Q. Wilson writes that "People matter, but organization matters also, and tasks matter most of all" (1989, 173). However, we know little about how tasks are allocated in an organization. What role do supervisors play in this allocation? In what capacity do supervisors serve as coordinators and facilitators of subordinates' work on different tasks? The traditional principal-agent model features agents' decision to work or shirk. (In Brehm and Gates 1997, we expanded the options available to subordinates to include sabotage, the act of producing negative output.) In this essay, we consider variation among the tasks available to subordinates, where the tasks vary in desirability to both the supervisor and to the subordinate. The key question is similar to that posed in our prior analysis: what accounts for the amount of time that subordinates devote to each task? We study the role of supervisors as coordinators by looking through the lenses of our two models, the enhanced principal-agent game and the imitative model, which are summarized here and presented elsewhere in greater detail (Brehm and Gates 1994, 1997).

The EPA Game

The EPA game models the strategic interactions between a supervisor and subordinates in order to provide greater insights into the problem of supervision and compliance in a bureaucratic setting. In terms of task allocation, the game begins with a supervisor deciding which subordinates should work on what policies. Every subordinate is given several assignments to maximize the supervisor's production goals. To do this, the supervisor attempts to match the "best" person to each task. In turn, each subordinate decides how he or she will allocate his or her time across these assigned tasks. Supervisors then must determine how and whom to supervise, given limited supervisory resources and the fact that some subordinates respond to supervision and some do not. The next time a supervisor decides who will do what, he or she takes into account how subordinates allocated their time and how responsive they were to supervision. A more formal overview of the EPA game is presented in our book (Brehm and Gates 1997, chap. 2).

One of the central assumptions of the EPA model is that the marginal effect of supervision on subordinate work inputs is a function of how amenable a subordinate is to supervision and the difference between the amount of work input by the subordinate toward a particular task and the amount of work input desired by the supervisor. (We assume that supervisors desire at least as much work from their subordinates as a subordinate would provide if left unsupervised.) Amenability is assumed to vary from subordinate to subordinate, but since we are not modeling responsiveness to supervision dynamically, we treat amenability to be given for an individual.[1]

To identify the equilibria for the EPA game, we explore the two main cases evident in the game. The first case arises when the difference between work desired by the supervisor and work provided without supervision is less than the marginal cost of supervision. This situation arises whenever the subordinate and the supervisor agree as to how much time to invest in a particular task, whether a great or little amount. This equilibria is pooling on type (meaning that there is no way to differentiate responsive from unresponsive subordinates).

The second case arises under the opposite condition, namely, when the difference between work desired and unsupervised work is greater than the marginal cost of supervision. If this condition holds, we can identify separating equilibria on type. Supervisors are able to identify responsive and unresponsive subordinates. If the subordinate is sufficiently unresponsive, no supervisory time is allocated. In other words, a supervisor will not waste time supervising an "insubordinate" subordinate, since the marginal costs exceed the marginal benefits associated with bureaucratic outputs.[2] Responsive subordinates, on the other hand, are supervised. Given that the marginal benefits of bureaucratic output are greater than the marginal costs, it is in a supervisor's interest to supervise a responsive subordinate.

From these equilibria generated by the EPA model we are able to derive the following testable propositions regarding subordinate work across tasks.

- If subordinates are indifferent between tasks (forms of work) and are homogeneously responsive to supervision (across tasks), then subordinates allocate greater work to tasks that receive greater supervision.

- If subordinates are not indifferent between tasks and are homogeneously responsive to supervision, subordinates allocate time to preferred activities, (largely) independent of supervision.
- If subordinates are indifferent between tasks but are heterogeneously responsive to supervision, subordinates allocate time to tasks for which they are more amenable to supervision.
- If subordinates are not indifferent between tasks and are heterogeneously responsive to supervision, subordinates will separate according to type. Nonresponsive subordinates will allocate time to preferred activities, and responsive subordinates will allocate time to activities for which they receive supervision.

The Imitative Model

The imitative model traces its origins to social psychological models of compliance and persuasion. Two central concepts are emphasized, social proof and consistency. Both concepts come into practice whenever a subordinate faces a request from a supervisor. In terms of consistency, the subordinate bureaucrat will ask "what have I done in the past?" As for social proof, the subordinate asks "what are others like myself doing?" If the subordinate is facing an uncertain or ambiguous situation, imitation serves as a simple and direct information shortcut.[3] These two central concepts are used to model how subordinates learn to respond to supervision by imitating one another.

Several factors influence the final distribution of work allocations by subordinates after many iterations of the simulated model. The adaptation of the bureaucrats is nonlinear and contingent on four discrete sets: the available set of responses to a policy, responses at first iteration, disposition toward a policy, and connections among bureaucrats and supervisors. We also consider supervisory tolerance for noncompliance and the level of sanction he or she applies to shape subordinate performance.

The imitative model simulation offers several conclusions for understanding task assignment and the role of a supervisor as coordinator. The central finding is that the more individual subordinates look to fellow subordinates for information about how to respond to a rule the greater the degree of conformity between subordinates. In turn, several factors influence the degree to which subordinates imitate one another, including the observability of subordinate actions, the level of uncertainty subordinates face, the frequency of contact between subordinates, the policy predispositions of subordinates (both the mean and variance), and the utility subordinates derive from working on a particular task. From these relationships, we are able to derive several propositions regarding the allocation of subordinate time across tasks.

- As the observability between subordinates increases, conformity of allocation of time across tasks increases (due to either information search or coercion by subordinates).
- As supervisors are better able to observe subordinate activity, subordinates allocate greater time to supervisors' preferred tasks (but see the "Data" section, which follows).
- Subordinates are more responsive to functional and solidary preferences than they are to supervisory coercion.
- As the number of tasks increases, uncertainty increases, and subordinates increasingly conform.

Data

The specific data at hand come from the 1977 Police Services Study, conducted by Elinor Ostrom, Roger Parks, and Gordon Whittaker in three cities (Rochester, St. Louis, and St. Petersburg). The study combined multiple methods, including observations of police officers' behavior during their shifts. The observational data provide an excellent opportunity to test our propositions about the allocation of time across tasks.[4] At the conclusion of each shift, the observer recorded the amount of time officers spent on a total of eleven tasks (the italicized phrase or word denotes our label in subsequent graphs).

1. Time spent on administrative duties (*Adm*)
2. Time spent on report writing (*Rept*)
3. Time spent out of car for foot patrol (not on an encounter or dispatched run) (*Foot Pat*)
4. Time spent on routine mobile patrol (*Mob Pat*)
5. Time spent at or en route to an encounter or dispatched run (*Run*)
6. Time spent on mobile traffic work (radar, vascar, etc.) (*Mob Traf*)
7. Time spent on stationary traffic work (radar, etc.) (*Stat Traf*)

8. Time spent on meals and other 10–7 breaks (*Meal*)
9. Time spent on mobile personal business (*Mob Pers*)
10. Time spent on stationary personal business (*Stat Pers*)
11. Time spent on other stationary police work (surveillance, stake-out, etc.) (*Other*)

We apply two different means for examining the amount of time an officer devotes to different tasks. The first of these, the "ternary" diagram, is most useful when one collapses the distribution of time across tasks into three categories. Here we consider time spent on personal business (meals and stationary and mobile personal business), time spent completing paperwork (administration, reports), and time spent policing (mobile and stationary traffic, runs, mobile and foot patrol, and other). The collapsing of time into three tasks corresponds nicely with a division into a police officer's principal responsibilities (policing and paperwork) plus a category denoting time not devoted to responsibilities. In our previous analysis (Brehm and Gates 1993, 1997), we facetiously referred to these as "donut shops" (shirking, here measured as time spent on personal business) and "speed traps" (working). In the present analysis, we divide time spent working between the categories of policing (speed traps) and paperwork.

If the amount of time spent on tasks is transformed into percentages of total time and total time is constrained to sum to 1, then the data are arranged on what is known as a simplex. One could produce a three-dimensional scatterplot of the data across the three dimensions of tasks, but all of the points would fall on the triangular plane intersecting the three axes at 1.0 (fig. 1). Instead, we focus solely on the triangular plane displayed in figure 2.[5]

The figure makes it quite clear that the majority of these officers' time is devoted to policing. The mode of the distribution is quite close to the extreme lower right corner, although there is a fair amount of dispersion throughout the lower right trident of the ternary diagram. Only five officers spent a plurality of their time on personal business, running counter to stereotypes about police behavior. Eight officers devoted a plurality of their time to paperwork, including one officer who spent the entire shift on paperwork. There are also some interesting edge conditions—officers who divided their time between either policing and paperwork or policing and personal business.

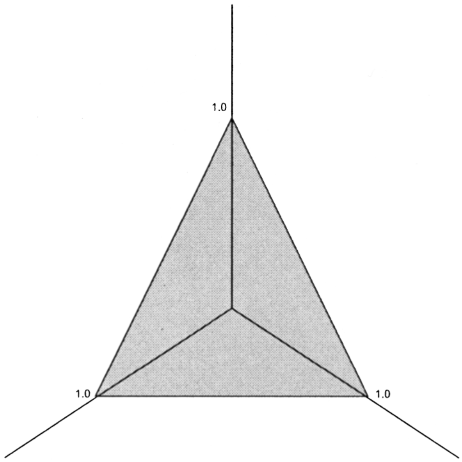

Fig. 1. Simplex for three dimensions

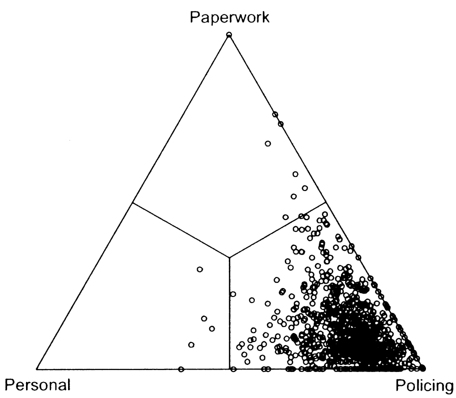

Fig. 2. Actual distribution of time on simplex, 1977 police data

The second graphical display (fig. 3) involves use of a novel technique called the checkerboard plot. Each officer is displayed as a vertical column of rectangles (here, quite thin—nearly lines—since we need to display more than nine hundred officers' shifts). Each row of rectangles corresponds to one of the eleven tasks (e.g., mobile patrol or meals). We shade each rectangle with a percentage of gray to denote the amount of time devoted by the officer to that task: rectangles that are completely

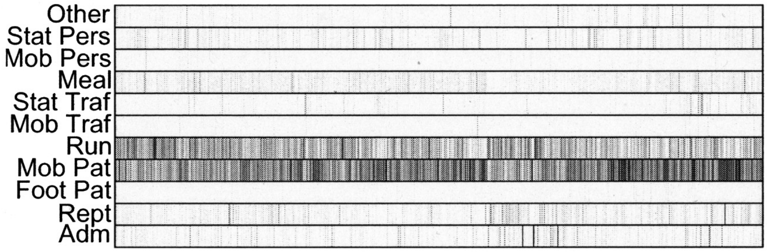

Fig. 3. Checkerboard plot of actual distribution of time spent on simplex, 1977 police data

white denote those tasks on which an officer spent zero time; rectangles that are completely black denote those tasks to which the officer devoted his or her entire shift; and those that are gray denote those tasks on which the officer spent some middling fraction of time. The darker the gray the more time was devoted to the task.

As is readily apparent from the checkerboard plot, police officers spend the majority of their days confined to two tasks: mobile patrol and on route to an encounter. Officers spend the least amount of their time on foot patrol and mobile and stationary traffic. Officers spend middling amounts completing reports or performing other administrative duties as well as on meals or stationary personal business. (The meals category is in third place, on average, although it is distantly behind runs and mobile patrol).

As is also apparent, these patterns are strikingly homogeneous across the more than nine hundred police officers in the three different cities. Although one can identify individuals who devote a plurality of time to administration and reports (the dark lines in those sections of the plot), as well as those who spend nearly twice as much time at meals than other officers, the general pattern here is one of uniformity, not variation.

This pattern of results is consistent with several of the conditions of the EPA and imitative models. The strong pooling equilibrium at work in the three cities (convergence on runs or mobile patrol) could be produced under the first and third hypotheses of the EPA model: it could be that the subordinates are either largely indifferent between tasks or fairly responsive to supervision. The mean behavior is also consistent with the logic of the imitative model, which generates a high degree of conformity as a result of intersubordinate contact and high solidary preferences.

To test the propositions of the EPA and imitative models, we require the following measures as explanatory variables. We need measures of the amount of supervision supplied. Here we use the number of contacts with the supervisor as the measure. Although we do not know how the supervisor is spending his or her time with the subordinate, the amount of contact is an excellent proxy for the amount of time available for supervision. (The EPA model explicitly treats supervisory time as a constrained budget, while the imitative model does not.) Note also that our use of the measure of the frequency of contact does not explicitly capture the duration of contact. One could reasonably suppose that the least responsive bureaucrats receive not only more frequent but longer periods of contact with their supervisors than more responsive bureaucrats do.

In addition, we include a dummy variable for radio contact to denote how observable the officer is to the supervisor. The imitative model explicitly models the supervisor's ability to observe subordinate output (although it predicts that observability is only weakly related to increased effort). We also need measures of the subordinates' preferences for work. Both the EPA and imitative models argue that subordinate preferences directly affect the amount of time spent on work. Here we use whether the subordinate expressed "functional" likes or dislikes (specifically commenting in the positive or negative about the responsibilities of officers). The imitative model argues that intersubordinate contacts influence performance. We employ a count of the number of contacts with fellow officers and whether the officer expressed a "solidary" like or dislike.

Inferential Methods for Analysis of Time Allocation

Suppose that one has a record for each bureaucrat i of the proportion of total time that the bureaucrat spends on each task j, denoted y_{ij}. Suppose further that the list of tasks $1 \ldots J$ are mutually exclusive and exhaustive and that a minimum amount of time is spent on each task. By definition, then, both of the following hold.

$$y_{ij} > 0, \qquad \forall j = 1 \ldots J \tag{1}$$

$$\sum_{j=1}^{J} y_{ij} = 1. \tag{2}$$

These two features mean that the allocation of time across tasks y_{ij} constitutes a simplex. Mathematical features of a simplex will mean that the distributions of the y_{ij} are not fully independent. For example, if one knows the values of y_{i1} up through $y_{i(J-1)}$, then one knows the value of y_j. In the simplest case, let y_{i1} describe the amount of time spent working; if "shirking" constitutes any time that is not spent working, then $y_{i2} = 1 - y_{i1}$. When the number of tasks increases beyond 2, the relationship is no longer as immediately straightforward—increased time spent on task 1 means less time for the remaining tasks but not necessarily any one particular remaining task—but it is still constrained. Similar problems are evident in any situation in which a variable is characterized by multiple outcomes that sum to unity for each observation.[6] Problems in this general class are referred to as "compositional data analysis" (Aitchison and Shen 1980; Aitchison 1986; Katz and King 1999). We consider one main variant of compositional data analytic strategies, the Dirichlet.

One relatively simple solution begins from an assumption that each stream of tasks is produced by an independent process. Suppose y^*_{ij} represents the hours in a week devoted by bureaucrat i to task j and that y^*_{ij} is distributed as J independent gamma random variates with shape parameters $\nu_1 \ldots \nu_J$. The probability density function for the gamma distribution is

$$y^*_{ij} = f_\gamma\,(Y^*_{ij}|v_j) \tag{3}$$

$$= y_{ij}^{v_{j-1}}\frac{\exp(-y_{ij})}{\Gamma(\nu_j)}. \tag{4}$$

The mean and variance for the gamma pdf are both ν_j, and when the shape parameters are integer the distribution is also known as the Erlang distribution.

The total hours in the week is $T_i = \Sigma y^*_{ij}$. The proportion of time devoted to each task is then $y_{ij} = y^*_{ij}/T_i$. The proportion of time devoted to each task is distributed according to a Dirichlet distribution.

$$(y_1 \ldots y_J) = f_D(Y_1 \ldots Y_J|\nu_1 \ldots \nu_J) \tag{5}$$

$$= \frac{\Gamma(\Sigma_{k=0}^{J}\,\nu_k)}{\Pi_{k=0}^{J}\Gamma(\nu_k)}\prod_{k=1}^{J} y_k^{\nu_{k-1}}, \tag{6}$$

where

$$\nu_j > 0, \quad \forall j = 1 \ldots J.$$

One can reparameterize the ν_j in terms of explanatory variables and coefficients with simple exponentiation.

$$\nu_j = \exp(X\beta_j),$$

where the effect parameters (β_j) vary by task and the X may or may not be the same set of explanatory variables (identification for the system is accomplished through covariance restrictions, detailed later, and through functional form). If one assumes that the observations are distributed identically and independently, then the log-likelihood for the reparameterized Dirichlet is

$$\ln L(\beta|X,y) = \sum_{i=1}^{N}\left[\ln\Gamma\left(\sum_{j=1}^{J} e^{X\beta_j}\right) + \sum_{j=1}^{J} e^{X\beta_j} \ln y_j - \sum_{j=1}^{J}\ln\Gamma\left(e^{X\beta_j}\right)\right].$$

(Note that we assume that intersubordinate influence is entirely captured by the intersubordinate contact and solidary preference measures in order to sustain the independence part of the IID assumptions). This log likelihood is easily optimized with a statistical package such as Gauss.

Several features of the Dirichlet lend themselves to some desirable properties for purposes of interpretation. The Dirichlet is a multivariate generalization of the Beta distribution (which we use extensively in our analysis of the allocation of time across two "tasks" (working and shirking) in Brehm and Gates 1997. As such, it is a highly flexible distribution permitting multiple modes and asymmetry. Further, the moments are easily found. Let $\nu^* = \sum_{k=1}^{j} \nu_k$. The mean of the amount of time spent on task j is

$$\mu_j = \frac{\nu_j}{\nu^*}. \tag{7}$$

The variance of time spent on task j is

$$\text{var}(y_j) = \frac{\nu_j(\nu^* - \nu_j)}{\nu^{*2}(\nu^* + 1)}, \tag{8}$$

and the covariance of time spent on tasks k and m is

$$\text{cov}(y_k, y_m) = \frac{-\nu_k\nu_m}{\nu^{*2}(\nu^* + 1)}. \tag{9}$$

Since all the ν_j are positive, this means that the covariance of time spent on any pair of tasks k and m is negative or that any increase (decrease) in time spent on one task necessitates a decrease (increase) in time spent on *every other* task.

This property of the Dirichlet distribution is the first sign that there are hidden assumptions in the Dirichlet that may warrant another selection of distributional assumptions. Aitchison (1986) writes:

> It is thus clear that every Dirichlet composition has a very strong implied independence structure and so the Dirichlet class is unlikely to be of any great use for describing compositions whose components have even weak forms of dependence. . . . This independence property, which holds for every partition of every Dirichlet composition, is again extremely strong, and unlikely to be possessed by many compositions in practice. For example, one implication of it is that each ratio x_i/x_j of two components is independent of any other ratio x_k/x_l formed from two other components. (60)

What remains to be seen, however, is just how sensitive the analysis of composite data is to this particular "strong" independence of irrelevant alternatives (IIA) assumption. As we will demonstrate, the Dirichlet estimates turn out to be quite adequate for general prediction of the amount of time the officers in this sample devote to the different tasks.

The irony is that the Dirichlet distribution, like the Beta distribution, is capable of considerable variation in potential distributions of allocation of the compositions. Figures 4–7 demonstrate simulated Dirichlet distributions for varying selections of the parameters. It is possible to generate, among other forms, Dirichlet distributions that are uniformly dispersed (fig. 4), unimodal and centered (fig. 5), unimodal and off center (fig. 6), or multimodal and skewed (fig. 7).

Our selection of the Dirichlet is motivated by two distinct rationales. One is that the flexibility of the Dirichlet and relatively easy optimization (even for many equation systems) allow us to explore the task allocation problem at a high degree of disaggregation. In this sense, we chose the Dirichlet because of its tractability and good fit to the data. The second motivation is that the Dirichlet can be thought of as an outcome of some k gamma processes. In a separate work, we document that one can view the task allocation problem as a cooperative game where the supervisor encourages subordinates to work on the basis of the allocation of

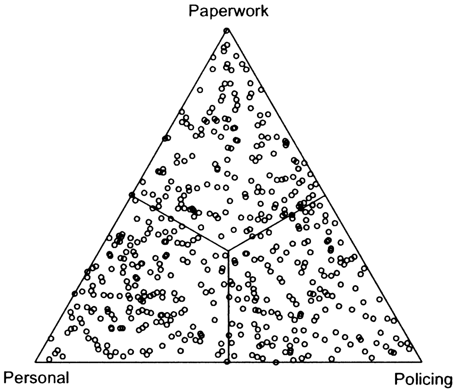

Fig. 4. Simulated Dirichlet, $\nu_1 = 1$, $\nu_2 = 1$, $\nu_3 = 1$

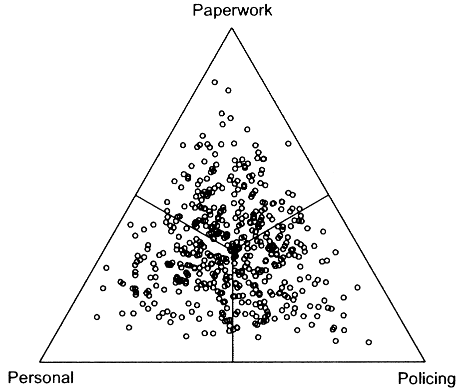

Fig. 5. Simulated Dirichlet, $\nu_1 = 3.5$, $\nu_2 = 3.5$, $\nu_3 = 3.5$

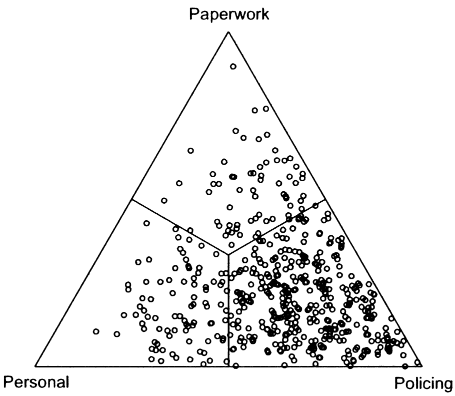

Fig. 6. Simulated Dirichlet, $\nu_1 = 1.5$, $\nu_2 = 1.5$, $\nu_3 = 3.5$

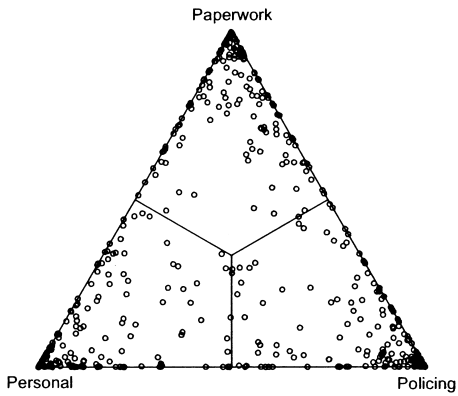

Fig. 7. Simulated Dirichlet, $\nu_1 = .25$, $\nu_2 = .25$, $\nu_3 = .25$

perks, represented as integer units (Brehm and Gates 1999). This stream of integer allocations would constitute a gamma process.

Results

Tables 1 and 2 display the maximum likelihood estimates for the Dirichlet distributions of the collapsed three-category and expanded eleven-category models, respectively. Positive signs on the coefficients indicate that an increase in the variable corresponds to an increase in the level of work. Frankly, more nuanced direct interpretation of the coefficients across even the simpler three-category model is cumbersome and across the eleven-category model even more so. Still, one should note that most of the variables are statistically significant for most of the equations, clearly so for the equations for "reporting" and "mobile patrol." Most of the time, the likes and dislikes are oppositely signed, which is sensible, although for the reporting equation all of the likes and dislikes are positive.

Because of the difficulties of providing more qualitative interpretation of the coefficients, our principal method for displaying the results of the Dirichlet analyses of time allocation is a form of computer simulation. We feature the Dirichlet analysis in this section. The process first requires that we generate estimates of the relevant parameters for each distribution based upon our estimated coefficients and selected values for the regressors and then generating vectors of random numbers drawn from the correct distributions with those parameters.

For the Dirichlet, the method works as follows (here using the mean values of the regressors for illustration). First compute the parameters ν_j^* for each task j from the equation

$$\nu_j^* = \exp(\hat{\beta}_{0j} + \hat{\beta}_{1j}\bar{x}_{1j} + \cdots + \hat{\beta}_{kj}\bar{x}_{kj}) \tag{10}$$

for each of the k regressors. Then draw 1,000 observations of y_j^* from the gamma distribution with shape parameter ν_j^*. Each of the y_j^* are independent. To scale the gamma variates to the simplex, simply divide by their sum.

$$y_j = y_j^* \Big/ \sum_{i=1}^{J} y_j^*. \tag{11}$$

TABLE 1. Dirichlet Estimates for Allocation of Time Devoted to Tasks (collapsed), 1977 Police Data

Variable	Paperwork	Policing	Personal
Constant	.51*	2.37*	.86*
Patrol contacts	.01	.02*	.02*
Supervisor contacts	.04*	−.00	−.02
Radio contact	.00	.02*	.01
Functional likes	−.22*	−.18*	.02
Functional dislikes	.03*	.09*	.02
Solidary likes	.02	.06*	.05*
Solidary dislikes	.05*	.16*	.14*

$N = 944$.
*indicates coefficient statistically significant at $p < .05$.

TABLE 2. Dirichlet Estimates for Allocation of Time Devoted to Tasks (expanded), 1977 Police Data

Variable	Adm	Stat Traf	Rept	Meal	Run	Other	Foot Pat	Mob Pat	Stat Pers	Mob Traf	Mob Pers
Constant	−0.49*	−1.06*	0.41*	−0.97*	−1.06*	−0.93*	−0.02	1.38*	−0.24*	0.87*	−0.36*
Patrol contacts	0.01	0.00	0.06*	0.00	0.04*	0.01	0.00	0.04*	0.03*	0.00	0.01
Supervisor contacts	0.02*	−0.01	0.08*	−0.01	−0.01	0.00	0.04	−0.04*	−0.07*	−0.02	−0.01
Radio contacts	0.01	0.03*	0.08*	0.03	0.02	0.02	0.01	0.09*	0.10*	0.03*	0.07*
Functional likes	−0.16*	−0.07*	0.20*	0.00	0.01	−0.06*	−0.25*	−0.11*	−0.01	−0.01	0.15*
Functional dislikes	0.02	0.05*	0.13*	0.02	0.03*	0.03	0.03	0.00	−0.07*	0.01	0.12*
Solidary likes	0.01	0.00	0.07*	0.02	0.00	−0.02	0.00	0.07*	−0.01	0.02	0.04*
Solidary dislikes	0.03*	0.03	0.08*	0.05*	0.03	0.05*	−0.02	−0.05*	0.05*	−0.01	0.06*

$N = 944$.
*indicates coefficient statistically significant at $p < .05$.

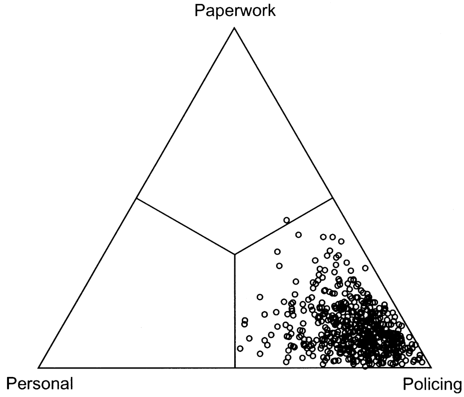

Fig. 8. Simulated Dirichlet at mean estimates

Figure 8 displays a simulated draw from the Dirichlet for the parameters computed at the mean. Clearly, this simulated distribution of time allocation comes quite close to the actual distribution (fig. 2). The mode of the simulated distribution falls approximately at the same location as the mode of the actual distribution (i.e., significantly skewed toward policing). The spread of the estimated distribution mimics the spread of the actual distribution, covering most of the lower-right trident. At the same time, the simulated distribution is missing some of the more striking features of the actual distribution. The simulated distribution does not capture the edge cases, neither the extreme outlier for time on paperwork, nor any of the cases that fall strictly between two tasks (either paperwork and policing or policing and personal business). Still, one would have to regard the Dirichlet distribution as one that replicates the actual distribution to a high degree.

Figure 9 presents the checkerboard plot of the simulated Dirichlet results, evaluated at the mean. While the ternary diagram leads one to conclude that officers spend the majority of their time policing, the checkerboard plot makes clear how officers spend time in the division of tasks within the category of policing.

As is consistent with the actual distribution of police officers' time-at-task, the simulation based on the mean values predicts that officers will spend the vast majority of their time on either runs or mobile patrol. Other forms of policing fall significantly behind. Unlike the actual distribution of time-at-task, the simulation based on means overpredicts the

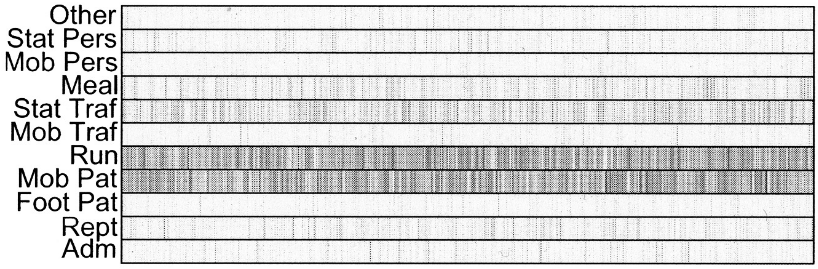

Fig. 9. Checkerboard plot of simulated Dirichlet at mean

amount of time that officers spend on static traffic duty. The Dirichlet estimates slightly underpredict the amount of time that officers spend on reports or administration, although it has the correct balance between the two forms of paperwork (i.e., more time is spent on reports). The model correctly predicts the amount of time that officers spend on meals, although it slightly overpredicts the amount of time spent on static personal business. Clearly, the Dirichlet estimates reproduce the actual distribution of time across the eleven tasks with great faithfulness.

The most interesting results will appear when we generate new simulations based upon selected values of the independent variables. Although we are not able to provide direct tests for the propositions developed here, we are able to examine the importance of supervisor observability (a key feature of principal-agency models is the extent to which action is hidden), subordinate observability (consistent with the ideas of the imitation model), and subordinate functional and solidary preferences. In the ternary plots of the simulated Dirichlets, we will show what happens at the maximum levels of the variables in question. With the checkerboard plots, we can simulate *increasing* levels of the variables, moving from left to right in the graphs.

What happens when the subordinate police officers are maximally observable to their supervisors? This situation ensues when the subordinates are in radio contact and when they have had the maximum number of contacts with the supervisor (twenty-one in this sample). (Note that we cannot ascertain the nature of the contact with the supervisor. This contact could have included specific requests for specific tasks to be completed, such as paperwork or policing; may have led to the imposition of formal or informal sanctions; or could have been strictly incidental.) Figures 10 and 14 display the simulated effect of being maximally observable.

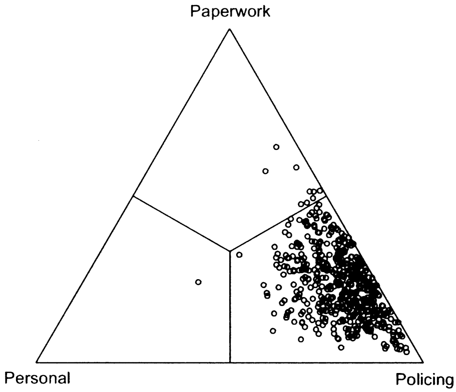

Fig. 10. Simulated Dirichlet when officers are most observable to supervisors

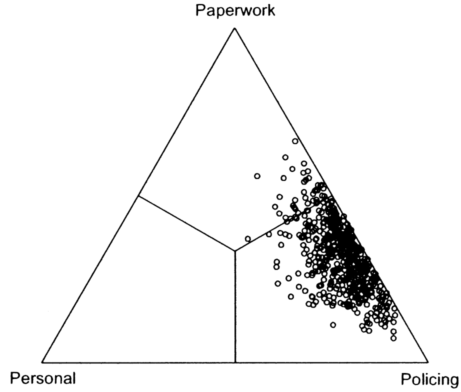

Fig. 11. Simulated Dirichlet when officers are most observable to subordinates

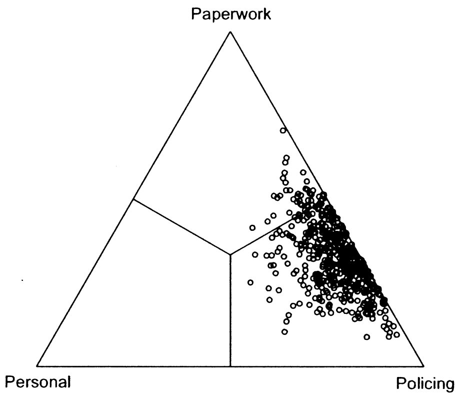

Fig. 12. Simulated Dirichlet when officers are most satisfied with squad

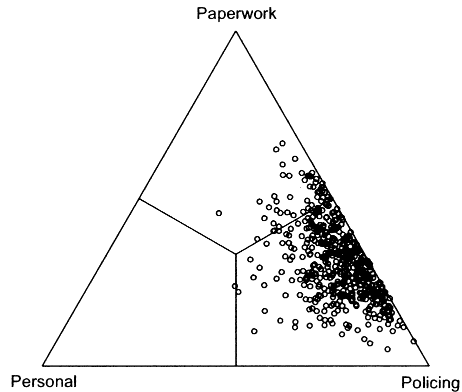

Fig. 13. Simulated Dirichlet when officers are most satisfied with job

The results of being under supervisory observation are hardly surprising. The variance of the distribution clusters away from time on personal business, suggesting a significant effect of supervisors on deterring shirking. Furthermore, the distribution of working tasks shifts away from policing and toward both paperwork and stationary traffic duty. Supervisors clearly have an influence on how subordinates allocate their time across policing and administration. Furthermore, as illustrated by the checkerboard plot, the greatest effect of increased supervision is upon in-

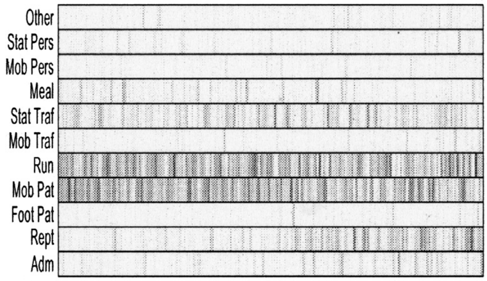

Fig. 14. Simulated Dirichlet when officers are most observable to supervisors

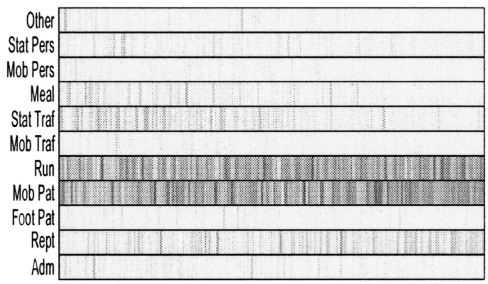

Fig. 15. Simulated Dirichlet when officers are most observable to subordinates

creased time spent completing reports, not on administration, and a significant reduction of the amount of time spent on personal business, meals especially.

The more complicated checkerboard plot reveals that supervisors have a significant role in encouraging subordinates to distribute time over a wide variety of tasks. In contrast to the simple work-shirk split, the supervisors exercise a high degree of influence over the subordinates' choices among different forms of work.

A similar figure may be produced to display the effect of being observable to fellow subordinates (figs. 11 and 15). In the present analysis, this entails raising the number of contacts with fellow officers to its maximum (thirty-one). As with the figure for officers maximally observable to supervisors, there is a pronounced shift toward increased time spent on paperwork. In fact, not only is there a greater increase in the amount of time spent on paperwork, but the dispersion of the distribution is

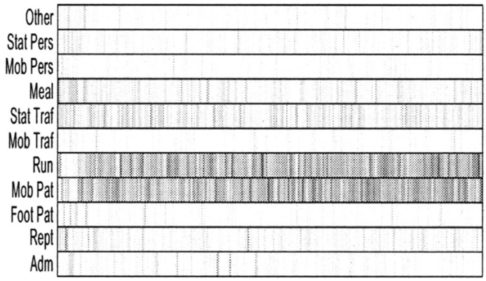

Fig. 16. Simulated Dirichlet when officers are most satisfied with squad

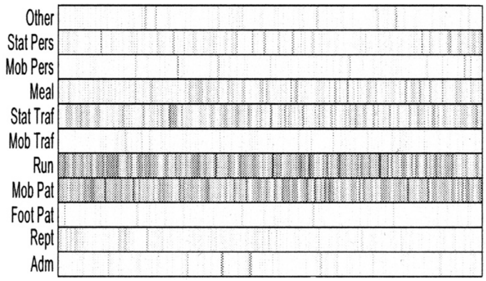

Fig. 17. Simulated Dirichlet when officers are most satisfied with job

noticeably more concentrated. This concentration is especially apparent in the checkerboard plot (fig. 15). There is an additional concentration that was somewhat obscured in the ternary plot: subordinates who have a great deal of contact with each other are much less likely to spend time on meals, which one would have to classify as a form of shirking.

Why would we see such strong effects? We think there are two classes of explanation. One follows from the first proposition, derived from the imitative model: under greater subordinate observability (presuming conditions of uncertainty about appropriate time allocation), subordinates are more likely to conform in their behavior. This explanation could account for the lesser dispersion but not the shift of the mean, however. The second explanation depends on a variety of the collective goods problem. Ask any bureaucrat, and he or she will tell you that the least desirable part of his or her job is completing paperwork. The police officers' condition not only requires completing reports on each meaningful activity during the day but that the task be completed for each partnership of officers. If an officer fails to complete the report, then the task falls to that officer's partner. What we believe we are demonstrating is that greater contact with fellow subordinates encourages officers to devote more time to the completion of mutually disliked tasks.

The other side of our propositions revolved around subordinate preferences. Although the present data collection permits only a modest evaluation of the effect of preferences, there is supporting evidence. During the shift, if the officer mentioned that he or she was satisfied (or unsatisfied) with the squad, we recorded this as a "solidary like" (or dislike). Similarly, if the officer mentioned that he or she was satisfied (or unsatisfied) with various functional aspects of the job (preventative checks at houses and businesses, maintaining visibility for residents, and satisfaction with the beat), we recorded this as a "functional like" (or dislike).

Figures 12 and 15 demonstrate the effect of strong solidary preferences on the allocation of work. As with the other figures, the officers spend the majority of their time on runs or mobile patrol. There is a very slight increase in the amount of time spent on paperwork and a very slight decrease in the amount of time spent on meals. What this implies is that intersubordinate contact exercises a different effect from solidary preferences—seeing one's fellow officers frequently is quite different from developing strong positive relations with them. The difference between the two plots also supports the idea that intersubordinate contact is a way for officers to resolve the collective production problem instead of copying the behavior of people with whom one has a positive relationship.

Similarly, figures 12 and 16 display the effect on the simulated distribution of time across tasks when officers are most satisfied with their jobs. Like the previous three, there is a marked increase in the amount of time spent on paperwork. In this simulation, however, there is also evidence of an increase in the amount of time that officers spend on personal business: the distribution is more dispersed, and there are even two (simulated) officers who devote a plurality of their time to personal business. There are also a greater number of officers who devote a greater fraction of their time to paperwork. The checkerboard plot (fig. 16) is the most evenly gray of all the plots, indicating a roughly equal allocation of time across many tasks. In other words, heterogeneous preferences lead to heterogeneous performance.

Given the structure of each of the propositions associated with the

EPA model, we first need to assess whether or not subordinates are indifferent between tasks and then determine whether subordinates are homogeneously or heterogeneously responsive to supervision. Figure 12 shows that police officers are not indifferent between tasks and heterogeneously responsive to supervision. While we do not find a clear differentiation of task allocation, we do see a good degree of differentiation when controlling for various other factors. These conditions indicate that we should be evaluating the fourth proposition derived from the EPA model: *if subordinates are not indifferent between tasks, and are heterogeneously responsive to supervision, subordinates will separate according to type. Nonresponsive subordinates will allocate time to preferred activities, and responsive subordinates will allocate time to activities for which they receive supervision.*

The EPA model, and this proposition in particular, make predictions about conditions for pooling and separating equilibria, which follow through in the four checkerboard plots. Two of the conditions generated pooling equilibria: both the high intersubordinate contact and high solidary preference plots were reduced to a small subset of the tasks. This is consistent with officers with relatively homogeneous preferences and responsiveness. Two of the conditions generated separating equilibria: both the high supervision and high functional preferences plots yielded a more diverse set of activities, especially for the high supervision plot. This is consistent with the surmise that officers have heterogeneous preferences and especially are heterogeneously responsive to supervision. Indeed, the EPA model emphasizes that those subordinates who are responsive to supervision will be the ones who are supervised. Those unresponsive to supervision will be left alone.

We are able to more directly evaluate one of the propositions derived from the imitative model. As noted earlier, we find very strong support for the first of these propositions: *as the observability between subordinates increases, conformity of allocation of time across tasks increases (due to either information search or coercion by subordinates).* Figures 11 and 15 show clear shifts in allocation of time by police officers when they are most observable to fellow officers.

Discussion

In our *Working, Shirking, and Sabotage* (Brehm and Gates 1997), we asked the question: Who, or what, controls the policy choices of bureaucrats? We have further explored this question in this essay by examining how subordinate police officers allocate their time across a variety of tasks. By further differentiating how officers devote their time, either across three dimensions (personal activities, administrative paperwork, or policing) or even more finely across eleven dimensions, we gain some understanding of the factors that shape subordinates' decisions as to how they spend their time.

Our analysis reveals that supervisory contact results in a shift in subordinate activities away from personal business and toward administrative paperwork. Contact with other subordinates results in an even more pronounced shift toward paperwork and a considerably more concentrated dispersion of time allocations. Satisfaction with the other members of the police squad also leads to shifts in work toward paperwork. Those officers most satisfied with their jobs tend to have the greatest dispersion of time allocation, with a shift away from policing toward personal business and administrative duties. We find relatively strong support for propositions derived from the imitative model. Frequency of contact and solidary between subordinates play a large role in shaping subordinates' decisions.

These results demonstrate how strongly fellow officers affect police behavior. The importance of solidary norms with respect to an officer's squad relates to this result as well, demonstrating that the imitative model offers powerful insights for understanding police bureaucracy. Our results also suggests that supervisors play an important role in shaping how police officers do their jobs.

What are the policy implications of this research? First and foremost, this research indicates that there is no homogeneous solution to the problem of monitoring task allocations in public bureaucracies, if only because some bureaucracies will prefer that bureaucrats concentrate on a narrow range of tasks, while others would prefer that bureaucrats complete a broad repertoire of tasks. Under Wilson's (1989) typology, "production organizations" (in which both outputs and outcomes are visible) such as the Social Security Administration (in the processing of checks) or the Postal Service (in the sorting of mail) may have a very limited set of tasks for bureaucrats to complete. Under these conditions, the paramount problem is selection of subordinates with the right mix of functional preferences or to design contracts and remuneration schemes that sufficiently reward the bureaucrat for completing tasks for which he or she does not have a preference. In contrast, "craft organizations" (in which neither outputs nor outcomes are visible) such as police forces or social work agencies may prefer that their bureaucrats complete a diverse set of tasks. Under these conditions, selection is also important but perhaps with an emphasis on recruiting bureaucrats who, as a group, have heterogeneous preferences in order to encourage separating equilibrium. Contact with the supervisor, especially for amenable subordinates, can further increase the diversity of task allocations. Restraining lateral contact with other subordinates may also be necessary to encourage production of a large number of tasks.

Since our main finding concerns subordinate preferences and observability, these characteristics can also be used to make comparisons across bureaucracies. Bureaucrats working in dispersed settings with relatively little contact with other subordinates, or with supervisors, should be expected to have more heterogeneous distributions of time allocations. A prototypical example of such a dispersed, low-connection bureaucrat would be the forest ranger. As Kaufman's 1960 classic detailed, forest rangers exhibited a relatively uniform level of effort, although they were quite heterogeneous in how they divided their efforts among such tasks as conservation, constituency relations, resource management, and (increasingly) policing. Bureaucrats who work in settings where intersubordinate contacts are quite high but contacts with supervisors are quite low should be expected to parallel the results obtained here for police officers. A second prototypical example would be the social worker.

We also think that there is general utility for the underlying logic of our models outside of the bureaucratic setting. Learning under uncertainty is a condition in which most humans find themselves most of the time. To the extent that multiple and competing demands create a range of necessary tasks for social actors, the diversity of tasks adds to the conditions of uncertainty. Our pooling equilibrium result confirms the power of imitation and social proof in such conditions of uncertainty.

We could not concur more with James Q. Wilson's aphorism about tasks ("People matter, but organization matters also, and tasks matter most of all" [1989, 173]). Tasks define what it is that bureaucrats, and hence bureaucracies, do. But instead of treating tasks as immutable routines or SOPs, our approach allows us to recognize the critical interactions between people, organizations, and tasks. The preferences that bureaucrats have and the degree to which they are seen by both their supervisors and fellow bureaucrats fundamentally alter the amount of time that officers devote to the range of tasks that they face.

Notes

An earlier version of this work was prepared for presentation at the annual meeting of the Midwest Political Science Association, Chicago, IL, April 23–25, 1998. Thanks to Frank Baumgartner, Tom Hammond, and Ken Meier for comments, to Jonathan Katz and Gary King for consultation, and to Jeff Gill for the idea of using checkerboard plots.

1. We do not attempt to model the psychology behind amenability. Obviously, amenability is affected by the relationship between a subordinate and a supervisor, which will change as the relationship changes, but such an analysis has been beyond the scope of our work.

2. We note in Brehm and Gates 1997 that in a dynamic environment it may be in the supervisor's interest to punish or even fire a recalcitrant subordinate.

3. Note that imitation follows much more closely with the social psychological literature of a dual-path model of persuasion, in which some are persuaded by low cognition heuristics like imitation, than it follows such game theoretic treatments as Lupia and McCubbins 1998 purporting to explain persuasion. The Lupia and McCubbins explanation explicitly regards the level of attention as a prerequisite for persuasion in a mass democratic context, a highly saturated information environment. That said, imitation can also be modeled as a problem of informational cascades (Bikhchandani, Hirshleifer, and Welch 1992) or in replicator dynamics (Friedman 1991), in which imitation can be shown to be a Nash equilibrium solution.

4. With such a data-gathering technique, a potential Hawthorne effect is possible. Presumably such an effect would lead to a bias against spending time on leisure activities. Yet our data demonstrate no significant compunction on the part of the police officers to refrain from such activities.

5. Our data do not provide detailed information regarding the allocation of overtime work. Such information would be valuable in assessing the effect of overtime on time allocation. Lacking such information, we must presume similar and uniform exogenous budgetary constraints for all individuals.

6. An example involves multiparty voting data (Katz and King 1999). In the United Kingdom, for example, the proportion of the vote can be seen to be divided across three parties, Labour, Tory, and Alliance. In turn, the vote total sums to unity.