

The Motivation and Bias of Bureaucrats

By CANICE PRENDERGAST*

Many individuals are motivated to exert effort because they care about their jobs, rather than because there are monetary consequences to their actions. The intrinsic motivation of bureaucrats is the focus of this paper, and three primary results are shown. First, bureaucrats should be biased. Second, sometimes this bias takes the form of advocating for their clients more than would their principal, while in other cases, they are more hostile to their interests. For a range of bureaucracies, those who are biased against clients lead to more efficient outcomes. Third, self-selection need not produce the desired bias. Instead, selection to bureaucracies is likely to be bifurcated, in the sense that it becomes composed of those who are most preferred by the principal, and those who are least preferred. (JEL D64, D73, D82)

Employees often have no monetary reason for doing what they do. Instead, the outcomes of their actions intrinsically matter to them: teachers value the education of their students, police officers care about reducing crime, firefighters risk their lives in order to save others, and so on. The focus of this work is on such employees—bureaucrats—individuals who assign treatments to clients, and who have weak monetary incentives. Their willingness to exert effort without significant monetary incentives is often attributed to such intrinsic motivation,¹ and there is a large literature on the importance of hiring bureaucrats with the right preferences.² A particular

motivation that recurs in the literature is the role of bureaucrats as advocates for their clients, and a popular prescription for improving efficiency is the recruitment of workers whose goals coincide with those of their clients.³ Yet it would be hard to claim that such coincidence of preferences characterizes all bureaucracies. Indeed, the defining characteristic of many bureaucracies seems to be indifference, or in some cases hostility, toward their clients' wishes. Such accusations are routinely leveled at institutions such as the Immigration and Naturalization Service (INS), the Internal Revenue Service (IRS), and many police forces.

* Graduate School of Business, University of Chicago, 1101 East 58th Street, Chicago, IL 60637 (e-mail: canice.prendergast@gsb.uchicago.edu). A version of this paper was previously circulated under the title "Selection and Oversight in the Public Sector." I am grateful to Tano Santos, Sendhil Mullainathan, the referees, and seminar participants at Harvard University, the Massachusetts Institute of Technology, National Bureau for Economic Research, World Bank, and Université Libre de Bruxelles for helpful comments, and to the National Science Foundation and the University of Chicago for generous support. Any errors are my own.

¹ As James Q. Wilson (1989) puts it, "what is surprising is that bureaucrats work at all ... rather than shirk at every opportunity" (156). He notes that the answer is that "bureaucrats have preferences ... among them is the desire to do the job" (159).

² This is not because researchers have not thought of the contractual design issues of the economics literature, mak-

ing pay and career prospects depend on performance measures. For instance, John Brehm and Scott Gates (1997) explicitly evaluate the relevance of recent work in agency theory and ultimately conclude that although monetary incentives are not irrelevant, "the overwhelming evidence ... indicates that the bureaucrat's own preferences have the greatest effect on behavior." They note that a weakness of the economics literature is an overreliance on "leisure- or budget-maximization, [which ignores] policy-oriented preferences" (196).

³ For instance, in perhaps the best known study of street-level bureaucrats, Michael Lipsky (1980) argues, "that street level bureaucrats should be advocates for their clients is articulated explicitly in the professional training and canons of lawyers, doctors, social workers, teachers and others" (66). For other work that stresses the importance of these kinds of selection concerns, see Martha Derthick (1979), Charles Goodsell (1981), Alfred Kadushin (1985), and Warner et al. (1963).

Understanding why the motivation of firemen, nurses, and social workers seems to be so different from that of IRS and INS officials surely has little to do with how they are paid; both typically have pay that varies little with their performance. The claim here is that these differences arise largely from what kind of employees these institutions hire. Bureaucrats are typically employed to implement the objectives of a "principal." Despite this, it is shown that bureaucrats typically perform better if they are biased, by not sharing the preferences of their principal. Furthermore, desired bias varies by the kind of activities they carry out. In some instances, they advocate for their clients more than would their principal; while in other cases, they are more hostile to their interests. These differences relate directly to the kinds of "goods" they allocate, such that the IRS desires employees who are hostile to their clients' interests, while social workers likely associate more with the interests of theirs.

To see why bias is useful, note that bureaucrats who allocate treatments to clients can make two kinds of mistakes: those that benefit clients and those that harm them. To allow for the possibility of bias, it is assumed here that members of society vary in how they trade off these errors. (For instance, people may disagree over the cost of incorrectly giving welfare benefits to an undeserving candidate compared to failing to give these benefits to a deserving one.) Society aggregates these preferences into an objective. Bureaucratic bias is then identified by how the bureaucrat's preference differs from that objective by weighting client welfare more or less heavily, and is shown to determine willingness to exert effort through the rather simple mechanism of whether clients *want* bureaucrats to work harder. In other words, do clients want more accurate diagnoses?

The first factor that affects the value of biased bureaucrats is how client preferences align with those of the principal. First, consider instances where a client wants a particular treatment only if it is efficient (from the perspective of the principal), such as a patient being considered for coronary bypass surgery: he wants that treatment only if his disease truly warrants it. As the interests of the

principal and the client are aligned, the client desires a *more* informed bureaucrat. It trivially follows that those who are biased toward clients (advocates) increase efficiency.⁴ But there are many situations where the interests of clients are opposed to those of the principal. For instance, society wishes guilty suspects to be arrested, but suspects have no such preferences. Indeed, those who should be denied benefits may value receiving them most—for example, those who have cheated on their taxes most desire not to be investigated. In this instance, the optimal policy is to hire workers who are biased against clients, because, on average, clients benefit from *less* informed bureaucrats, and those who associate more with clients internalize these preferences by exerting less effort. In this way, the kind of employees desired by the bureaucracy depends on whether clients want what society wants.

A wide variety of bureaucratic benefits involve those in need valuing them more than those who are not in need. For example, the truly destitute likely value welfare benefits more than the undeserving, and the task of a benefit official is to determine need. In this instance, there are two conflicting effects. As clients value benefits more when they truly deserve them than when they do not, this effect *alone* leads to a demand for more accuracy and, hence, a desire for client advocates. However, desired bureaucratic bias is also shown to depend on how mistakes are corrected. Most bureaucracies have some oversight mechanisms in place to correct mistakes, and a common feature of such oversight is that it is biased in favor of clients. This arises for two separate reasons. First, clients typically point out bureaucratic error when it harms them, but say nothing when they benefit from an error. For example, those who have been incorrectly arrested can often make this clear, whereas those who have incorrectly been set free will hardly do so. Second, society often views the cost of incorrectly denying benefits to a deserving case to be greater than that of incorrectly

⁴ As Wilson (1989) puts it, this sense of mission "enables administrators to economize on other kinds of incentives" (95). See Mathias Dewatripont, Ian Jewitt, and Jean Tirole (1999) for formal work on the role of missions in the public sector.

giving benefits to the undeserving. For instance, convicting the innocent is perceived a worse error than letting the guilty go free. For either of these reasons, oversight tends to be biased in favor of clients. Client-favored oversight implies that these clients benefit from *less* informed agents, as they can (at least partially) correct mistakes that harm them but leave unchanged those that benefit them. This effect *alone* leads to the principal preferring hostile bureaucrats. Combining this with the fact that the benefit is more valuable when truly deserved—which leads to a demand for advocates—shows that the desirability of client advocates depends on three factors. First, how much more does the consumer value the good when the principal believes it to be needed? Second, how capable are clients at pointing out errors that harm them? Third, how easy is it to correct these errors?

These observations merely illustrate that certain kinds of employees are desired. Unfortunately, these workers may not end up being employed. If agent preferences are easily identified, hiring and firing decisions can be based on these. This seems improbable in many occupations; instead, possible employees likely have private information on their preferences. Accordingly, the issue of whether bureaucrats with the appropriate bias can be induced to self-select is considered. It is shown that with sufficient dispersion of preferences over desired outcomes, selection to bureaucratic positions becomes bifurcated, in that it consists of (a) those who are *most* preferred by the principal, and (b) those who are *least* preferred by the principal. In other words, the people who are likely to be deterred from applying for bureaucratic positions are those of intermediate preferences; those with extreme preferences become bureaucrats.

Section I begins by constructing a simple model of bureaucratic effort exertion. Section II illustrates how the relationship between client and principal preferences leads to a demand for bias,⁵ and Section III considers the sources of client-favored oversight. The modeling work

is extended in Section IV to deal with self-selection, and Section V considers a pair of relevant occupations. Section VI concludes with a brief discussion.

I. The Model

An allocation A must be made to a client, where A can take on a value of 0 or 1. This allocation affects the welfare of (a) the client, and (b) the remaining members of society. Individuals in society vary in how they weight these, and the objective of the principal is to maximize the welfare of the median member of (b).⁶ The utility of the allocation depends on a parameter α measuring the state of nature, which also takes on a value of 0 or 1. There are three actors in the model—the client, the principal, and the bureaucrat—whose incentives and actions are described in turn.

The Client.—Let $B(A, \alpha)$ be the benefit obtained by the client if his type is truly α and the allocation is A . For notational convenience, let $\beta_0 = B(0, 0) - B(1, 0)$ be the benefit attained by the client when correctly given allocation 0 relative to incorrectly being given allocation 1, and let $\beta_1 = B(1, 1) - B(0, 1)$ be the benefit attained by the client when correctly given allocation 1 relative to incorrectly being given allocation 0. Without loss of generality, let $\beta_1 \geq 0$. Thus, in a police setting, β_1 is the benefit of being set free when innocent compared to being arrested, while β_0 is the benefit of being arrested when guilty compared to being set free.

The Principal.—Assume that the direct utility to each member of the rest of society is given by $\Sigma(A, \alpha)$. However, they also value the effect of the allocation on the client: for instance, society values that its poor be fed. Members of society vary in how they weight these, where utility from the allocation for a person of type v is given by $U = \Sigma(A, \alpha) + vB(A, \alpha)$, where the parameter v measures how that person weights benefits to the client relative to the rest of society. Assume that the political process results in

⁵ This demand for bias differs from recent work by Philippe Aghion and Tirole (1997), who build a model of delegating authority based on the congruence of preferences between agent and principal.

⁶ The results are unchanged if the principal's objective also partly weights the utility of the client, as would be the case if, say, there were a random chance of each member of the population being the client.

the maximization of the expected utility of the median member of society,

$$(1) \quad S(A, \alpha) = \Sigma(A, \alpha) + v_m B(A, \alpha),$$

where v_m is the median value of v . The optimal allocation (from the perspective of the principal) involves matching the allocation to the state of nature ($A = \alpha$), so it is assumed that allocation 1 is preferred in state 1, and allocation 0 in state 0. For simplicity, begin by assuming that the surplus from the allocation is symmetric:

$$(2) \quad S(A, \alpha) = \begin{cases} 1 & \text{if } A = \alpha, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, social surplus is positive only if A is properly matched to the underlying environment, α , and the cost of error is independent of which error is made: the marginal cost is 1; the case of asymmetric error costs is considered below. The true value of α is unknown and takes values $\alpha = 1$ or $\alpha = 0$ with equal probability.

The Agent or Bureaucrat.—The principal chooses an agent to help make the assignment. The agent exerts unobserved effort $e \geq 1/2$, which affects the precision of her estimate of the correct assignment, α . She observes a signal $a \in \{0, 1\}$, which is correct with probability e , where she chooses effort e at a cost $d(e)$, where $d(1/2) = 0$, $d' > 0$, $d'' > 0$, and $d(1) \rightarrow \infty$. The agent's signal is *hard* information and is observed by all parties.⁷

There are typically two ways to induce agents to exert effort: (a) through the use of monetary incentives, and (b) through who is hired. In most bureaucracies, monetary incentives play a very limited role (see John Brehm and Scott Gates 1997; Avinash Dixit 2002; and Ronald Johnson and Gary Libecap 1989). In order to focus on selection concerns, it is assumed here that monetary incentives are nonexistent by assuming that the principal hires the agent at an (initially) exogenous wage w .

⁷ This important assumption, discussed in the conclusion, allows me to concentrate on how bias affects effort decisions without worrying about agents lying about their signals. See the conclusion for how the results are changed when information is soft.

The principal chooses which bureaucrat to hire and may not choose the median voter. Net of a fixed wage, w , her preferences depend on the allocation via (1), and how much effort she exerts. Her utility (net of the wage) is given by $U_a = \Sigma(A, \alpha) + vB(A, \alpha) - d(e)$, which can be rewritten as

$$(3) \quad U_a = S(A, \alpha) + (v - v_m)B(A, \alpha) - d(e).$$

I assume that the alternative wage of all possible agents is independent of v and normalized to 0. The population distribution of v is given by F , with density f that ranges from \underline{v} to \bar{v} . I assume initially that v is observed by all parties.

Mistakes.—Bureaucracies usually have some mechanism available for correcting mistakes. Accordingly, assume that when the agent's signal in state α is wrong, the principal receives convincing evidence of this with probability $\rho(\alpha)$. Thus, for instance, in the policing example, $\rho(1)$ ($\rho(0)$) is the likelihood of a mistake by the police officer being detected when the suspect is innocent (guilty). Assume initially that these probabilities of error correction are exogenous, an issue that is returned to in Section III.

DEFINITION: Let oversight be unbiased if $\rho(0) = \rho(1)$, and let oversight be biased if $\rho(0) < \rho(1)$.

Thus, unbiased oversight means that mistakes are equally likely to be detected, while biased oversight means that mistakes where the client truly should be given the benefit ($\alpha = 1$) are more likely to be detected than when he should be denied ($\alpha = 0$). In Section III, I explore two reasons why oversight may be biased: (a) the ability of clients to complain, and (b) asymmetries in the principal's preferences.

Timing.—The timing of the game is as follows. First, the principal offers w to an agent. If accepted, he is employed; otherwise, the good is allocated without an agent. Second, nature as-

signs α to the client and the agent exerts effort and is paid. Third, all parties observe the agent's signal, a . Fourth, the principal receives his information on whether an error has been made and makes an assignment. I now address the Bayesian Nash equilibria of the game.

Note, as a result of these timing and commitment assumptions, that the principal is limited in his credible strategies: in effect, the only ex ante choice variable to affect outcomes is through who is hired.

II. The Demand for Bureaucratic Bias

To begin, consider who the principal would like to hire. To address this, it is necessary to know how the principal allocates the good at the final stage of the game. As a result of the symmetry assumption, the principal optimally allocates the good according to the agent's signal unless he receives information that the allocation is incorrect, in which case he changes the "decision." As a result, the probability that the good is correctly allocated in state α is $e + (1 - e)\rho(\alpha)$.

The bureaucrat exerts effort to affect outcomes. To understand how bias affects effort, note that the client is affected by the accuracy of the allocation. Accuracy is affected by two features: (a) how hard the agent works, e , and (b) how often her mistakes are corrected, ρ . The expected private nonwage benefit for an agent of type v is given by

$$\begin{aligned}
 (4) \quad & e + (1 - e) \frac{\rho(0) + \rho(1)}{2} + (v - v_m) \\
 & \times \left\{ e \frac{B(0, 0) + B(1, 1)}{2} + (1 - e) \right. \\
 & \times \left[\rho(1) \frac{B(1, 1)}{2} \right. \\
 & \left. \left. + (1 - \rho(1)) \frac{B(0, 1)}{2} \right] \right. \\
 & \left. + (1 - e) \left[\rho(0) \frac{B(0, 0)}{2} \right. \right. \\
 & \left. \left. + (1 - \rho(0)) \frac{B(1, 0)}{2} \right] \right\} - d(e).
 \end{aligned}$$

Thus, in state α , there is a probability $e + (1 - e)\rho(\alpha)$ that the agent receives the principal's preferred outcome, and probability $(1 - e)(1 - \rho(\alpha))$ that the incorrect allocation is given. The agent's choice of effort e^* is then given by

$$\begin{aligned}
 (5) \quad & (1 - \rho(0)) + (1 - \rho(1)) + (v - v_m) \\
 & \times [\beta_0(1 - \rho(0)) + \beta_1(1 - \rho(1))] \\
 & \geq 2d'(e^*),
 \end{aligned}$$

where she trades off the effects on social surplus and the client relative to effort costs. For $v - v_m$ sufficiently small, this constraint binds; for larger v it can be the case that the constraint does not bind for some agents, where they all exert the lowest possible level of effort, $e = 1/2$.⁸ Unless otherwise stated, I will treat this constraint as binding.

To see how this affects the demand for biased bureaucrats, note that increasing v increases effort only if

$$(6) \quad \beta_1[1 - \rho(1)] + \beta_0[1 - \rho(0)] \geq 0.$$

This condition generates all the results of the first half of the paper. If (6) holds, more altruistic agents (advocates) induce greater effort, while if it is violated, more hostile agents work harder.

The demand for bias depends on two factors. First, how aligned are the interests of the principal and client—the β 's? Second, how often is the bureaucrat pivotal; in other words, what is the likelihood that her mistakes are not corrected, $1 - \rho(\cdot)$? By backward induction, all that remains to be determined is the choice of agent, i.e., the v that maximizes the principal's welfare. As wage and worker participation are (thus far) exogenous, the unique equilibrium is

⁸ Consider the case where $[\beta_0(1 - \rho(0)) + \beta_1(1 - \rho(1))]$ is large and negative. Then agents with large enough positive v prefer effort of $1/2$ to any higher level. This is because they prefer less accurate assessments, and accuracy is bounded below at $1/2$. Consequently, a range of types exerts the lowest possible level of effort. In instances where a range of agents exerts the minimum effort, the propositions below, which denote a preference for a certain type of worker, should be replaced with a weak preference.

where the principal chooses the agent who exerts most effort, as this increases expected surplus for the median member of society. Proposition 1 immediately follows.

PROPOSITION 1:

- If $\beta_i \geq 0$, effort increases in client advocacy (*v*).
- If $-\beta_0 > \beta_1$, then advocacy decreases effort.
- If $\beta_1 = -\beta_0$, then, if oversight is biased, client advocates exert less effort. If oversight is unbiased, all agents exert the lowest level of effort, and the principal is indifferent over whom to hire.
- If $\beta_1 > -\beta_0 > 0$, then, if oversight is unbiased, effort is increasing in altruism. If oversight is biased, effort is increasing in altruism only if

$$(7) \quad \frac{\beta_1}{\beta_0} \geq \frac{1 - \rho(0)}{1 - \rho(1)}.$$

Consider each of these cases in turn.

When Client and Principal Preferences Align: $\beta_0 \geq 0$.—In some settings, the interests of the client and principal coincide. A good example of this would include a patient being evaluated for heart surgery: only if the client truly needs the surgery is he likely to want it. This implies that $\beta_1 > 0$ and $\beta_0 > 0$. Then (6) always holds for any oversight levels, and the uniquely optimal choice of agent is the most altruistic to the client. The idea is trivial; when clients have the same ranking of outcomes as the principal, they value more accurate estimates of their needs, and the agent who cares most about the client exerts the most effort.

When Client Preferences Are Opposed to Those of the Principal: $\beta_0 < 0$.—A wide range of bureaucratic environments involve clients who have preferences that relate little to those of the principal. To see how this affects desired bias, first consider the case where clients desire the good at least as much when they do not deserve it as when they do, where $-\beta_0 \geq \beta_1$. (For example, those who have cheated on their taxes wish to avoid an audit at least as much as

the honest.) Then, if the inequality is strict, $-\beta_0 > \beta_1$, the client benefits from a less informed agent, independent of whether oversight is biased or not. This can be seen from (6) and simply reflects the fact that the client cares more about avoiding investigation when guilty than avoiding investigation when innocent.

Next, consider the case where $\beta_1 = -\beta_0$. For instance, this could be a traffic fine or a prison sentence, which has a similar cost to the client whether he is guilty or innocent. Begin with the case where oversight is unbiased so (6) becomes $\beta_0 + \beta_1 \geq 0$. When $\beta_1 = -\beta_0$, all the client cares about is the average probability of receiving the good, as he weights each error equally. If $\rho(0) = \rho(1)$, the probability of receiving the benefit is $1/2$, independent of effort, and so the bureaucrat chooses $e = 1/2$. Consequently, agent attributes have no effect on effort and the choice of bureaucrat is irrelevant. But when the client, on average, benefits from oversight ($\rho(0) < \rho(1)$), (6) is always violated if $\beta_1 = -\beta_0$, and incentives are maximized by hiring agents who are most hostile to the client's wishes. Here the client strictly prefers inaccurate decisions, because he can (at least partially) correct the mistakes that harm him. Hence, the optimal response by the principal is to hire the agent most hostile to the client.

Finally, consider the case where the client wants the benefit ($\beta_0 < 0$), but values it more when he is deserving: $\beta_1 > -\beta_0 > 0$. This case likely covers many benefits offered by bureaucrats. For example, those who are truly ill want medical tests more than those who are using them as a precautionary measure, or those who are truly destitute need welfare assistance more than those whose needs are less great. In the case of symmetric oversight, the reason for altruistic bureaucrats is simply the mirror image of the case where $\beta_0 + \beta_1 < 0$: here the agent cares more about receiving the benefit when warranted than when not warranted, $\beta_0 + \beta_1 > 0$, and so, on average, values more accuracy. With biased oversight, the decision on whether to hire advocates depends not simply on the difference in the benefits β —which would lead to advocates—but also on the ρ 's, where biased oversight leads to a demand for hostile agents. The logic for this is simple: as clients can correct

mistakes that harm them and benefit from those that work in their favor, they prefer less accuracy. As a result, the choice of agent depends on which effect is stronger, an issue that is addressed in Section III, when oversight is endogenized.

A. Wages

Thus far, all that has been shown is that certain types of agents work harder than others. As it has been assumed that wages are set exogenously and supply of the desired agent is forthcoming, the principal hires the agent who exerts the most effort. A plausible alternative assumption, however, would be that the principal can choose the wage, and will do so to match the agent's reservation utility.

In this case, the principal becomes the residual claimant on the utility of the bureaucrat, as wages simply change to offer the bureaucrat her reservation utility. This extension strengthens the value of hiring the hardest working agent because, holding all else constant, the agent who exerts the most effort gains the most utility from the job. As a result, he can be offered a lower wage than someone who exerts less effort.⁹ This does not change the result above, as illustrated in Lemma 1, whose proof is in the Appendix.

LEMMA 1: *Assume that the principal can choose agent wages and (5) binds. Then, if effort is increasing (decreasing) in v , the required wage to employ a worker is decreasing (increasing) in v . As a result, the principal hires the agent who exerts the most effort.*

In other words, because agents derive utility from affecting clients, and those who care most about this are willing to exert more effort, they accept lower wages for the opportunity to do so. Hence, the results generalize to a more traditional assumption on wage determination.

⁹ Note that this is only true *because* those with more desired preferences exert more effort. The agent values the utility of the client, independent of whether she works: if effort were no higher than that of an alternative, she would get as much utility via vB without being employed.

III. Oversight

Oversight matters only in the (important) case where clients benefit from receiving the allocation, but at least as much when deserved: $\beta_1 \geq -\beta_0 > 0$. To address this further, I now endogenize the oversight decision with the specific purpose of offering two reasons why oversight is likely to be biased. From (7), it should be clear that two issues affect how much the principal intervenes to correct mistakes: (a) how much she knows about mistakes being made, and (b) how much she cares about these mistakes. I address each in turn.

A. Client Interests

Principals typically rely on relevant information to help them intervene. A primary source of such information is complaints by clients. Information that mistakes have been made typically comes from interested parties. Specifically, those who have been wronged will often voice their concerns.¹⁰ To model this, consider a simple extension of the model above where the principal may observe some useful information before deciding whether to investigate.

Assume that in the absence of the principal receiving disconfirming evidence from the client, the principal receives convincing evidence of an error with probability ρ^* . In addition, however, now assume that when an error has been made, the client observes hard evidence of that error with probability s . In all other instances, the client observes no more than the agent. For simplicity, assume that the principal's and agent's signals are conditionally independent. The signal received by the client is

¹⁰ By definition, whenever a poor decision has been made by a bureaucrat, there is someone who has been wronged who could conceivably reveal this information. However, a central issue that arises is the diffusion or concentration of losers. For instance, take the case of a social worker who makes an incorrect decision in assigning benefits to an applicant. If the applicant is incorrectly given these benefits, society at large loses. But the losers typically do not know, as it is in the interest of no individual taxpayer to monitor these cases. By contrast, if an applicant is incorrectly denied benefits, this is more likely to come to light as the client has a concentrated interest in revealing this mistake to the authorities. As a result, there is a likely asymmetry in mistakes coming to light.

private information, but the client can send this information to the principal if he chooses.¹¹

The client's information can now possibly be used to correct bureaucratic mistakes. His choice is either to report his disconfirming evidence (if he has any) or simply to report the agent's signal. It should be clear that the client will point out errors that harm her. As a result, in the state where $\alpha = 1$, if the agent's signal is $a = 0$, the agent reveals any information she has that an error has been made. As a result, it is simple to see that

$$(8) \quad \rho(1) = \rho^* + (1 - \rho^*)s.$$

By contrast, when the agent benefits from an error—when $\alpha = 0$ but $a = 1$ —the client cannot be induced to reveal information that an error has been made. Instead, she reports the same information that is contained in the agent's signal.¹² Hence, $\rho(0) = \rho^*$.

As a result, the ability of clients to complain can lead to biased oversight. If this bias is sufficiently strong, the principal will respond by hiring agents who are hostile to their interests. Rather than use this section simply to illustrate this point, an additional purpose is to offer additional insights into understanding the factors that determine bias in the case where $\beta_1 > -\beta_0 > 0$, in which case the condition for client advocates being desired—(7)—becomes

$$(9) \quad \frac{\beta_1}{\beta_0} \geq \frac{1 - \rho^*}{1 - \rho^* - (1 - \rho^*)s}.$$

What Characteristics of the Bureaucracy Lead to More Altruistic Agents?—Equation (9) illustrates how the optimal makeup of the bureaucracy depends on three factors: (a) the effect of being “deserving” on client welfare from

¹¹ Formally, this changes the timing of the game from that above to the following: “Fourth, the principal and agent observe their information on whether an error has occurred and the agent transmits a report to the principal. Fifth, the principal makes an allocation.”

¹² What is important here is that the agent does not always get *additional* confirming evidence that the decision was correct. If that were the case, the usual unraveling can occur when $\alpha = 1$, where the principal can infer from not receiving the additional confirming evidence that an error must have been made in this case.

the benefit, (b) the likelihood of oversight, and (c) how the client affects oversight.

1. The left-hand side of (9) measures the value of the benefit when “deserved” *relative* to “undeserved.” The greater it is, the more likely it is that altruistic agents will be hired, for the reasons described above.
2. The choice of agent depends on the ability of the principal to observe (and correct) the bureaucrat's mistakes, where outcomes that are more difficult to verify (ρ^* decreasing) result in an increase in the value of more altruistic agents. More informally, the reason why clients want less informed agents is loosely based on the fact that they can correct mistakes that harm them, but ultimately benefit from mistakes that are in their favor. But this works only if their complaints result in these bad decisions being likely overturned. When oversight is more costly, this is less likely.
3. Finally, altruistic agents are less likely when clients are better informed. Again, consider the outcome when $s \rightarrow 0$, in which case clients are uninformed. (For instance, a cancer patient may know little about whether he is suitable for chemotherapy.) In this case, the right-hand side of (9) tends to 1, and again (9) is satisfied. The reason is similar to that above: well informed clients can effectively focus investigations on cases where they have been incorrectly denied, better than their less informed counterparts.

B. Asymmetries in the Principal's Preferences

So far, it has been assumed that the principal's preferences are symmetric. A second source of biased oversight can arise when the principal cares about some errors more than others, and so the surplus function is not symmetric, as in (2). For example, it is a tenet of the justice system that it is better for the guilty to go free than for the innocent to be convicted. Similarly, many would argue that the benefit that a truly sick person receives from a procedure exceeds the cost of giving the procedure to one who is not as deserving. Such asymmetries in the utilitarian preferences of the principal have not been addressed thus far.

To see how this affects the results, ignore client feedback, so that the principal receives information on all errors with probability ρ^* . However, return to the more general specification of the principal's preferences in (1) and define $\sigma(1) = S(1, 1) - S(1, 0)$ and $\sigma(0) = S(0, 0) - S(0, 1)$ as the costs of error in the two states. In the previous section, these costs were both 1. In keeping with the discussion above, most of these instances concern cases where the costs of denying deserved benefits to clients are greater than the costs of incorrectly assigning those benefits. Consequently, assume that the cost of mistakenly denying the benefit ($A = 0$ when $\alpha = 1$) is higher than that of mistakenly giving the benefit: $\sigma(1) > \sigma(0)$.

Dropping the symmetry assumption has two effects on the model. First, the agent's signal may no longer affect the allocation, in which case there is no need to hire an agent. For instance, if the cost of mistakenly denying benefits to a client is large enough (compared to the cost of mistakenly giving the benefit) the principal will give the benefit even if the agent's signal says otherwise. In this case, the allocation problem is no longer interesting, as there is no role for the agent. For this not to be the case, it is necessary that $e\sigma(0) \geq (1 - e)\sigma(1)$. Therefore, assume that this is the case.

Second, the return to correcting bureaucratic mistakes that harm the client exceeds that from correcting the ones that benefit him. How does this affect bureaucratic bias? As the model is currently construed, it makes no difference because changing the decision is costless. However, a natural extension of the model would be to allow costs of overturning decisions made by the bureaucrat, as realistically an investigation is often required by the principal. Assume, accordingly, that the principal faces a random cost c of "changing" the allocation from that recommended by the agent's signal. Let the distribution of costs be given by $G(c)$.

In this case, the principal will overturn decisions when (a) he receives information that an error has been made, and (b) the cost of overturning the decision is less than its benefit. As a result, it is simple to show that

$$(5.5q) \quad \rho(1) = \rho^*G(\sigma(0)),$$

and

$$(5.5q) \quad \rho(0) = \rho^*G(\sigma(1)),$$

so that $\rho(1) \leq \rho(0)$. The preferred choice of bureaucrat is still given by (6), but oversight is now weighted in favor of investigating instances where the benefit is denied to the client for the simple reason that the principal cares more about these mistakes. As a result, altruistic bureaucrats are now preferred if

$$(12) \quad \beta_1[1 - \rho^*G(\sigma(1))] + \beta_0[1 - \rho^*G(\sigma(0))] \geq 0.$$

For the relevant case where $\beta_0 < 0$, this condition is more difficult to satisfy than (6), and so those who are biased against clients are more likely to be preferred.¹³

IV. The Self-Selection of Bureaucrats

The previous section identifies a demand for biased bureaucrats. But the bureaucracy will exhibit this bias only if appropriate supply is forthcoming. It seems oftentimes unrealistic to assume that preferences will be apparent in applicants; instead, the principal may be reliant on other mechanisms. This section deals with the case where individuals hold private information on their preferences and apply for the position based on those preferences. For instance, it is likely that child-care workers have more information on how much they like children than do potential employers.

First, note that there is a simple reason why the most desired agents are likely to apply for the position: by exerting more effort than would an alternative, they affect clients in the way desired by themselves and by the principal. Despite this, if preferences are sufficiently disperse, the bureaucracy becomes comprised of both the most desired agents and those with preferences *least* desired. The reason is intuitive: while those who have the most desirable preferences affect clients in the way desired by the principal, those agents who have the least desirable preferences affect outcomes in the way that *they* desire, even though it is discordant with the principal's desire.¹⁴

¹³ Note that effort exertion also changes from (5) when surplus is asymmetric, but this does not change the selection condition from (6).

¹⁴ For example, many could disagree with a policy that restricts health or welfare benefits to illegal immigrants. On

In order to illustrate these influences, assume that v is private information for the agent.¹⁵ Furthermore, in order to reduce complexity, a pair of simplifying changes are made to the basic model.

A. Simplified Model

In this section, I make the following assumptions:

- The equilibrium level of oversight by the principal is ρ^* for all α .
- The agent chooses between two levels of effort, \bar{e} and \underline{e} , where $\bar{e} > \underline{e} \geq 1/2$. The cost of the lower effort is 0 and the cost of the higher effort is d .

To avoid detailing each of the cases above, attention is restricted to the situation where $\beta_i > 0$. Following Proposition 1, the principal prefers the most altruistic agent.¹⁶ Unlike in the previous section, the principal cannot now observe the altruism of the applicant. Instead, he can observe only those who apply for the position at wage w , and as he has no more information, randomly chooses among the applicants. As a result, the only instrument available to the principal is the wage w .

As oversight is ρ^* , the probability that the allocation is correct is where the agent is either correct (with probability e), or is incorrect but an investigation overturns the decision (with probability $(1 - e)\rho^*$). Hence the principal chooses w to maximize

$$(13) \quad E[e + (1 - e)\rho^* - w],$$

subject to the incentives of possible agents. The effort level exerted by the agent depends on (a) what effort to exert, conditional on being

the basis of the previous section, this would naturally appear to be an occupation where the principal would hire agents who are hostile to client interests. However, it may also be the case that those who care *most* about these groups will select into these occupations in order to thwart the desires of the principal.

¹⁵ Note here that the intrinsic motivation of any given worker is invariant to wages, unlike Bruno Frey and Reto Jegen (2001) or Roland Bénabou and Tirole (2003).

¹⁶ This restriction is purely for notational simplicity. The results for the case where preferences are not aligned is a mirror image of the results below.

hired—Incentive Compatibility; and (b) who applies for the position—Individual Rationality. Each is considered in turn.

The Incentive to Exert Effort.—If employed (the only relevant condition for the incentive compatibility constraint), the utility of the client is $(e + (1 - e)\rho^*)[B(1, 1) + B(0, 0)]/2 + (1 - e)(1 - \rho^*)[B(1, 0) + B(0, 1)]/2$. As a result, simple computations show that the agent exerts high effort only if

$$(14) \quad [\bar{e} - \underline{e}](1 - \rho^*) + \frac{v - v_m}{2} [\bar{e} - \underline{e}](1 - \rho^*)(\beta_0 + \beta_1) \geq d.$$

This is the discrete analog of (5) and, if it binds for some agents, defines the critical value of v , called \hat{v} , above which agents exert effort, defined by

$$(15) \quad \hat{v} = v_m + \frac{d - [\bar{e} - \underline{e}](1 - \rho^*)}{[\bar{e} - \underline{e}](1 - \rho^*)} \frac{\beta_0 + \beta_1}{2}.$$

This condition is intuitive; those who feel greater altruism toward the client have greater utility from exerting effort, where this incentive is increasing in the value that the client received from increased accuracy $(\beta_0 + \beta_1)/2$, and from the likelihood that their effort is pivotal in the allocation (i.e., $1 - \rho^*$).¹⁷

Who Becomes a Bureaucrat?—The novelty of this section arises from understanding who becomes a bureaucrat. An agent applies for the position if her expected utility from doing so exceeds her reservation utility. Consider the agent's reservation utility: this consists of her (zero) wage and the welfare she gets from the utility of the client. Because of this, her reservation utility depends on how hard alternative hires work. Let μ be the fraction of

¹⁷ Note that the result that those who are more altruistic exert more effort arises only because β_i is positive; if the client benefited more in states where it was undeserved (the IRS example), then the more altruistic agents would be less likely to exert effort.

agents who exert high effort, and define $\varepsilon = \mu\bar{e} + (1 - \mu)\underline{e}$ as the expected level of effort. Then the type v agent's reservation utility is

$$(16) \quad (\varepsilon + (1 - \varepsilon)\rho^*) \times \left[1 + \frac{v - v_m}{2} (B(0, 0) + B(1, 1)) \right] + (1 - \varepsilon)(1 - \rho^*) \frac{v - v_m}{2} \times (B(0, 1) + B(1, 0)).$$

If, instead, she is hired, her utility from effort e is

$$(17) \quad w + (e + (1 - e)\rho^*) \times \left[1 + \frac{v - v_m}{2} (B(0, 0) + B(1, 1)) \right] + (1 - e)(1 - \rho^*) \frac{v - v_m}{2} \times (B(0, 1) + B(1, 0)).$$

Now, consider the net utility from being employed as a bureaucrat. If the agent exerts effort \bar{e} , her rents are

$$(18) \quad w + (\bar{e} - \varepsilon)(1 - \rho^*) \times \left(1 + \frac{v - v_m}{2} (\beta_0 + \beta_1) \right) - d.$$

Thus, the agent gains utility from working harder than an alternative hire, but at the cost of that effort. This defines the critical value of altruism, $\tilde{v}(\bar{e})$, above which the agent applies for the position, given by

$$(19) \quad \tilde{v}(\bar{e}) = v_m + \frac{2(d - w - (\bar{e} - \varepsilon)(1 - \rho^*))}{(\bar{e} - \varepsilon)(1 - \rho^*)(\beta_0 + \beta_1)}.$$

Again, this is intuitive; those who are more altruistic value the increased utility that they provide to the client, as they exert at least as

much effort as the alternative hire. In this sense, self-selection operates in a way favorable to the principal, because their higher effort affects clients in the preferred way.

Of more interest is the low-effort agent. If the agent exerts effort \underline{e} , her rents from the job are

$$(20) \quad w + (\underline{e} - \varepsilon)(1 - \rho^*) \times \left(1 + \frac{v - v_m}{2} (\beta_0 + \beta_1) \right).$$

This then defines a critical value of $\tilde{v}(\underline{e})$, given by

$$(21) \quad \tilde{v}(\underline{e}) = v_m + \frac{2((\underline{e} - \varepsilon)(1 - \rho^*) - w)}{(\underline{e} - \varepsilon)(1 - \rho^*)(\beta_0 + \beta_1)}.$$

Here, $\tilde{v}(\underline{e})$ identifies the *most* altruistic (low-effort) person who applies for the position; all those who are more altruistic do not do so (for $\mu > 0$).

The reason should be clear. By exerting effort of \underline{e} , the agent reduces the utility of the client. But those who are less altruistic care least about this cost (indeed, it could be a benefit), and so the least altruistic find the low-effort outcome more appealing than do more altruistic individuals. As a result, if $\tilde{v}(\underline{e}) > \underline{v}$, bureaucrats are drawn from those most hostile to the client's and principal's interests. (Another interpretation of this could be that they value the wage w and care little about the cost that their lack of effort imposes on the client.) Note also that if the distribution of preferences is sufficiently dispersed, i.e., $\bar{v}(\underline{v})$ is high (low), then this bifurcation arises if some choose not to participate.

Two possible outcomes can arise here, depending on the wage offered. First, for high enough wages, all apply for the position. Then there is random selection of bureaucrats from the population, and the equilibrium is characterized by (15): all those with $v \geq \hat{v}$ exert high effort, and those with lower levels of altruism exert low effort. Alternatively, some may choose not to apply, in which case the equilibrium is characterized by (19) and (21). To determine which outcome arises, what matters is whether the participation constraint is satisfied for type \hat{v} . (Not surprisingly, \hat{v} is identical to the

point where $\tilde{\alpha}(\bar{e}) = \tilde{\alpha}(\underline{e})$.) Straightforward computations show that this arises for all wages above \hat{w} , where $\hat{w} = [(\varepsilon - \underline{e})d - 2(\underline{e} - \varepsilon)(\bar{e} - \varepsilon)(1 - \rho^*)]/(\bar{e} - \underline{e})$. Proposition 2 follows.

PROPOSITION 2: *If the wage exceeds $\hat{w} = [(\varepsilon - \underline{e})d - 2(\underline{e} - \varepsilon)(\bar{e} - \varepsilon)(1 - \rho^*)]/(\bar{e} - \underline{e})$, all agents apply for the position and only those above \hat{v} exert high effort. Otherwise, those who select to become bureaucrats are those with $v \leq \tilde{\alpha}(\underline{e})$ who exert effort of \underline{e} , and $v \geq \tilde{\alpha}(\bar{e})$, who exert effort of \bar{e} , where $\tilde{\alpha}(\underline{e}) < \tilde{\alpha}(\bar{e})$.*

In other words, for low enough wages, bureaucrats are drawn disproportionately from both ends of the distribution, so long as preferences are intense enough: $\tilde{\alpha}(\underline{e}) > \bar{v}$ and $\tilde{\alpha}(\bar{e}) > \underline{v}$.

Will Selection be Bifurcated?—The principal can affect which equilibrium arises through the choice of wage. There is little that is unintuitive here. The principal's objective function is linear in the expected effort of the agent and is given by $\varepsilon + [1 - \rho^*](1 - \varepsilon) - w$. If wages are set above \hat{w} , all apply for the position. But offering a wage this high has two problems. First, it is costly. Second, it may be that a lower wage offers higher average effort, as higher wages induce more high-effort and low-effort applicants. As a result, the principal may choose a wage such that the participation constraint binds for some workers, and so there is nonrandom selection.

Given this, which bureaucracies are most likely to exhibit this tendency? First, it depends on bureaucrats caring sufficiently about the activity. As a result, it would not be surprising to see that those who assign visas or who provide social welfare benefits could come from both ends of the altruism distribution, but it would be a less plausible explanation for selection at a department of motor vehicles, where it is hard to see the strength of preferences being great enough. In terms of the model, β is too small. Another important issue is the extent to which the agent is pivotal— $(1 - \rho^*)(e - \varepsilon)$ in the model. Again, from (19) and (21), note that as $1 - \rho^*$ tends to zero, the only factor that affects supply is wages rather than v . An important component of the influence of bureaucrats is the extent to which they are allowed to exercise discretion. In many bureaucracies, employees do little more than rubber stamp documents

provided by the client; again, the DMV would be a good example. Only in instances where bureaucracies must delegate decision making to agents do these issues become important: relevant examples here include police officers, immigration officials, and social welfare workers at a department of children and family services, who must use judgment in deciding whether to take a child from a family.

V. Examples

As mentioned in the introduction, the literature on achieving efficiency in public bureaucracies largely considers monetary incentives as infeasible or counterproductive, and instead emphasizes the empathy (or lack of empathy) that bureaucrats feel toward their clients. This section offers examples from two occupations.

A. Social Workers

Researchers interested in understanding why social workers perform as they do routinely emphasize their empathy and altruism toward clients. For instance, in a survey on the preferences of social workers, Robert Peabody (1964) notes that “by far the most dominant organizational goal perceived as important ... is service to clientele” (66), where 83 percent of survey respondents view such service as important, compared to only 9 percent who see “obligation to taxpayers” or “assistance to the public in general” as important concerns affecting their decisions. Similar evidence is found in Richard Weatherly et al. (1980), who note that “a feeling of wanting to help others” was the main reason cited by workers in their job satisfaction. Brehm and Gates (1997) survey social workers and find that “social workers ranked their customers as the most important influence over how they spend their time,” with “helping others” as the single most common reason for liking their jobs. Goodsell (1981) describes the altruism of social workers toward their clients in documenting that the greatest complaint that social workers have is “a scarcity of resources with which to help clients.” Along these lines, Kadushin (1985, 233) notes the absence of such altruism as a sign of trouble in a social worker. This is supported by a series of studies predicting the turnover behavior of social workers, such as

Lee and Ashforth (1993), Laura Blankertz and Susan Robinson (1997), Michal E. Mor Barak, Jan Nissly, and Amy Levin (2001), and Alberta Ellett (2001), who illustrate a negative relationship between the turnover of social workers and their reported commitment to clients.

The primary job of social workers is to provide advice to individuals to improve their lives, so this would naturally seem a case where $\beta_i \geq 0$. As a result, it should not be surprising, in the light of the model above, that there is a demand for client advocates. Yet there is one notable exception to this emphasis on social workers as client advocates. Specifically, social workers are sometimes charged with the task of determining whether clients qualify for certain treatments. In these instances, the use of altruistic agents can cause problems. This is noted by Derthick (1979) in her study of the US Social Service Administration (SSA). After its founding, the SSA was given the mission of administering retirement benefits to people who had paid into the Social Security trust fund. These decisions were based on clearly defined national standards, and its recruitment policies emphasized a "client-serving ethic." During the 1970s, however, the SSA was given additional responsibilities, namely for allocating Supplemental Security Income and Disability Insurance. By contrast to retirement benefits, the SSA now had a considerable role in determining who should receive these benefits. As Derthick describes it, the organization found it almost impossible to cope with these additional responsibilities, largely as the desired characteristics in workers for these new tasks conflicted with those needed beforehand. In effect, their empathy with clients, which led them to go to considerable lengths to ensure they received their deserved benefits, also served to make them less capable of (efficiently) denying them undeserved benefits.

B. Police Officers

Many goods allocated by the public sector are benefits enjoyed at least as much by the underserving. Police officers may be one such example. There is a considerable amount of work on the attitudes of police officers, though this work is more ethnographic than the statistical evidence on social workers. It would be hard to claim from the results of this literature, how-

ever, that police officers exhibit attributes similar to those of social workers. Instead, they seem to treat their primary clients (suspects) with little empathy, and are almost certainly less sympathetic than the average person. As Goodsell (1981) describes it, "police officers have a cynical view of human nature," while "social workers display just the opposite attitudes" (51). Also notable is how different police officers perceive themselves in relation to the general citizenry. According to Jerome Skolnick (1968), "If any one finding can be said to characterize all of the recent studies of police, it is the fact of social isolation ... from the citizenry" (17). Specifically, compared to the general public, police officers are "extremely ambivalent about the rights of others" (Peter Manning and John Van Maanen 1978, 19) and view lack of support from the public in convicting suspects as their greatest occupational concern. Furthermore, "the policeman feels that criminal procedure has been unfairly weighted against him" (Skolnick 1966, 199), where "the state requires the policeman to work in a milieu filled with extraneous and to him needless restrictions" (202).

VI. Conclusion

To many people, the observation that a police force attracts a different sort of person than a community welfare department may be so obvious as to hardly warrant stating. Yet these differences ultimately relate to the nature of policing compared to being a social worker, and the objective is to identify *why* police officers may share characteristics with INS or IRS officials, but have a radically different outlook from firemen, social workers, or employees in a hospital. People often joke about how effectively institutions like the IRS and INS select employees characterized by unpleasantness and a lack of willingness to help. This is often attributed to the absence of competition, and so they have leeway to be unpleasant to clients without fear of "losing their business." Yet many doctors, nurses, firemen, social workers, and teachers also face little competition, yet do not seem to be as vilified. Indeed, a more common moniker for many of these individuals is "do-gooders." If nothing more, the purpose of this paper has been to argue that these differences warrant

study, and to offer a theory of how such attributes affect the efficiency of bureaucracies.¹⁸

As with any model, the plausibility of the assumptions used determines its credibility. With this in mind, consider the interpretation and plausibility of some of these.

Bias against Clients.—So far, the possibility that oversight is biased against clients— $\rho(0) > \rho(1)$ —has been ignored. This is largely because its empirical relevance seems limited.¹⁹ In the interests of completeness, consider the outcome when oversight is biased against clients. These largely are a mirror image of those above. Specifically, when incentives are not so misaligned, where $\beta_1 > -\beta_0 > 0$, altruistic agents are now always preferred—unlike in Proposition 1 where it depended on parameter values. This is because, here, the bias in oversight against clients exacerbates the desire for advocates. By contrast, when preferences are highly misaligned, where $-\beta_0 > \beta_1$, the desired bureaucrat now depends on parameter values, unlike in Proposition 1. This is because the client's desire for a less informed bureaucrat is traded against the fact that oversight harms him on average. As such, the results are simply the mirror image of those above.

Rare Events.—In the modeling section, the two states have been assumed to be equally likely. This implies that without oversight by the principal, each allocation is equally likely, *independent of effort*. But this is not true if one state is more likely than the other. Suppose, for instance, that the probability of $\alpha = 1$ is $\phi > 1/2$. Then if the agent exerts effort of less than ϕ , her signal is less valuable than the prior, and so $A = 1$ is implemented unless the principal investigates. As a result, for effort levels below ϕ , the distribution of the optimal allocation based on the agent's signal is $A = 1$. For effort levels higher than ϕ , the agent's signal dominates, and

the ex ante distribution of the recommended state is $A = 1$ with probability ϕ . Thus, the ex ante expected outcome depends on the effort level. This implies that in equilibrium no agent exerts effort between $e = 1/2$ and $e = \phi$, because it does not change the allocation. This changes the results above by adding another factor to the principal's choice, namely, by giving a reason to hire an agent who values the change in the client's welfare in the rare state. Thus, if the likelihood of arresting a suspect is low, the principal prefers police officers who care especially about errors when the guilty are not arrested, even though rare.²⁰

Truth-telling.—Throughout the paper, it has been assumed that the agent's information is *hard*, in the sense that her evidence can be observed by the principal. This assumption is important, as it means that the agent cannot lie about her signal by, for instance, suggesting that a client should be given the benefit when her signal says otherwise. In a previous version of the paper, allowing this information to be *soft*—such that it could not be verified—was considered. If the agent's information is soft, the principal has to rely on the agent revealing it truthfully rather than lying. This leads to a possible need to limit the bias of the agent compared to the results above in order to guarantee truth-telling, so that those who are hired are less biased than outlined here.

Formally, it adds a truth-telling constraint, where the return to lying in one of the two states is increasing as the agent moves farther from the median voter. In state α , this condition is given by

$$(22) \quad S(\alpha, \alpha) + (v - v_m)B(\alpha, \alpha) \\ \geq S(\alpha', \alpha) + (v - v_m)B(\alpha', \alpha).$$

As the agent cares about surplus in addition to

¹⁸ In this sense, the paper shares some similarity with Dewatripont and Tirole (1999), who show that offering agents objectives that differ from that of the principal can increase effort exerted.

¹⁹ One case where it may have some relevance is where the bureaucracy incurs costs from treating clients and wishes to avoid these costs more than it values the client's welfare: for instance, a costly medical treatment that the hospital would like to avoid even if socially beneficial, and so is reluctant to overturn decisions that save them money.

²⁰ The reason is simple. By exerting low effort, the client is more likely to experience the ex ante, more common, state than with higher effort, because the agent's information is not enough to outweigh the prior. If the agent values the client receiving the benefit of the more common state, this becomes another reason to exert no effort. To counteract this, the principal desires an agent who especially cares about errors in the rare state, i.e., letting the guilty go free.

the client's welfare, the benefit of lying is small enough for agents close to the median voter to guarantee truth-telling—as $v \rightarrow v_m$, this condition is always satisfied for any effort level—but rises (for one of the two α 's) until at some point truth-telling is violated. As a result, there is a limit to how biased agents can be to guarantee truth-telling.

Involvement.—In the model, it is assumed that the potential bureaucrat cares about the client's welfare. In some instances, however, what may matter is her *involvement* in the welfare of the client. An alternative assumption, addressed in Tim Besley and Maitreesh Ghatak (2005) and George Akerlof and Rachel Kranton (2003) is that it is the agent's involvement in the outcome that affects their utility, which can result in some different conclusions.²¹ As a result, it is important to note that the results relate best to situations where what matters to the potential applicant is the utility of the agent, whether the agent is involved in the decision or not.

The effect of allowing the agent's utility to depend on her involvement now implies that the decision of who to hire depends not only on the agent who exerts most effort, but also on the agent who values the *level* of utility attained by the client, rather like any compensating differential. So, for instance, if the client attains positive utility on average, it becomes a reason to hire the most altruistic, whereas a bureaucracy that harms clients will want to hire those hostile, as they value that harm most.

Endogenous States.—In many of the examples above, the state of nature can be controlled by the client: for instance, individuals may choose whether to commit crimes, and patient's activities may affect their health. Ex post outcome for the guilty or sick can obviously affect the extent to which clients engage in these activities. A standard result in many agency settings is that better monitoring tends

to improve the efficiency of providing incentives. So, for example, a technology that identifies guilt better tends to deter criminal activities. The objective of the principal in the model above is to render such monitoring as efficient as possible to avoid ex post errors, but there is a likely additional benefit of improving ex ante incentives to attain desired states. Indeed, the source of S could be precisely that improved monitoring ex post gives rise to better ex ante incentives rather than a utilitarian concern for ex post errors in themselves.

The Information and Actions of Agents.—Agents have been assumed here to know their preferences before applying for the position. In some instances, this is reasonable. In others, it seems implausible. For instance, police officers may know little about how they trade off the welfare and rights of suspects before they join; instead, the experience of being a police officer reveals their latent preferences. This changes the timing of events in the model, where the effects above take place from the point at which this information is revealed, but leaves the qualitative implications unchanged. A more serious concern is if the job determines the preferences of workers. For instance, perhaps all police officers ultimately become hostile to suspects. Yet it seems unlikely that all employees share the same beliefs and desires; the evidence on the turnover of social workers above would certainly suggest this, and the paper continues to offer suggestions on who is likely to be most efficient. Alternatively, this paper can offer a more basic observation on the attributes of bureaucrats, *wherever they come from*; in other words, is it a good thing that social workers who give out benefits show altruism, or that INS officials are suspicious of entrants' motives?

Other Activities by Bureaucrats.—I have addressed only how the preferences of bureaucrats affect the accuracy of their allocation decisions. But most "street-level bureaucrats"—the term typically used in the literature for the occupations considered here—typically do more. For example, they aid clients in understanding the application process. The desired attributes of a bureaucrat in terms of effort exertion described above can easily conflict with these other tasks carried out by him. For instance, a lack of common courtesy or

²¹ To take an extreme example where my model seems implausible, consider the decision of a pacifist to join the army in wartime. In any objective sense, his actions harm the "clients" that he faces; he shoots them. Yet the assumptions of the model above imply that the army has an attraction to pacifists, because their actions may involve less killing than would an alternative person hired into the army.

willingness to help in explaining requirements may be a substantial price to pay for hostile bureaucrats exerting more effort. On the other hand, the altruism that may be desired by social workers in terms of inducing effort may make them unwilling to check up on items that may make clients ineligible for benefits. As a result, the benefits described above should be seen as one input to the decision on whom to hire when there are aspects of bureaucrats' jobs other than effort exertion, as above.

Finally, it is worthwhile to relate these results back to the two economics literatures. First, the literature on compensating differentials has long recognized that employees have job-related preferences. But unlike the usual model of compensating differentials, the premise here is that even after controlling for wages, employers care about who takes the job. Second, consider the relationship to the literature on incentive provision. Most organizations use compensation partly for the purpose of selecting workers with appropriate characteristics, and much of the recent empirical work on incentives has emphasized these selection effects (see Prendergast 1999 for details). Perhaps the most important finding of this paper is that selecting on intrinsic motivation is complex and need not lead to efficient outcomes.²² As such, this paper offers another reason why the efficiency of bureaucracies is likely to be limited.

APPENDIX: PROOFS OF RESULTS

PROOF OF LEMMA 1:

Let U^* be the expected utility to the agent if the agent who is offered the position turns it down. As the good is assigned without an agent, this is equivalent to an agent who exerts effort of $e = 1/2$. (Allowing the alternative mechanism, where the client proposes an allocation, which is preferred when interests are aligned, does not change the result.)

Given this, the wage, $w^*(v)$, offered to an agent of type v is chosen such that her utility equals $U^*(v)$, as this constitutes her reservation utility. The utility of the agent is given by

$U_a = S(A, \alpha) + (v - v_m)B(A, \alpha) - d(e)$,
so that

$$(23) \quad w^*(v) = U^*(v) - E[S(A, \alpha) - (v - v_m)B(A, \alpha) + d(e)].$$

Then, note that

$$(24) \quad w^{*'}(v) = E\left[U^{*'}(v) - \frac{dS}{de} \frac{de}{dv} - (v - v_m) \frac{dB}{de} \frac{de}{dv} + d'(e) \frac{de}{dv} - B(A, \alpha)\right].$$

But if (5) binds,

$$(25) \quad E\left[\frac{dS}{de} + (v - v_m) \frac{dB}{de}\right] = d'(e),$$

so (24) simplifies to

$$(26) \quad w^{*'}(v) = E[U^{*'}(v) - B(A, \alpha)].$$

Let $\Omega(e)$ be the expected welfare of the client if the agent exerts effort of e . Then $U^{*'}(v) = \Omega(1/2)$, as no agent is equivalent to one who exerts effort of $1/2$. As a result,

$$(27) \quad w^{*'}(v) = \Omega\left(\frac{1}{2}\right) - \Omega(e(v)).$$

First consider the case where $e'(v) > 0$. In this instance, the principal generates higher effort from more altruistic agents: but so also is client welfare higher with more effort and so $w^{*'}(v) = \Omega(1/2) - \Omega(e(v)) < 0$. Hence the principal hires the most altruistic agent. Next consider the case where $e'(v) < 0$ and the effort is highest for the least altruistic. Then $w^{*'}(v) = \Omega(1/2) - \Omega(e(v)) > 0$. Hence the principal hires the least altruistic agent.

REFERENCES

- Acemoglu, Daron, and Thierry Verdier.** 2000. "The Choice between Market Failures and Corruption." *American Economic Review*, 90(1): 194–211.

²² In this sense, it follows Daron Acemoglu and Thierry Verdier (2000) and Prendergast (2003) in its focus on the particular difficulties of attaining bureaucratic efficiency.

- Aghion, Philippe, and Jean Tirole.** 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy*, 105(1): 1–29.
- Akerlof, George A., and Rachel E. Kranton.** 2005. "Identity and the Economics of Organizations." *Journal of Economic Perspectives*, 19(1): 9–32.
- Bénabou, Roland, and Jean Tirole.** 2003. "Intrinsic and Extrinsic Motivation." *Review of Economic Studies*, 70(3): 489–520.
- Besley, Timothy, and Maitreesh Ghatak.** 2005. "Competition and Incentives with Motivated Agents." *American Economic Review*, 95(3): 616–36.
- Blankertz, Laura, and Susan Robinson.** 1997. "Turnover Intentions of Community Mental Health Workers in Psychosocial Rehabilitation Services." *Community Mental Health Journal*, 33(6): 517–29.
- Brehm, John, and Scott Gates.** 1997. *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Public* (Michigan Studies in Political Analysis). Ann Arbor: University of Michigan Press.
- Derthick, Martha.** 1979. *Policymaking for Social Security*. Washington, DC: Brookings Institution Press.
- Dewatripont, Mathias, Ian Jewitt, and Jean Tirole.** 1999. "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies." *Review of Economic Studies*, 66(1): 199–217.
- Dewatripont, Mathias, and Jean Tirole.** 1999. "Advocates." *Journal of Political Economy*, 107(1): 1–39.
- Dixit, Avinash.** 2002. "Incentives and Organizations in the Public Sector: An Interpretative Review." *Journal of Human Resources*, 37(4): 696–727.
- Ellett, Alberta.** 2001. "The Role of Human Caring in Employee Retention in Child Welfare: Practice Implications from a Two-State Quantitative Study." Unpublished.
- Frey, Bruno S., and Reto Jegen.** 2001. "Motivation Crowding Theory." *Journal of Economic Surveys*, 15(5): 589–611.
- Goodsell, Charles T.** 1981. *The Public Encounter: Where State and Citizen Meet*. Bloomington: Indiana University Press.
- Johnson, Ronald N., and Gary D. Libecap.** 1989. "Bureaucratic Rules, Supervisor Behavior, and the Effect on Salaries in the Federal Government." *Journal of Law, Economics, and Organization*, 5(1): 53–82.
- Kadushin, Alfred, and Daniel Harkness.** 1985. *Supervision in Social Work*. New York: Columbia University Press.
- Lee, Raymond, and Blake Ashforth.** 1993. "A Further Examination of Managerial Burnout: Towards an Integrated Model." *Journal of Organizational Behavior*, 14(1): 3–20.
- Lipsky, Michael.** 1980. *Street-Level Bureaucracy*. New York: Russell Sage Foundation Publications.
- Manning, Peter, and John Van Maanen.** 1978. *Policing: A View from the Street*. Santa Monica: Goodyear Publishing.
- Mor Barak, Michal E., Jan A. Nissly, and Amy Levin.** 2001. "Antecedents to Retention and Turnover among Child Welfare, Social Work, and Other Human Service Employees: What Can We Learn from Past Research? A Review and Metanalysis." *Social Service Review*, 75(4): 625–61.
- Peabody, Robert L.,** 1964. *Organizational Authority: Superior-Subordinate Relationships in Three Public Service Organizations*. New York: Atherton Press.
- Prendergast, Canice.** 1999. "The Provision of Incentives in Firms." *Journal of Economic Literature*, 37(1): 7–63.
- Prendergast, Canice.** 2003. "The Limits of Bureaucratic Efficiency." *Journal of Political Economy*, 111(5): 929–58.
- Skolnick, Jerome H.,** 1966. *Justice without Trial: Law Enforcement in Democratic Society*. New York: John Wiley and Sons Inc.
- Skolnick, Jerome H.,** 1968. *The Police and the Urban Ghetto*. Chicago: American Bar Foundation.
- Warner, W. Lloyd, Paul Van Riper, Norman Martin, and Orvis Collins.** 1963. *The American Federal Executive: A Study of the Social and Personal Characteristics of the Civilian and Military Leaders of the United States Federal Government*. New Haven: Yale University Press.
- Weatherly, Richard, Claudia Kottwitz, Denise Lishner, Kelly Reid, Frank Roset, and Karen Wong.** "Accountability of Social Service Workers at the Front Line." *Social Service Review*, 33(4): 556–71.
- Wilson, James Q.** 1989. *Bureaucracy: What Government Agencies Do and Why They Do It*. New York: Basic Books.

Copyright of *American Economic Review* is the property of American Economic Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.