

The Motivation and Bias of Bureaucrats

May 2004

Abstract

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 four 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. Finally, incentive contracting can reduce effort exerted, by affecting who chooses to become a bureaucrat.

1 Introduction

Employees often have little monetary reason for doing what they do. Among these are bureaucrats like social workers, teachers, nurses, police officers, and INS officials, whose actions typically have little or no effect on pay, yet things usually get done. As Wilson, 1989, puts it, “what is surprising is that bureaucrats work at all..rather than shirk at every opportunity” (p.156). In these and other occupations it is non-financial objectives - making a difference, helping those in need, keeping the nation safe, etc. - that act as motivation. The purpose of this paper is how such “intrinsic” preferences affect efficiency, and to address whether those hired are likely to exhibit the kind of preferences that increase efficiency.

Four primary results are shown on the nature of bureaucracies. First, bureaucrats - those who intermediate between the objectives of a principal and the clients they affect - should be biased, not sharing the preferences of those whose objectives they implement. Second, this 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. Third, self-selection need not produce the desired kind of bureaucrat. Instead, if there is sufficient disagreement among potential applicants about desired outcomes, bureaucracies become composed of those who are *most* preferred by the principal, and those who are *least* preferred. Finally, more intensive use of monetary incentive contracts can reduce effort exerted, by having an adverse effect on who chooses to become a bureaucrat.

There is a large literature in political science and public administration on improving the efficiency of public bureaucracies.¹ This literature tends to focus on employee selection - hiring and retaining workers with the right motivations - rather than the contractual design issues of the economics literature, where pay depends on some measure of performance.² Bureaucrats typically carry out some action that affects the welfare of clients - police officers arrest suspects, INS officials deny visas to applicants, benefit officers give welfare to those in need, and so on. Specifically, the role of bureaucrats as advocates for their clients is often emphasized, and perhaps the most popular prescription for improving the quality of public services is the recruitment and retention of workers

¹A notable feature of the economics literature on improving the efficiency of public bureaucracies is its emphasis on “macro” features, such as political accountability of principals, or the role of competition, rather than on the agency concerns of low level functionaries, those who typically make the front-line decisions.

²This is not because researchers have not thought of making pay depend on performance measures. At times, the issue of using compensation to align incentives is simply ignored as infeasible or counterproductive. Yet recent work by Brehm and Gates (1997), which explicitly evaluates the relevance of recent work in agency theory, ultimately concludes that although monetary incentives are not irrelevant, “the overwhelming evidence..indicates that the bureaucrat’s own preferences have the greatest effect on behavior” (p.196), and notes that a weakness of the economics literature is an over-reliance on “leisure- or budget-maximization, and ignore policy-oriented preferences”.

whose goals coincide with those of their clients.³ Yet it would be hard to claim that such coincidence of preferences characterizes all bureaucrats. In fact, a more commonplace characterization of many bureaucracies would be indifference to their clients, or in some cases an outright hostility towards their clients' wishes. Such accusations are routinely leveled at institutions such as the Immigration and Naturalization Service, the Internal Revenue Service, and many police forces. The alignment of interests between clients and bureaucrats plays a central role here.

Oversight of bureaucrats rarely occurs through the use of explicit contracts relating rewards to observable outcomes.⁴ Instead, objectives are (at least partially) aligned by the twin mechanisms of intrinsic motivation and sporadic oversight. Consider the case of a social worker charged with giving benefits to clients. Welfare departments do not offer explicit pay for performance plans; instead, the focus is on selecting (and retaining) employees with intrinsic characteristics that lead them to make better decisions.⁵ Yet there is also oversight of social workers, where supervisors oversee how caseloads are being handled. But this oversight is typically not random: instead it tends to be triggered by signals that errors have been made.

In this context, biased bureaucrats are typically preferred to those who share the principal's preferences.⁶ The task of a typical bureaucrat is to make a decision on allocating a benefit to a client. The agency problem considered here is where the bureaucrat or agent exerts effort to identify more accurate treatments for a client (for example, a doctor must incur costs diagnosing appropriate treatment for a patient, or a police officer spends costly time accumulating evidence against a suspect). I assume that potential bureaucrats vary in how they view allocative inefficiencies: some are more bothered by errors which favor clients whereas others are more affected by those which harm them.⁷ This is formally modeled here by assuming that individuals vary in their altruism towards the client. The principal would like accurate assessments from the bureaucrat. But altruism affects effort decisions because it causes them to internalize client preferences. Altruism then affects incentives through the rather simple mechanism of whether clients *want* more accuracy: in other

³For instance, in perhaps the best known study of street-level bureaucrats, 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", and "those who recruit themselves for public sector work are attracted to some degree by the prospect that their life will gain meaning through helping others" (p.66). For other work that stresses the importance of these kinds of selection concerns, see Warner et al., 1963, Goldsmith, 1991, Goodsell, 1981, Kadushin, 1985, and Derthick, 1979.

⁴This is, of course, because there are typically few objective measures of performance that one would want to readily use to reward these agents. See Johnson and Liebcap, 1989, and Dixit, 2002, for more details on the compensation of bureaucrats.

⁵See Section 3 for the relevant research.

⁶This differs from recent work by Aghion and Tirole, 1997, who build a model of delegating authority based on the congruence of preferences between agent and principal.

⁷For instance, some see the costs of incorrectly giving welfare benefits to an undeserving candidate as less important than failing to give these benefits to a deserving case, whereas others would disagree.

words, do clients want bureaucrats to offer more accurate diagnoses?

There are two issues that determine how intrinsic preferences affect efficiency. The first concerns the alignment of the interests of the client and the principal. Begin by considering instances where a client only wants a particular treatment if it is efficient from the perspective of the principal. As a concrete example, consider a patient being considered for coronary by-pass surgery: he only wants that treatment 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 more altruistic bureaucrats increase efficiency.⁸ But there are many situations where the interests of clients are opposed to those of the principal. There are some instances where preferences are so opposed that those who should be denied a benefit derive the most benefit from receiving it - for instance, 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.⁹

The second issue that affects desired bureaucratic preferences is the asymmetry of oversight; more simply, how likely are mistakes to be corrected? In Prendergast, 2003, it is argued that some bureaucratic mistakes are more likely to be corrected than others. Take a trivial but conceptually relevant example - receiving change from a retail assistant, and suppose that a retail assistant randomly makes errors in giving change to customers. That randomness, per se, has no effect on customer welfare, except that the customer can point out errors when she received less than she should.¹⁰ Because clients can partially correct errors from which they lose, they benefit from *less* informed agents. In effect, clients (and sometimes other parties) can operate as whistleblowers. This asymmetry implies that those who are less altruistic are preferred.

This second effect is relevant in two situations. First consider the case where the client values the allocation equally whether he deserves it or not: for instance, avoiding a traffic fine. Then, with symmetric oversight, the client is (on average) indifferent about the bureaucrat's accuracy

⁸As Wilson, 1989, puts it, this sense of mission "enables administrators to economize on other kinds of incentives" (p.95). See Dewatripont et al, 1999, for formal work on the role of missions in the public sector.

⁹To see this, note that greater accuracy means that it is more likely that a taxpayer be investigated when guilty of evasion (which he prefers not to have happen) but less likely when innocent (which he would like to occur). But as the client benefits more from the error of not being investigated when guilty more than being investigated when innocent, he prefers a less accurate agent.

¹⁰This is not an isolated case. Perhaps the typical agency problem takes the form of someone intermediating between a principal and a consumer. For instance, a waiter intermediates between the restaurant and a patron, a police officer intermediates between the objectives of her department and a suspect, an auto mechanic intermediates between a car owner and his employer, and so on. In the absence of explicit pay for performance contracts, probably the most important method of oversight is to use consumers to identify mistakes. For example, patrons in a restaurant complain to the manager if a waiter's performance is poor, a car owner will return the car to the shop if it is not fixed, suspects can file suit if a police officer uses excessive force, and so on.

(as this simply adds randomness). But if clients can help to overturn those decisions that harm them, this implies that the benefit more from inaccurate assessments, and so hostile agents are preferred. Second, there are instances where clients interest are somewhat aligned, in that they value a treatment more when it is truly needed, but value in anyway. One example is being approved for medical tests, which patients would like even if not truly ill (“just in case”), but value them more when they are ill. As the client benefits more from correctly getting benefits than he gains from incorrectly getting them, this would suggest that more altruistic agents would be preferred. However, the second influence on bureaucratic bias, the role of client oversight, that can overturn this conclusion. The desirability of client advocates in this case depends upon 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 which harm them? Third, how easy it is to correct these errors? This framework is used to argue that the preferences of those who work for the INS, IRS, and benefits offices are likely to differ from those of medical personnel, social workers, and teachers.

These results illustrate a preference for biased bureaucrats. But bureaucracies will only exhibit this bias if the supply of individuals for these positions matches the principal’s demands. If the principal can easily identify agent preferences, he could base hiring and firing decisions on these. This seems unlikely in many occupations. Accordingly, in Section 4, the issue of whether bureaucrats with the appropriate bias can be induced to self-select is considered. As mentioned above, incentives for most bureaucrats are crude, and often take the form of being fired for malfeasance. This is addressed here by assuming that incentives are provided by efficiency wages; the desire to carry out appropriate actions is based on the higher likelihood of retaining a valued job. In this case, it is shown that with sufficient dispersion of preferences, selection to bureaucratic positions becomes bifurcated, in that the bureaucracy consists of (i) those who are *most* preferred by the principal, and (ii) 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 who become bureaucrats.

Furthermore, there is a tension between the ability of wages to induce effort exertion and their ability to aid recruiting. Specifically, conditional on being employed, more incentives result in more effort, as is standard. But incentive pay can harm recruitment, as agents with less intrinsic motivation choose this occupation; in effect, they do it “for the money.”¹¹ As a result, average intrinsic motivation can fall, thereby limiting the value of monetary contracts.

¹¹This point has also recently been made by Delfgaauw and Dur, 2003.

Section 2 begins by constructing a simple model of effort exertion, where a bureaucrat chooses how hard to work on behalf of a client. Section 3 illustrates how the relationship between client and principal preferences leads to a demand for bias. The modeling work is extended in Section 4, dealing with self-selection and the effect of wages on the intrinsic motivation of the burueacracy. Section 5 concludes with a brief discussion.

2 The Model

An allocation A must be made to a client, where A can take on a value of 0 or 1. The surplus from the allocation depends on a parameter α and is given by

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

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 true value of α is unknown and can take two values $\alpha = 1$ or $\alpha = 0$ with equal probability. There are three actors in the model, 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 . The client observes a signal σ , which correctly identifies α with probability $s \geq \frac{1}{2}$. The only action taken by the client is to send a binary message m to the principal, where message c means a complaint, and message n means that no complaint has been sent. This message is sent after the agent proposes a treatment, a , for the client. There is a (monetary) cost $\delta \geq 0$ to making a complaint. His objective is to choose m to maximize his expected benefits $EB(A, \alpha)$ minus the cost of complaint.

For notational convenience, let $\beta_0 = B(0, 0) - B(1, 0)$ be the benefits attained by the client when correctly given allocation 0 relative to being given allocation 1, and let $\beta_1 = B(1, 1) - B(0, 1)$ be the benefits attained by the client when correctly given allocation 1 relative to being given allocation 0. Without loss of generality, let $\beta_1 \geq 0$. Thus, in the police example, 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 Agent or Bureaucrat An agent exerts effort to determine the appropriate assignment. To do so, she exerts effort e , which affects the precision of her estimate of the correct assignment, α . She observes a signal $\phi \in \{0, 1\}$ which is correct with probability e , where she chooses between high effort \bar{e} or low effort \underline{e} , where $1 \geq \bar{e} > \underline{e} \geq \frac{1}{2}$. The cost of the low effort is normalized to 0 and the high effort costs d . Let a be the assignment recommended by the agent's information. The information of the agent is observable (but not contractible) to the principal; as a result, the principal can always require that the agent provide whatever information she has.¹²

To render the problem interesting, assume that the effort of the agent cannot be observed. Her preferences depend on: (i) her wage, (ii) her effort decision, and (iii) the welfare of the client. Her utility is given by

$$U_a = w - dI_e + vB(A, \alpha), \quad (2)$$

where I_e is an indicator equal to 1 if effort is high, and 0 otherwise. Unlike most agency problems, it is assumed that the agent has preferences toward the client, where v measures the intensity with which he associates with the client. (Altruism should be seen as a reduced form for a more general set of preferences, as described in Section 5.) Begin by assuming that v is observable, and that it is drawn from a distribution F , with density f that ranges from \underline{v} to \bar{v} . (The problem of unobservable bias is considered in Section 4.) I assume that reservation wages are independent of v and equal to 0.

The Principal The principal plays two roles: (i) he hires an agent and designs a contract, and (ii) he monitors the recommendation of the agent.

Contracts The principal hires an agent. Assume that incentives to that agent are provided by an efficiency wage contract, where a worker earns a wage w unless fired. Specifically, the agent has a reservation wage of 0, but faces a liquidity constraint whereby the wage must be non-negative. If the agent is fired, she is instantaneously replaced by another.¹³ I assume that the principal finds it sufficiently beneficial to induce effort exertion from his preferred type that he will incur the monetary costs to do so.

¹²Note that this assumption rules out the possibility that the agent recommends a course of action not suggested by her information, as was studied in Prendergast, 2003.

¹³It is assumed that the agent is replaced in order to avoid what I believe to be an unimportant incentive in these contexts; namely, to investigate cases in order to save on wages, where the principal can pocket the efficiency wage by finding evidence of misdoing by a bureaucrat. Realistically, investigations in the public sector are not carried out for this purpose, as the fired worker is immediately replaced and no wage savings arise. As a result, assume that the worker is replaced simply to avoid this incentive. Allowing this additional return to investigation does not, however, change any of the results; instead, it simply offers an additional reason for investigating.

Oversight After the agent has proposed an action and the client has sent his message, the principal can monitor that decision. He is assumed to be uninformed (beyond the information provided by the agent and client) unless he carries out an investigation. To model a role for investigations, assume that the overseer chooses a probability of observing α at some cost. Specifically, the overseer chooses a probability ρ with which he observes α , at a cost $\kappa(\rho)$, which for simplicity, is assumed to be quadratic: $\kappa(\rho) = \frac{\gamma\rho^2}{2}$, where $\gamma \geq 1$.¹⁴ Complaints allow oversight to potentially be based on the client's message; accordingly, let $\rho(a, m)$ be the probability that the principal investigates when the agent recommends allocation a and the client responds with a message m . If an investigation turns up evidence that the agent made a mistake, the principal overturns the agent's decision. If the investigation concurs with the agent's findings, or there is no investigation, the decision remains that was suggested by the agent.¹⁵ Oversight occurs in the subgame perfect fashion given available information.

When choosing wage contracts and agent, the principal maximizes expected social surplus minus the ex ante cost of providing incentives and the cost of investigations.¹⁶ The ex ante objective of the principal is therefore to choose v^* and w^* to maximize

$$E_{\alpha,a,m}[S - \frac{\gamma\rho^2(a, m)}{2} - w] \quad (3)$$

subject to the effort incentives of the agent, the client's propensity to complain, and the principal's ex post incentive to investigate the case.

Timing The timing of the game is as follows. First, the principal hires an agent and chooses the wage to maximize ex ante expected surplus minus wage and investigation costs. Second, nature assigns α to the client and the agent exerts effort. Third, the customer and the agent privately observe their signals. Next, the agent provides his information a . Following this, the customer send a message m (i.e., he complains or not). The principal observes the relevant information and then monitors. If he observes α , he allocates $A = \alpha$. Otherwise, the agent's recommendation is implemented. Finally, the agent is paid or replaced.

¹⁴In order to cut down on notation, it is assumed that the principal's signal is perfect. Allowing for imperfect observations by the principal does not change the qualitative results.

¹⁵For some bureaucratic decisions, this is not the correct allocation mechanism. For instance, patients have the right to refuse surgery even if a doctor recommends it. These are typically only for those instances where the interests of client and principal align; for instance, suspects cannot refuse to be arrested. In instances where incentives align, the results of this section continue to hold with this allocation mechanism.

¹⁶Note that I have not included the cost of complaint here. This would be the case if the cost of complaint to the client was a monetary transfer to the principal. However, allowing this cost to be deadweight loss does not change the results.

Equilibrium I characterize the Perfect Bayesian Nash equilibrium which yields highest surplus to the principal. To avoid multiplicity of equilibria, attention is restricted to cases where messages are meant literally, in the sense that a complaint weakly increases the likelihood of investigation, $\rho(a, c) \geq \rho(a, n)$. This implies that those instances where complaints are informative are identified (if they exist) as the principal would like to make oversight depend on all available information. Consequently, the focus is on equilibria where client complaints reveal most information on the decision of the bureaucrat. The formal program is defined in the Appendix.

3 The Demand for Bureaucratic Bias

To begin, consider the preferences of the principal for biased agents. To begin, consider the optimal way of designing wages to induce effort exertion. This is described in Lemma 1.

Lemma 1 *The principal (weakly) prefers to (i) hire the agent who receives highest private benefits from exerting effort \bar{e} and (ii) offer a contract where the agent is fired only if investigated and incorrect, or not investigated and a complaint has been made.*

The logic for this result is simple. The principal wishes the agent to exert high effort. There are two ways of doing so: (i) changing the attributes of the agent, or (ii) providing monetary incentives. But monetary incentives are costly, so the principal will (at least weakly) choose agents whose incentives are greatest in the absence of monetary incentives.¹⁷

3.1 How Altruism Affects Incentives

Agents exert effort to affect the welfare of clients, which they value at a rate v . The welfare of clients is affected by the accuracy of the decisions made by the agent, where more effort causes more precise allocations. Accuracy is affected by two features: (i) how hard the agent works, e , and (ii) how often her mistakes are corrected, ρ . Let $\underline{\rho}(\alpha)$ be the probability that an incorrect decision by the agent is corrected when the true state is α . The expected private (i.e., non-wage) benefit that an agent of type v then receives is given by

$$v(e \frac{B(0,0) + B(1,1)}{2} + (1-e)[\underline{\rho}(1) \frac{B(1,1)}{2} + (1-\underline{\rho}(1)) \frac{B(0,1)}{2}] + (1-e)[\underline{\rho}(0) \frac{B(0,0)}{2} + (1-\underline{\rho}(0)) \frac{B(1,0)}{2}]) - dI_e. \quad (4)$$

¹⁷In 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.

Private benefits arise from two endogenous sources; how hard the agent works, and how often mistakes are corrected. The change in private benefits from exerting effort is then given by

$$v(\bar{e} - \underline{e}) [\beta_0(1 - \underline{\rho}(0)) + \beta_1(1 - \underline{\rho}(1))] - d. \quad (5)$$

To see how this affects the optimal selection of agents, note that increasing v only increases marginal private benefits if

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

This condition generates all the results of the first half of the paper. If (6) holds, more altruistic agents induce greater effort for any fixed wage contract; while if it is violated, more hostile agents maximize effort incentives.

It is worthwhile deliberating on this condition. The change in the client's welfare from more effort is $(\bar{e} - \underline{e}) [\beta_0(1 - \underline{\rho}(0)) + \beta_1(1 - \underline{\rho}(1))]$. This depends on two factors. First, more effort makes it more likely that the agent gets the benefits β_1 and β_0 . Of course, these need not both be positive, so the first factor that affects the outcome is how much the client desires the allocation when he “should” get it ($A = 1$) compared to when he should not ($A = 0$). Second, it depends on how often the bureaucrat is pivotal; in other words, the likelihood that her mistakes are not corrected, $1 - \underline{\rho}(\cdot)$. To see this, consider the following cases.

1. When client and principal preferences align. 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 optimal choice of the agent is the most altruistic. 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 that cares most about the client exerts the most effort. Proposition 1 immediately follows.

Proposition 1 *Assume that $\beta_i \geq 0$. The principal (weakly) prefers to hire the agent whose preferences most closely align with those of the client, type \bar{v} .*

An Example: Social Workers Researchers interested in understanding why social workers perform as they do routinely emphasize their empathy and altruism towards clients. For instance, in a survey on the preferences of social workers, Peabody, 1964, notes that “by far the most dominant organizational goal perceived as important..is service to clientele” (p.66), where 83% of

survey respondents view such service as important, compared to only 9% who see “obligation to taxpayers” or “assistance to the public in general” as important concerns affecting their decisions. Similar evidence is found in 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, 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 job. They further note that those who hold these beliefs are more likely to do unpleasant aspects of their jobs (such as taking paperwork home). 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”. (Note the difference below to the complaints of police officers.) Along these lines, Kadushin, 1985, p.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, Mor Barak et al., 2001, Ellett, 2001, and Blankertz and Robinson, 1997, who illustrate a negative relationship between the turnover of social workers and their reported commitment to clients.

2. When client preferences are opposed to those of the principal. The reason for altruism above is the alignment of client and principal preferences. But, as mentioned above, a wide range of bureaucratic environments involve clients who have preferences which relate little to those of the principal. Accordingly, next consider the case where clients desire the good at least as much when they do not deserve it as when they do, where $-\beta_0 \geq \beta_1 > 0$. (For example, those who have cheated on their taxes wish to avoid an audit at least as much as the honest.) In this case, the desired characteristics of the bureaucrat depend on any difference between $\underline{\rho}(1)$ and $\underline{\rho}(0)$; more simply, are some mistakes more likely to be corrected than others? But how mistakes are corrected depends centrally on the role of clients as overseers, which is now addressed.

Client Oversight After the bureaucrat has made a recommendation, and the client has possibly complained, the principal chooses $\rho(a, m)$ to maximize surplus minus investigation costs; i.e., to maximize

$$E_\alpha[S - \frac{\gamma \rho^2(a, m)}{2} | a, m], \quad (7)$$

given the information in the client’s message.¹⁸

First note that if the bureaucrat recommends giving the benefit ($a = 1$) to the client, the client

¹⁸Note the absence of wages here, as the worker is instantaneously replaced if fired.

never complains as $\delta \geq 0$. In that case, the principal investigates with probability $\rho^*(1, n) = \frac{1-e}{\gamma}$, as the unconditional likelihood of an error is $1 - e$. If the agent recommends $a = 1$, the likelihood that an error is corrected is therefore

$$\underline{\rho}(0) = \frac{1 - e}{\gamma}. \quad (8)$$

If the benefit to the client is denied, he may complain. But complaints only affect oversight if they are informative, which means that it needs to be the case that the client complains only if denied *and* his information disagrees with the allocation. If complaints are informative, then the ex post efficient probability of oversight with recommendation $a = 0$ and message m is given by

$$\rho^*(0, n) = \frac{(1 - s)(1 - e)}{\gamma[es + (1 - s)(1 - e)]} \quad (9)$$

and

$$\rho^*(0, c) = \frac{s(1 - e)}{\gamma[e(1 - s) + s(1 - e)]}. \quad (10)$$

As $s \geq \frac{1}{2}$, this implies that $\rho^*(0, c) \geq \rho^*(0, n)$, so that clients can focus investigations. The ex ante likelihood of an incorrect allocation of the good being corrected is then

$$\underline{\rho}(1) = s\rho^*(0, c) + (1 - s)\rho^*(0, n). \quad (11)$$

This exceeds $\underline{\rho}(0)$ for $s > \frac{1}{2}$; in words, there is an asymmetry in oversight that favors clients. As shown below, this affects the desired bureaucrat.

For complaints to be informative, there must be limits on the cost of complaint, δ ; if δ is too large, the client never complains, while if too small, he always complains when denied the benefit (in any putative equilibrium where complaints are informative). Routine calculations show that complaints are informative only if

$$\frac{s(1 - e)}{e(1 - s) + s(1 - e)} \geq \frac{\delta}{[\rho^*(0, c) - \rho^*(0, n)]\beta_1} \geq \frac{(1 - s)(1 - e)}{es + (1 - s)(1 - e)} \quad (12)$$

Note that such a δ always exists, so that if for example, the principal could choose the cost of complaint, he could always guarantee that it is informative.¹⁹

On the other hand, if complaints are not informative, the principal cannot garner extra information on the bureaucrat's behavior from the action of the client. In this case, $\underline{\rho}(1) = \frac{1-e}{\gamma}$, which

¹⁹It is here that the assumption that the principal observes the state of nature perfectly becomes quantitatively important. When the principal does not observe the state perfectly, there is an additional incentive for the client to complain when denied, even when he believes that the correct decision has been made, as there is now the chance that the principal will *incorrectly* overturn the decision. This makes it more difficult to make complaints informative.

is equal to $\underline{\rho}(0)$, so oversight is unbiased.

First consider the case where $\beta_1 < -\beta_0$: the guilty prefer being released more than the innocent. Then the consumer always benefits from a less informed agent, even when the client's complaints are uninformative. 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. Proposition 2 follows.

Proposition 2 *Assume that $-\beta_0 > \beta_1 > 0$. Then the principal prefers to hire the agent who is most hostile to the client.*

Next consider the case where $\beta_1 = -\beta_0$. For instance, this could be a traffic fine, which has a similar cost to the client whether he is guilty or innocent. Proposition 3 immediately follows.

Proposition 3 *Assume that $\beta_1 = -\beta_0$. Then if complaints are informative, the principal hires the agent who is most hostile to the client, \underline{v} . Otherwise, the principal is indifferent.*

To understand the intuition behind Proposition 3, begin with the case $\underline{\rho}(0) = \underline{\rho}(1)$, where complaints are uninformative. Since the client's preferences are independent of the social optimum, all that he cares about is the average probability of receiving the good. If $\underline{\rho}(0) = \underline{\rho}(1)$, the probability of receiving the benefit is $\frac{1}{2}$, independent of effort. Consequently, agent attributes have no effect on effort and the choice of bureaucrat is irrelevant. But when the client can affect oversight, (6) is always violated 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 least altruistic.

Examples: Police Officers, IRS Officials Many goods allocated by the public sector are benefits which are enjoyed independent of whether they are deserved. In some instances, the undeserving enjoy these benefits as much if not more than the deserving. 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 that is available on social workers. It would be hard to claim from the results of this literature, however, that they exhibit similar attributes to social workers. Instead, they seem to treat their primary clients (suspects) with little empathy, and were 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” (p.51). Such descriptions are commonplace, with Skolnick, 1966, concluding that “the policeman is generally a ‘suspicious’ person” (p.44), which Goodsell believes to be “an occupationally useful trait” (p.51).

Also notable is how different police officers perceive themselves to the general citizenry. According to 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” (p.17). Specifically, compared to the general public, police officers are “extremely ambivalent about the rights of others” (Manning and Van Maanen, 1978, p.19) and view lack of support in convicting suspects from the public as their greatest occupational concern. Furthermore, “the policeman feels that criminal procedure has been unfairly weighted against him” (Skolnick, 1966, p.199), where “the state requires the policeman to work in a milieu filled with extraneous and to him needless restrictions” (p.202).

3. When clients value the benefit, but more so when deserved Finally consider the case where the client wants the benefit ($\beta_0 < 0$), but values it more when he is deserving: $\beta_1 > -\beta_0 > 0$. So, for instance, 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 lower. In this instance, oversight is again given by (8), (9), and (10) when complaints are informative. Proposition 4 follows.²⁰

Proposition 4 *Assume that $\beta_1 > -\beta_0 > 0$. If consumer complaints are uninformative, the principal hires the most altruistic agent, \bar{v} . If complaints are informative, the principal hires type \bar{v} if*

$$\frac{\beta_1}{\beta_0} \geq - \left(\frac{1 - \frac{1-e}{\gamma}}{1 - \frac{s^2(1-e)}{\gamma[e(1-s)+s(1-e)]} - \frac{(1-s)^2(1-e)}{\gamma[es+(1-s)(1-e)]}} \right). \quad (13)$$

Otherwise, he hires the least altruistic agent, \underline{v} .

First consider the case of uninformative complaints. The reason for altruistic bureaucrats when clients play no role in oversight is simply the mirror image of that above; here the agent cares more about receiving the benefit when warranted than when not warranted, and so values more accuracy. It follows that the most altruistic bureaucrats satisfy this need. Unlike the previous cases, the outcome with informative complaints here depends on parameter values.

²⁰To see this, note that $\underline{\rho}(1) = s\rho^*(0, c) + (1-s)\rho^*(0, n)$ and $\underline{\rho}(0) = \frac{1-e}{\gamma}$. Substituting these into (6) yields (13).

What characteristics of the bureaucracy lead to more altruistic agents? How does the nature of goods transacted by bureaucrats leads to their selection? Accordingly, (13) illustrates how the optimal make-up of the bureaucracy depends on three factors: (i) the effect of being “deserving” on client welfare from the benefit, (ii) how the client affects oversight, and (iii) the cost of oversight.

1. The left hand side of (13) 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 to investigate the bureaucrat’s decisions, where outcomes that are more difficult to verify (γ increasing) result in more altruistic agents. To see this, consider the outcome as $\gamma \rightarrow \infty$, in which case the right hand side of (13) tends to 1. But then $\beta_1 > -\beta_0 > 0$ guarantees that (13) is satisfied, and so altruistic agents should be hired. More informally, the reason why clients want less informed agents is loosely that they can correct mistakes that harm them, but ultimately benefit from mistakes that are in their favor. But this only works if their complaints result in these bad decisions being likely to be overturned. When oversight is costly, clients cannot rely on complaint mechanisms to correct the inaccurate decisions made that harm them, and their incentives become more in line with those of the principal.
3. Finally, altruistic agents are less likely when clients are better informed. Again, consider the outcome when $s \rightarrow \frac{1}{2}$, 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 (13) tends to 1, and again $\beta_1 > -\beta_0 > 0$ guarantees that (13) is satisfied. Consequently, altruistic agents should be hired. The reason is similar to above; well informed clients can effectively focus investigations on cases where they have been incorrectly denied, better than their less well informed counterparts.

Examples: Doctors, INS Officials, Benefit Officers There are many goods that clients want to receive, but where their demands are greater when they are “deserved.” For example, insured patients want (costly) tests done on them, but they want them more when they are truly ill. Similarly, while all applicants for visas would like to be admitted to the US, those who are truly likely to be persecuted at home are more likely to desire immigration status. Finally, welfare workers who allocate means tested benefits realize that applicants always want to be approved, but desire approval more when they are truly in dire straights.

Consider each in turn. With doctors, there is substantial variation in how aligned are the preferences of the patient and the principal. At one extreme, few patients desire surgery when not needed, so diagnosis for surgery would be better served by altruistic doctors. Yet this congruence seems less true for other interventions. For example, many patients wish to receive antibiotics for illnesses for which they are useless, and anecdotal wisdom suggests that doctors routinely do so. Similarly, patients desire unnecessary tests if they do not pay their marginal costs. These instances would suggest constraints on the altruism of doctors. These concerns would be particularly important if patients are effective advocates in pointing out doctor error (s large), where poor decisions could be easily overturned. But patients are often bamboozled by the details of their diagnosis, and are often ineffective at pointing out diagnostic errors (except in unusual cases where courts become involved). Similarly, investigation of medical malpractice is costly. As a result, the analysis above would suggest that hiring altruistic doctors is optimal. By contrast, the calculus for INS officials and Benefits officers seems different, largely because the undeserving receive such large returns from being given unemployment assistance or a visa. As a result, this analysis would suggest less reason to choose altruistic INS Officers or Benefits Officers than doctors.

To summarize, this section offers a very simple framework for understanding the role of client advocacy on incentives. Four issues were highlighted. First, how does more accuracy affect the client? Second, what role does the client have in oversight? The greater is that role, the more likely are hostile bureaucrats. Third, how informed are clients about the correct treatment that they should receive? Somewhat paradoxically, when clients are sufficiently well informed, the optimal strategy is to hire agents who are indifferent to their clients. Finally, how easily can the principal investigate whether the agent made the correct decision? When the principal can relatively easily check outcomes, the more likely it is that the principal again chooses the least client-friendly agents.

To a lay person, 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 purpose of this section is to simply lay some groundwork that ultimately may be useful in identifying *why* police officers may share characteristics with INS or IRS officials, but have a radically different outlook to firemen, social workers, or employees in a hospital. Specifically, the claim here is that the selection of bureaucrats depends on a complex interaction between how client advocacy affects effort decisions, and the ability of losing parties to correct bureaucratic mistakes. This work raises these issues, which occupy center stage elsewhere in the literature, to the attention of economists.

4 The Self-Selection of Bureaucrats

The previous section identifies a demand for biased bureaucrats. But the bureaucracy will only exhibit this bias 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. Two possibilities naturally arise. First, no one may know the preferences of the agent at the outset. This would seem a plausible interpretation of a police officer, who has little experience with suspects and whose preferences will only be revealed through the process of doing the job. If these preferences ultimately become clear not only to the agent but also to the principal, the results of the previous section continue to hold qualitatively, but at a later point. Alternatively, it may be that the bureaucrat holds private information on her preferences. For instance, it is likely that child care workers have more information on how much they like children than do potential employers. This section is concerned with this case.

This extension yields the following results. There are two influences that result in more desirable types self-selecting, which both result from the fact that more desired types work harder. Specifically, they benefit from receiving (i) more utility from affecting the fate of clients, and (ii) higher monetary compensation. 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 desired preferences affect clients in the way desired by the principal, those who have the least desirable preferences affect it in the way that *they* desire, even though it is discordant with the principal's desire. In effect, agents with discordant preferences pay some monetary penalty to exercise their preferences, and it is the most discordant who are willing to pay most.²¹

Another implication explored here is that incentive provision can backfire, in that increased incentives reduce effort exerted. This arises for selection reasons. In the standard incarnation of the efficiency wage model, higher wages induce agents to work harder, as their jobs become more valuable. While this effect is evident here, there is a selection effect which can offer a countervailing outcome. Greater efficiency wages increase the expected wage. This means that those who are not intrinsically motivated now find the (higher paying) job more attractive, and so higher wages induce more undesired agents to apply. In effect, monetary contracts induce people who are “in it for the money” to apply, a point also made by Delfgaauw and Dur, 2003, which is ultimately harmful for

²¹For example, many disagree with a policy that restricts welfare benefits to the indigent or illegal immigrants. On 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.

efficiency.

In order to illustrate these influences, assume that v is private information for the agent.²² Furthermore, in order to reduce complexity, a number of simplifying changes are made to the basic model.

4.1 Simplified Model

In this section, I make the following assumption:

- The equilibrium level of oversight by the principal ρ^* is independent of the (perceived) effort level of the agent and whether the client complains. Specifically, there is a (low) cost $\underline{\kappa}$ for all $\rho \leq \rho^*$ and a (high) cost $\bar{\kappa}$ for all $\rho > \rho^*$, such that the principal always monitors with probability ρ^* .

To avoid detailing each of the cases above, attention is restricted to the situation where $\beta_1 > 0$, and $\beta_0 = 0$. Following Proposition 1, the principal prefers the most altruistic agents.²³ Unlike the previous section, the principal cannot now observe the altruism of the applicant. Instead, he can only observe who applies for the position, and randomly chooses among the applicants.

As complaints are uninformative, the principal cannot improve on a contract where the agent is fired unless he is investigated and correct - he is indifferent between this contract and that outlined in Lemma 1. For notational convenience, I assume this simpler contract. Note that as oversight is optimally set at ρ^* , the probability that the allocation is correct is where either the agent is correct (with probability e), or is incorrect but an investigation overturns the decision (with probability $(1 - e)\rho^*$). The cost of investigating is given by $\underline{\kappa}$, and so the principal chooses w to maximize

$$E[e + (1 - e)\rho^* - w - \underline{\kappa}], \quad (14)$$

where the expectation is over the effort level exerted by the agent. The effort level exerted by the agent depends on (i) what effort to exert, conditional on being hired - Incentive Compatibility, and (ii) who applies for the position - Individual Rationality. Each is considered in turn.

The Incentive to Exert Effort Both monetary and non-monetary returns affect effort exertion. The agent receives an expected wage of $e\rho^*w$, while her decisions affect the utility of the client,

²²Note here that the intrinsic motivation of any given worker is invariant to wages, unlike Frey and Jegen, forthcoming, or Benabou and Tirole, 2003.

²³This restriction is purely for notational simplicity. The results for the case where preferences are not aligned is a mirror image of the results below.

which she values at a rate v . If employed (the only relevant condition for the incentive compatibility constraint), the utility of the client from her actions is $(e + (1 - e)\rho^*)\frac{B(1,1)}{2} + (1 - e)(1 - \rho^*)\frac{B(1,0)}{2}$.²⁴ As a result, simple computations show that the agent exerts high effort only if

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

This defines the critical value of v , called \hat{v} , above which agents exert effort, defined by

$$\hat{v} = \frac{d - [\bar{e} - \underline{e}]\rho^*w}{[\bar{e} - \underline{e}](1 - \rho^*)\frac{\beta_1}{2}}. \quad (16)$$

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 (β_1), and the likelihood that their effort is pivotal in the allocation (i.e., $1 - \rho^*$).²⁵

Who Becomes a Bureaucrat? The novelty of this section arises from understanding who applies for the position as agent. A agent applies for the position if her expected utility from doing so exceeds her reservation utility. First consider the agent's reservation utility: this consists of her (zero) wage and the welfare she gets from the utility of the client. (Note here that the agent cares about the outcome even if she does not work in the job, unlike Besley and Ghatak, 2004). Because of this, her reservation utility depends on how hard alternative hires work. Let μ be the fraction of agents who exert high effort, and define $\epsilon = \mu\bar{e} + (1 - \mu)\underline{e}$ as the expected level of effort. Then the agent's reservation utility is

$$\frac{v}{2}[(\epsilon + (1 - \epsilon)\rho^*)B(1, 1) + (1 - \epsilon)(1 - \rho^*)B(1, 0)]. \quad (17)$$

If instead the agent is hired, her utility from effort e is

$$e\rho^*w + \frac{v}{2}[(e + (1 - e)\rho^*)B(1, 1) + (1 - e)(1 - \rho^*)B(1, 0)]. \quad (18)$$

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

$$\bar{e}\rho^*w + \frac{v}{2}(\bar{e} - \epsilon)(1 - \rho^*)\beta_1 - d. \quad (19)$$

²⁴With probability $\frac{(e + (1 - e)\rho^*)}{2}$ the agent is correctly given the assignment $A = 1$, while with probability $\frac{(1 - e)(1 - \rho^*)}{2}$ he is incorrectly given the benefit.

²⁵Note that the result that those who are more altruistic exert more effort only arises because β_1 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.

This defines the critical value of altruism, $\tilde{v}(\bar{e})$, *above* which the agent applies for the position, given by

$$\tilde{v}(\bar{e}) = \frac{2(d - \bar{e}\rho^*w)}{(\bar{e} - \epsilon)(1 - \rho^*)\beta_1}. \quad (20)$$

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

$$\underline{e}\rho^*w + \frac{v}{2}(\underline{e} - \epsilon)(1 - \rho^*)\beta_1. \quad (21)$$

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

$$\tilde{v}(\underline{e}) = \frac{-2\underline{e}\rho^*w}{(\underline{e} - \epsilon)(1 - \rho^*)\beta_1}. \quad (22)$$

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 $\mu > 0$). All those who are less altruistic apply, although they are precisely the type that the principal does not want to hire.

The reason should be clear. By exerting effort of \underline{e} , the agent reduces the utility of the client relative to a random alternative person who would get the job. But it is the least altruistic who care least about this cost and so selection also comes from those most hostile to the client's and principal's interests. Note also that if the distribution of preferences is sufficiently disperse: i.e., \bar{v} (v) is high (low), then this bifurcation arises if some choose not to participate.

Thus far, I have characterized two possible outcomes, one where all apply and where the equilibrium is characterized by (16), or where some choose not to apply, in which the equilibrium is characterized by (20) and (22). If all apply, then the bureaucracy is a random sample of the labor force, and so, there is no selection of the type of agent preferred by the principal. On the other hand, if some choose not to apply, then self-selection operates in a more interesting fashion. To determine which outcome arises, what matters is whether the participation constraints binds for some workers.

First, consider the case where the participation constraints are not binding, so that all apply. The relevant metric for how much effort is exerted is \hat{v} , the critical level of indifference for the incentive compatibility constraint. (Not surprisingly, \hat{v} is identical to the point where $\tilde{v}(\bar{e}) = \tilde{v}(\underline{e})$.) In that case, if \hat{v} is interior, the usual positive effect of incentives on performance holds, in that

$$\frac{d\hat{v}}{dw} = \frac{-2\rho^*}{(1 - \rho^*)\beta_1} < 0, \quad (23)$$

so that higher wages reduce the critical level of altruism above which high effort ensues. If \hat{v} is not interior, then contracts have no role on effort. The outcome where all apply arises for wages high enough; i.e, where type \hat{v} finds it desirable to apply. Straightforward computations show that this arises for all wages above \hat{w} , where $\hat{w} = \frac{(\epsilon - \underline{\epsilon})d}{(\bar{\epsilon} - \epsilon)(\bar{\epsilon} - \rho^*(\bar{\epsilon} - \epsilon))}$.

Next consider the case where the wage is below \hat{w} , where the participation constraint binds for some workers. Then the group that applies for the position is given by (i) all those with altruism greater than $\tilde{v}(\bar{\epsilon})$, and (ii) all those with altruism less than $\tilde{v}(\underline{\epsilon})$.²⁶ In other words, the employees of bureaucracies often have radically different preferences; it is not simply that preferences are randomly distributed with respect to the population (as would be the case when the participation constraints do not bind), but rather are drawn disproportionately from both ends of the distribution.²⁷

The Effect of Increased Incentives The other main observation of this section is that the use of more high-powered incentives can result in less effort. This arises simply because more incentives can disproportionately select those with little intrinsic motivation to apply for the position. The idea is quite intuitive: if firms increase wages, those who exert high effort benefit - by $\bar{\epsilon}\rho^*dw$, but so also do those who exert low effort - by $\underline{\epsilon}\rho^*dw$. There is no necessary reason why this will increase average efficiency.

A difficulty for comparative static exercises here is that the equilibrium of the model need not be unique for any wage. This should not be surprising, as the decision of an agent to join depends on who else is believed to join. As the objective of this section is simply to show a situation where increased incentives can backfire, I ignore this here by assuming that the density f is small; the implication of this is that changes in the fractions of different types of agents has little effect on a given agent's decision. Instead, her decision is dominated by wages and her own preferences.

To see this, note that when the participation constraints bind, expected effort is given by $\epsilon = \frac{[1 - F(\tilde{v}(\bar{\epsilon}))]\bar{\epsilon} + F(\tilde{v}(\underline{\epsilon}))\underline{\epsilon}}{1 - F(\tilde{v}(\bar{\epsilon})) + F(\tilde{v}(\underline{\epsilon}))}$. Changing the wage then affects incentives via changes in the \tilde{v} 's: does a higher wage induce more supply from high effort agents or low effort agents? Note that if the

²⁶If type \hat{v} does not want to participate, it is guaranteed that $\tilde{v}(\bar{\epsilon}) > \tilde{v}(\underline{\epsilon})$.

²⁷It is worth giving an example to illustrate the influences at work. Consider the recruitment of officials who allocate social welfare benefits. For the reasons described in Section 3, the principal is likely to desire agents who are least altruistic towards the client, as these agents exert more effort. Those who are indeed least altruistic will self-select, for the reason that it satisfies their greater desire to harm these clients (which high effort does). Yet, so also does this occupation select the most altruistic, as they see their lower effort increasing their utility, which they would like to see. It is only those who are relatively indifferent who are deterred from applying. In this sense, those who allocate social welfare benefits may have radically different preferences to each other.

equilibrium is unique,

$$\frac{d\tilde{v}(e)}{dw} = - \left(\frac{e\rho^* - \tilde{v}(e) \frac{(1-\rho^*)\beta_1}{2} \frac{d\epsilon}{dw}}{(e - \epsilon) \frac{(1-\rho^*)\beta_1}{2}} \right). \quad (24)$$

As a result, routine calculations show that

$$\frac{d\epsilon}{dw} = [\bar{e} - \underline{e}] \frac{\frac{2F(\tilde{v}(\underline{e}))f(\tilde{v}(\bar{e}))\bar{e}\rho^*}{(\bar{e}-\epsilon)(1-\rho^*)\beta_1} + \frac{2[1-F(\tilde{v}(\bar{e}))]f(\tilde{v}(\underline{e}))\underline{e}\rho^*}{(\underline{e}-\epsilon)(1-\rho^*)\beta_1}}{[1 - \frac{F(\tilde{v}(\underline{e}))f(\tilde{v}(\bar{e}))\tilde{v}(\bar{e})}{(\bar{e}-\epsilon)} - \frac{[1-F(\tilde{v}(\bar{e}))]f(\tilde{v}(\underline{e}))\tilde{v}(\underline{e})}{(\underline{e}-\epsilon)}][1 - F(\tilde{v}(\bar{e})) + F(\tilde{v}(\underline{e}))]^2}, \quad (25)$$

which has indeterminate sign. There are two sources of indeterminacy. The denominator contains a term $1 - \frac{F(\tilde{v}(\underline{e}))f(\tilde{v}(\bar{e}))\tilde{v}(\bar{e})}{(\bar{e}-\epsilon)} - \frac{[1-F(\tilde{v}(\bar{e}))]f(\tilde{v}(\underline{e}))\tilde{v}(\underline{e})}{(\underline{e}-\epsilon)}$, which itself has indeterminate sign. Only if this term is positive will the equilibrium be continuous and unique. A sufficient condition for this is that $f(\tilde{v})$ small, as I assume.

In this case, the outcome is determined by the numerator of (25), and the expected level of effort is increasing in wages only if

$$\left(\frac{f(\tilde{v}(\bar{e}))}{1 - F(\tilde{v}(\bar{e}))} \right) \frac{\bar{e}}{\bar{e} - \epsilon} > - \left(\frac{f(\tilde{v}(\underline{e}))}{F(\tilde{v}(\underline{e}))} \right) \frac{\underline{e}}{\underline{e} - \epsilon}. \quad (26)$$

In effect, higher efficiency wages induce supply from both high effort and low effort agents, and this term illustrates the conditions under which the high effort supply dominates. First, it depends on the hazard rates, reflecting the relative proportions of high-effort and low-effort agents on the margin. Second, it depends on the extent to which the agents gain from the increased efficiency wage, which is proportional to effort exerted by the high and low effort agents- this is the numerator of the second part of each side. This feature acts in favor of the high effort agents, as they benefit from the higher wage more often. Finally, it depends on the effort exerted by others in the position, because labor supply is affected by who may get the position, as this affects client welfare. To see this, note that the alternative hire exerts expected effort ϵ , so the marginal private benefit from being employed is proportional to $e - \epsilon$. What this implies is that the marginal incentive to supply high (low) effort is reduced if it is perceived that others do likewise.

Will selection be random or 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 $[1 - \rho^*]\epsilon + \rho^* - w - \underline{\kappa}$. If wages are set above \hat{w} , all apply for the position, and the marginal effect of wages on effort is positive. But offering a wage this high has two problems. First, it is costly in the usual way of efficiency wages. Second, it may be a lower wage offers higher average effort. As a result, the principal may choose a wage such that the participation constraint binds for some workers, and so there is non-random

selection.

Given this, which bureaucracies are most likely to exhibit this tendency? First, potential bureaucrats must care sufficiently about the activity. As a result, it would not be surprising to see that those who assign unemployment benefits or who provide social welfare benefits would exhibit this form of selection, but it would be a less plausible explanation for selection at the Department of Motor Vehicles, where it is hard to see the strength of preferences being strong enough. In terms of the model, β is too small. To see this, note from (20) and (22) that $\tilde{v}(\underline{e})$ and $\tilde{v}(\underline{e})$ tend to $-\infty$ as $\beta \rightarrow 0$, so long as the standard efficiency wage condition $(\bar{e} - \underline{e})\rho^*w \geq d$ holds. In words, when personal preferences are not important, the only thing that affects recruitment is wages, not altruism. Another important issue is the extent to which the agent is pivotal - $(1 - \rho^*)(e - \epsilon)$ in the model. Again, from (20) and (22), note that as $1 - \rho$ tends to zero, the only factor that affects supply is wages rather than v . Perhaps the most important aspect of this is the extent to which bureaucrats are allowed to exercise discretion in their choices. In many bureaucracies, employees do little more than rubber stamp documents which are provided by the client; again, the Department of Motor Vehicles would be a good example. Only in instances where bureaucracies must delegate decision-making to agents do these issues become important: relevant examples here would include police officers, immigration officials, and social welfare workers at the Department of Children and Family Services, who must use judgment in deciding whether to take a child from a family.

5 Conclusion

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 fact that they give out rents, and so have leeway to be unpleasant to clients without fear of “losing their business.” Yet so also do doctors, nurses, firemen, social workers, and teachers give rents to clients, yet do not seem to be as villified.²⁸ 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.

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.

²⁸A common concern among economists regarding public bureaucracies is the absence of competition. While it is true that there is little competition for the IRS or police force, it must be said that social welfare agencies, fire departments, and hospitals (in many countries) also hold a monopoly over their services, with different perceptions in the population.

The Numeraire for Altruism Throughout the paper, the preferences of the agent have been described as altruism towards the client. This is a reduced form for how bureaucrats trade off allocative mistakes. The actions that the agent takes typically affect parties other than the “client” of the model; for instance, taxpayers lose from giving benefits to an undeserving applicant. A feature of many goods allocated by bureaucrats is that people disagree over their importance. For example, some see income support programs as an important crutch for the poor, while others see them as little more than an institutionalized mechanism to avoid work. What altruism in the model refers to is how the agents value errors which act in favor of the affected client relative to the losses which accrue to those who suffer from this mistake. So, for instance, v refers to the weight that an agent places on a welfare recipient being given a dollar compared to the value that she places on the loss of the dollar from general revenues. The numeraire, then, is the cost incurred on other parties.

Complaints In the model, client complaints serve to focus particular attention on mistakes where a “client” loses. More generally, this is meant as a metaphor for any asymmetry in oversight in favor of the concentrated interests of losers. This is formalized through δ in the model, which if low, allows the voice of the client to be heard louder than other (unmodeled) parties. Again this need not be literally interpreted: instead, it represents any reason why a principal may intervene more when there is a concentrated losing client than the losses which accrue to the population at large. It could reflect the inherent preferences of the population; for instance, the judicial system is designed to allow the innocent to go free, even at the expense of allowing some guilty to also go free. Alternatively, it could represent an asymmetry in the way that bureaucrats are evaluated. For instance, a doctor who inefficiently fails to give a treatment to a patient is often more likely to be under much greater scrutiny than one who gives out treatments too freely. Or, finally it could simply illustrate the older Stiglerian view of the ability of concentrated interests to lobby, such as in many regulation settings. As such, the complaints metaphor should be seen as representing no more than any asymmetry in favor of concentrated clients.

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; such as where people have instincts about how they feel about those on welfare before applying for a position allocating benefits. In others, this seems implausible. For instance, it seems likely that police officers 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. Perhaps all police officers ultimately become hostile to suspects, and all social workers become altruistic to those in need. Then the job, rather than demand and supply conditions emphasized here, determine the observed preferences of bureaucrats. There seem to be two natural responses to this criticism. First, 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 remain, and who will ultimately leave. Second, this paper can offer a view on whether the observed attributes of bureaucrats, *wherever they come from*, are efficient; 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? As a result, if indeed the act of being a doctor leads to altruism (or hostility) towards patients, is this something that is desirable?

Involvement In the model, it is assumed that the potential bureaucrat identifies the change in the client's welfare if she is employed when deciding to apply for the position. However, in some instances, what may matter to the agent is her *involvement* in the decision. To take an extreme example, 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. The assumptions of the model 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. This is not so if it is the agent's involvement in the outcome that affects their utility, as addressed in Besley and Ghatak, 2004. So a very altruistic person may not become a soldier, as his actions harm the client in some absolute sense, even though the alternatives for the client would be much harsher. 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.

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. Specifically, they prefer those who exhibit the right kind of altruism over those who carry out the job simply for money. The implication of this is the tension between selection and incentives above. Second, consider the relationship to the literature on incentive provision. Most organizations use compensation plans for the joint purposes of increasing effort and selecting workers with appropriate characteristics;

not only does incentive pay increase effort, but it also helps to select the most able or motivated. Much of the recent empirical work on incentives (at least for workers in simple jobs where output can be measured) has emphasized these joint themes. (See Prendergast, 1999, for details.) Perhaps the most important finding of this paper is that neither of these avenues is as straightforward in resolving agency problems for bureaucrats.²⁹ As such, this paper offers another reason why the efficiency of bureaucracies is likely to be limited.

²⁹In this sense, it follows Acemoglu and Verdier, 2000, and Prendergast, 2003, in its focus on the particular difficulties of attaining bureaucratic efficiency.

APPENDIX: PROOFS OF RESULTS

The Formal Program: The formal program is described here for the case where $B(1,0) - B(0,0) > 0$. The logic is identical for the case where the incentives of the client and principal align and is not included for reasons of brevity.

The Incentive to Investigate Begin by considering the incentive to investigate. In any pure strategy equilibrium, for any allocation proposed by the agent, complaints either reveal that the information of the client differs (the informative case), or reveal no information (the uninformative case). Let $z(a, m)$ be the probability that the allocation made by the agent a is incorrect upon receiving message m . Then maximizing (7) implies that the principal sets

$$\rho^*(a, m) = \frac{z(a, m)}{\gamma}. \quad (27)$$

If complaints are informative $\rho^*(a, c) = \frac{s(1-e)}{\gamma[e(1-s)+s(1-e)]}$, $\rho^*(a, n) = \frac{(1-s)(1-e)}{\gamma[es+(1-s)(1-e)]}$, while if complaints are uninformative, $\rho^*(a, \cdot) = \frac{1-e}{\gamma}$.

The Incentive to Complain Next, let $\mu(a)$ be the probability that the agent complains when the state allocated is a . If complaints are informative, this is given by $(1-e)s$, while if uninformative, $\mu(a)$ equals 0 or 1. Note that a complaint increases the probability of the truly efficient state being implemented, so that the agent complains if his signal disagrees when the recommendation is a if

$$\frac{s(1-e)}{e(1-s)+s(1-e)}[\rho^*(a, c) - \rho^*(a, n)][B(j, j) - B(a, j)] \geq \delta. \quad (28)$$

Similarly, the agent does not complain if his signal agrees only if

$$\delta \geq \frac{(1-s)(1-e)}{es+(1-s)(1-e)}[\rho^*(a, c) - \rho^*(a, n)][B(j, j) - B(a, j)]. \quad (29)$$

If $B(1,0) - B(0,0) > 0$, then as $\delta > 0$, when $a = 1$, complaints are never informative, as complaining increases the probability of the correct state being introduced. Then if (12) holds, complaints are informative when $a = 0$, and otherwise uninformative when $B(1,0) - B(0,0) > 0$.

The Incentive to Exert Effort First consider the case where complaints are informative. The agent exerts effort for two possible reasons: (i) wages, and (ii) private benefits. First consider the agent's wages. She is given w^* if she is investigated and found to have been correct. This implies that expected wages from an allocation $a = 0$ is

$$W^* = e(s\rho^*(0, c) + (1-s)\rho^*(0, n))w^*, \quad (30)$$

and from an allocation of $a = 1$ expected wages are

$$W^{**} = e\rho^*(1, n)w^*. \quad (31)$$

Total expected wages are then given by $\frac{W^*+W^{**}}{2}$. Expected private benefits, EB^* , are given by

$$\frac{EB^*}{v} = [e + (1-e)(s\rho^*(0,c) + (1-s)\rho^*(0,n))] \frac{B(1,1)}{2} + [e + (1-e)\rho^*(1,n)] \frac{B(0,0)}{2} + [(1-e)(1 - (s\rho^*(0,c) + (1-s)\rho^*(0,n)))] \frac{B(0,1)}{2} + (1-e)(1 - \rho^*(1,n)) \frac{B(1,0)}{2}, \quad (32)$$

and total incentives are the sum of the two, where effort is exerted to maximize $\frac{W^*+W^{**}}{2} + EB^*$.

Next consider the case where complaints are uninformative, so that they are always ignored if offered. As a result, $\mu = 0$ as $\rho(a,c) \geq \rho(a,n)$ and $\delta \geq 0$. Here the agent's expected monetary rewards are W^{**} and expected private benefits are given by

$$\frac{EB^{**}}{v} = [e + (1-e)\rho^*(1,n)] \frac{B(1,1) + B(0,0)}{2} + (1-e)(1 - \rho^*(1,n)) \frac{B(0,1) + B(1,0)}{2}. \quad (33)$$

Here the agent exerts effort to maximize $W^{**} + EB^{**}$.

The choice of agent and contract Finally, consider the principal's choice. First, consider the case where complaints are informative. His objective is to maximize expected surplus minus investigation and wage costs, which is given by

$$e + (1-e)\rho^*(1,n) + (1-e)s\rho^*(0,c) + (1-e)(1-s)\rho^*(0,n) - \frac{\gamma\rho^*(1,n)^2}{2} - s\frac{\gamma\rho^*(0,c)^2}{2} - (1-s)\frac{\gamma\rho^*(0,n)^2}{2} - \frac{W^* + W^{**}}{2}. \quad (34)$$

When complaints are uninformative, the principal's return is given by

$$e + (1-e)\rho^*(1,n) - \frac{\gamma\rho^*(1,n)^2}{2} - W^{**}. \quad (35)$$

The program of the principal then as follows.

1. If (12) holds, the principal maximizes (34) subject to the agent exerting effort to maximize $\frac{W^*+W^{**}}{2} + EB^*$, and the principal overseeing with probabilities $\rho^*(0,c) = \frac{s(1-e)}{\gamma[e(1-s)+s(1-e)]}$, $\rho^*(0,n) = \frac{(1-s)(1-e)}{\gamma[es+(1-s)(1-e)]}$, and $\rho^*(1,n) = \rho^*(1,c) = \frac{1-e}{\gamma}$.
2. If (12) is violated, the principal maximizes (35) subject to the agent exerting effort to maximize $W^{**} + EB^{**}$, and the principal overseeing with probabilities $\rho^*(a,c) = \rho^*(a,n) = \frac{1-e}{\gamma}$.

Note that in both instances the only effect of v on the principal's return is through the effort decisions that are made. Since $w^* > 0$, it immediately follows that the principal chooses the agent type that (weakly) exerts the most effort, as carried out in the text of the paper.

Proof of Lemma 1:

Two issues arise: (i) how costly is it to induce the agent to participate- individual rationality?, and (ii) how costly is to induce effort exertion- incentive compatibility? Each are considered in turn.

Individual Rationality In this model, agents earn non-negative monetary rents. Hence, if private benefits are greater than zero, there is always supply of labor from the principal's most preferred type. But the preferred agent receives zero marginal private benefits from being employed in equilibrium. This is because the principal always induces effort exertion from the employee. As a result, it makes no difference (in terms of private benefits) whether she works there, or someone else does.

Monetary Incentives If employed, the agent is paid w . If fired, the agent is replaced. Consequently, the principal chooses the lowest wage, w , such that the agent exerts effort. Consider when the agent should be retained. The agent can propose $a = 1$ or $a = 0$. If $a = 1$ is proposed, the agent never complains as $\delta > 0$. It should be immediately apparent that if investigated and wrong, the agent is fired. Similarly, the principal cannot improve on a contract where the agent is retained if not investigated; the wage will be paid anyway, and this state has no effect on the agent's incentives. It should also be clear that the agent should be retained if investigated and correct, as this increases incentives to exert effort.

Next, the agent can propose $a = 0$. Here the agent may complain, and so there are (possibly) 6 states: for each message sent by the client, there is a state (i) when not investigated, (ii) when investigated and correct, and (iii) when investigated and incorrect. Again, it should be clear that if investigated and incorrect, the agent should be fired. If not investigated and there is a complaint, then the agent should also be fired, as probabilistically this reveals an incorrect decision by the agent. Then it should be clear that if investigated and correct, the agent should be retained.

What remains is whether the agent should be retained when not investigated but there is no complaint. If complaints are uninformative, then the principal is indifferent about whether to retain the worker. If complaints are informative, then probabilistically the absence of a complaint reveals information that the agent is correct, and so the principal can offer at least as low a wage by retaining the worker in this state than firing her with any positive probability.

REFERENCES

- Acemoglu, Daron, and Thierry Verdier**, “The Choice between Market Failures and Corruption,” *American Economic Review*, 2000, 90(1), 194-211.
- Aghion, Philippe, and Jean Tirole**, “Formal and Real Authority in Organizations,” *Journal of Political Economy*, 1997, 105, 1-29.
- Akerlof, George, and Rachel Kranton**, “Identity and the Economics of Organizations,” 2003, mimeo, UC Berkeley.
- Benabou, Roland, and Jean Tirole**, “Intrinsic and Extrinsic Motivation,” *Review of Economic Studies*, 2003, 70(3), 488-520.
- Besley, Tim, and Maitreesh Ghatak**, “Competition and Incentives With Motivated Agents,” 2003, mimeo, LSE.
- Blakertz, Laura, and Susan Robinson**, “Turnover Intentions of Community Mental Health Workers in Psychosocial Rehabilitation Services,” *Community Mental Health Journal*, 1997, 33(6), 517-29.
- Brehm, John, and Scott Gates**, *Working, Shirking, and Sabotage: Bureaucratic Responses to a Democratic Public*, 1997, University of Michigan Press, Ann Arbor.
- Delfgaauw, Josse, and Robert Dur**, “Signaling and Screening of Workers’ Motivation,” 2003, mimeo, Erasmus University.
- Derthick, Martha**, *Policymaking for Social Security*, 1979, Brookings Institute, Washington.
- Dewatripont, Matthias, Ian Jewitt, and Jean Tirole**, “Missions and the Accountability of Government Agencies,” *Review of Economic Studies*, 1999, vol. 66, issue 1, pages 199-217.
- Dewatripont, Matthias, and Jean Tirole**, “Advocates,” *Journal of Political Economy*, 1999, vol. 107, issue 1, pages 1-39.
- Dixit, Avinash**, “Incentives and Organizations in the Public Sector: An Interpretive Review,” *Journal of Human Resources*, 2002, 37(4): 696-727.
- Ellett, Alberta**, “The Role of Human Caring in Employee Retention in Child Welfare: Practice Implications from a Two-State Quantitative Study,” mimeo, 2001, University of Georgia.
- Frey, Bruno, and Reto Jegen**, “Motivation Crowding Theory: A Survey of Empirical Evidence,” forthcoming, *Journal of Economic Surveys*.
- Goldsmith, Andrew**, *Complaints Against the Police*, 1991, Clarendon Press, Oxford.
- Goodsell, Charles**, *The Public Encounter*, 1981, Indiana University Press, Bloomington.
- Johnson, Ronald, and Gary Libecap**, “Bureaucratic Rules, Supervisory Behavior, and the Effect of Salaries in the Federal Government,” *Journal of Law, Economics, and Organizations*, 1989, 5, 53-83.
- Kadushin, Alfred**, *Supervision in Social Work*, 1985, Columbia University Press, New York.

- Lee, Raymond, and Blake Ashforth**, “A Further Examination of Managerial Burnout: Towards an Integrated Model,” *Journal of Organizational Behavior*, 1993, 14(1), 3-20.
- Lipsky, Michael**, *Street Level Bureaucrats*, 1980, Russell Sage Foundation, New York.
- Manning, Peter, and John Van Maanen**, *Policing: A View from the Street*, 1978, Goodyear Publishing, Santa Monica.
- Mor Barak, Michal, Jan Nissly, and Amy Levin**, “Antecedents to Retention and Turnover among Child Welfare, Social Work and Other Human Service Employees,” *Social Service Review*, 2001, 625-661.
- Peabody, Robert**, *Organizational Authority*, 1964, Atherton Press, New York.
- Prendergast, Canice**, “The Provision of Incentives in Firms,” *Journal of Economic Literature*, 1999, 37, 7-63.
- Prendergast, Canice**, “The Limits of Bureaucratic Efficiency,” *Journal of Political Economy*, October 2003, 929-959.
- Skolnick, Jerome**, *Justice Without Trial*, 1966, John Wiley, New York.
- Skolnick, Jerome**, *The Police and the Urban Ghetto*, 1968, American Bar Foundation, Chicago.
- Warner, W. Lloyd, Paul Van Riper, Norman Martin, and Orvis Collins**, *The American Federal Executive*, 1963, Yale University Press, New Haven.
- Weatherly, R., C. Kottwitz, D. Lishner, K. Reid, G. Roset, and K. Wong** “Accountability of Social Service Workers at the Front Line,” *Social Service Review*, 1980, 33, 556-571.
- Wilson, James**, *Bureaucracy*, Basic Books, New York, 1989.