NBER WORKING PAPER SERIES

SELECTION AND OVERSIGHT IN THE PUBLIC SECTOR, WITH THE LOS ANGELES POLICE DEPARTMENT AS AN EXAMPLE

Canice Prendergast

Working Paper 8664
http://www.nber.org/papers/w8664

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2001

I would like to thank the Los Angeles Police Department for kindly sharing their data, Raghu Rajan and Tano Santos for helpful comments, Lan Shi for excellent research assistance, and the NSF and the University of Chicago for generous support. Any errors are my own. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

Selection and Oversight in the Public Sector,
With The Los Angeles Police Department as an Example

Canice Prendergast

NBER Working Paper No. 8664
December 2001
JEL No. H8, L3, R5

ABSTRACT

I offer theoretical and empirical observations on the oversight of public sector employees. I argue that it is unreasonable to expect that the solutions typically considered in the literature will be effective with public sector employees, because bureaucrats are especially difficult to monitor. To offset this weakness, agencies tend to hire bureaucrats who are biased against consumers, where such bias increases incentives. I then address how bureaucrats should be overseen and offer a choice between internal monitoring of public agencies, with overseers who are biased against consumers, or external monitoring, where bureaucrats become excessively worried about the prospect of an investigation and may change their behavior to attain that goal. I provide evidence from the Los Angeles Police Department to show that officers appear to have responded to increased oversight by reducing crime-fighting activities in an attempt to avoid investigation.

Canice Prendergast
Graduate School of Business
The University of Chicago
1101 E. 58th Street
Chicago, IL 60637
and NBER
canice.prendergast@gsb.uchicago.edu

1 Introduction

Concerns about the behavior of public sector employees are widespread. Accusations of malfeasance include the use of excessive force by police officers, immigration officials making decisions based on nationality rather than merit[1], the indifference of the Internal Revenue Service employees towards their clients, and so on. Yet despite these concerns, little work has been done in economics on how public sector agencies should be monitored.[2] An important topical issue concerns the role of external parties in overseeing public agencies. Much oversight in the public sector is internal, where for example the performance of a police officer is evaluated by a superior. Yet external constituencies, such as the media and congress, have become increasingly involved in evaluating performance. Two well known examples are (i) the 1997 congressional hearings on the way that IRS officials treat tax payers, that (among other things) resulted in the creation of the Internal Revenue Service Oversight Board, and (ii) the media attention on the Los Angeles Police Department that has ultimately resulted in almost 10% of all officers being discipined in the last three years. This paper addresses the role of internal and external oversight for the kind of activities typically carried out in the public sector.

One feature that seems to characterize government agencies is a lack of trust in their willingness to credibly investigate the actions of their employees.[3] I argue that this occurs partly because their employees have different preferences to the population whose objectives they implement. These differences in preferences make them more effective when dealing with clients but tend to make them poor overseers. For example, I show that there is a value to choosing police officers who see the claims of innocence by suspects as less important than does the average member of the population. While this helps to catch criminals, genuine claims of innocence tend to fall on deaf ears.

External monitoring, if based on public opinion, is less likely to suffer from such bias. Yet outsiders tend to be poorly informed. As I have argued in Prendergast, 2001, their attention is

[1] See New York Times, 2000, and Hall, 2001, for examples.

[2] It is well known that the kinds of pay-for-performance contracts used in the private sector are unlikely to be effective in the public sector, mainly for multi-tasking reasons (Dixit, 2001). But this still leaves open the issue of how oversight *should* be done.

[3] An obvious example is the perceived reluctance of police departments to investigate wrongdoing by their officers (for example, www.policeabuse.org), heightened by a reluctance of officers to report on each other, known as the "Thin Blue Line". Yet the police are hardly alone in this: one of the primary reasons for the proliferation of malpractice suits in the medical profession is the historical resistance of hospitals to effectively monitor doctors. Similarly, the INS and IRS are hardly known for taking the side of consumers in interactions with their officials. This is not to say that private sector firms can be always trusted to rely on internal monitoring. For instance, the media and courts have played an important role in the many cases of auto defects and drug problems, such as with Thalidomide in the 1950s.

often piqued only when there is evidence that a mistake has happened. This typically comes from a disgruntled consumer; a suspect incorrectly arrested, a patient wrongly turned down for a medical procedure, a legitimate immigrant turned down at the border. But if consumers play this role of pointing out mistakes, officials may simply accede to the demands of consumers, even when not warranted, to avoid complaints.[4] As a result, the paper offers a choice between internal monitoring of public agencies, with potentially biased overseers, or external monitoring, where bureaucrats may inefficiently allocate benefits to avoid investigation. I also provide empirical evidence on the Los Angeles Police Department which suggests that increased external oversight has caused officers to inefficiently avoid situations that could lead them into trouble.

It is worth emphasizing that I do not address whether an activity should be carried out in the private or public sector: this is not a theory of ownership. Instead, my objective is to understand oversight when (i) consumer complaints are used to monitor employees, and (ii) consumers benefit from some inefficient allocations. I argue below that many public sector activities satisfy these criteria. I make no claim that the problems below arise *because* they are done in the public sector, nor that there are not some private sector examples. Instead, these problems arise in the public sector only because they often satisfy these two criteria.

I begin in Section 2 by arguing that oversight for many public sector activities is especially difficult because signals on performance are typically unreliable. A key issue in solving many agency problems is whether consumers, the primary source of information on performance, are willing to point out employee incompetence. Because many goods allocated by the public sector are benefits (citizenship, unemployment insurance, etc.), their consumers can have very different incentives to overseers and are unwilling to point out known errors. For example, guilty suspects have no reason to point out a police officer's error by not arresting him, nor do unqualified immigrants own up when incorrectly allowed into the country.

I show that the ensuing weakness of contracts to solve public sector agency problems has important implications for the selection of bureaucrats. A feature of many goods allocated by the public sector 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. Equally, there is disagreement over how the police should be allowed to investigate a crime, such as the conditions under which searches should be allowed. Some people are more concerned with catching criminals, whereas others worry more

[4]For instance, a police officer could avoid confrontations and reduce arrests, the IRS could cut down on auditing difficult cases which could land the official in trouble, and the INS could allow unqualified applicants to enter the country.

about the privacy rights of citizens. I consider where bureaucrats should lie on this dimension: should they have the preferences of the average member of the population, or should society hire biased bureaucrats?[5]

I provide a reason why bureaucrats should have biased preferences, where they see the claims of consumers as less important than does society. The bias improves their incentives to make good decisions. As with other members of society, bureaucrats care about appropriately allocating treatments to clients; doctors care about curing patients, police officers care about controlling crime, immigration official are concerned with denying access to the unqualified, and so on.[6] Consider then the decision to hire one of two police officers: one who is largely concerned with avoiding the innocent being convicted, and another who worries most about the loss from letting the guility go free. There is one central difference between these two cases. Specifically, there are many mechanisms in place to correct outcomes where the innocent are arrested, such as complaints, trials, appeals, and so on. These mechanisms reduces the ultimate likelihood of the innocent being convicted. But the knowledge that their mistakes will be corrected reduces incentives, as an officer is partly motivated by the desire to avoid mistakes. By contrast, an officer who mistakenly releases the guility can expect little likelihood of this being brought to light, as the suspect has nothing to gain from pointing out the error. Then an officer who is particularly concerned with the guility being set free becomes more useful, as she has greater incentives to find appropriate evidence on the client. Thus, more generally bureaucrats will be disproportionately selected from those who are sceptical of the demands of consumers.

Given this tendency to select biased bureaucrats, I then consider their oversight; internal monitoring is addressed in Section 3. Insiders typically know more than do those who are more removed from the proceedings, especially about routine cases. This allow overseers to reward employees for taking hard decisions, rather than only responding when a mistake occurrs. But internal monitoring decisions largely reflect the preferences of insiders: for instance, they rely on the willingness of co-workers to report errors. Bureaucratic bias then becomes a problem, as oversight is not targeted in the way that society would like: because bureaucrats are, on average, biased against consumers, there is too little responsiveness to the legitimate concerns of denied consumers, and too much intervention when consumers are allocated benefits. In effect, the characteristics that make for good bureaucrats also make for poor overseers.

[5]Selection concerns are well known in economics, and the role of contracts in (for example) attracting more able workers has become part of the lexicon of agency theory. I do not consider selection in this usual sense of getting talented workers: instead I am more interested in the role of biased evaluators.

[6]This "professionalism" of bureaucrats repeatedly arises in discussions of public sector agencies, such as Wilson, 1989.

This leads to a role for external monitoring, considered in Section 4, which has the advantage that external monitors do not exhibit the bias that characterizes internal monitoring. But as in Prendergast, 2001, external monitors are typically less informed than are their internal counterparts, and their attention tends to be focused the actions of spurned consumers. It is usually the case where the drug-dealing suspect was incorrectly arrested or beaten up that brings attention, not the case where the drug-dealer was incorrectly set free.[7] Investigations by external evaluators tend to be focused not on all cases where the bureaucrat erred, but where she erred *and* harmed a consumer. Consequently, bureaucrats are often faced with the temptation to largely avoid oversight by capitulating to consumers. The problem with external monitoring is that it becomes difficult to induce bureaucrats to deny benefits to consumers. The paper therefore offers a cautionary tale for increasing external oversight (such as the "zero-tolerance" programs of some police departments) as the incentive to simply give consumers what they want becomes too tempting.

Section 5 provides empirical evidence on the particular difficulty of monitoring police officers, by providing time series data on the Los Angeles Police Department. I provide evidence that is consistent with officers responding to increased oversight by being less likely to harm consumers, in this case suspects, as doing so is now more likely to result in disciplinary measures. Specifically, I show that after an increase in oversight in 1998, officers have been more reluctant to engage in "aggressive" policing on a variey of dimensions such as use-of-force, officer-involved shootings, assaults *on* officers, and arrest rates. This change in behavior is also correlated with a large recent increase in gang crime.

Section 6 considers some extensions. Foremost among these is to consider why these problems only arise for activities where consumers are known to benefit from particular allocations, even when they do not deserve them. To take an example, consider a patient who may qualify for a costly procedure that can improve his health. If the patient does not pay for the procedure, the patient always wants the procedure, and the problems described above arise. If, by contrast, the patient is required to pay for the procedure, I show that the payment of prices can make consumer feedback more credible, which eliminates the need to bias the selection of agents. As a result, it is the nature of the goods allocated, namely, that they are *benefits*, that induces the selection and agency issues that are the focus of the paper.

Before beginning, it is worth emphasizing that this paper is not meant to justify oversight difficulties for all public sector employees. As discussed in Dixit, 2001, there many reasons why oversight is difficult in the public sector. Instead, the contribution of this paper is restricted to those

[7]While the police department itself may be criticized for its inability to control crime, this rarely reaches the individual officer.

cases where (i) a bureaucrat allocates benefits to a consumer, and (ii) feedback from the consumer reveals information on her performance. This would include many officials of institutions like the IRS, INS, and Department of Motor Vehicles, in addition to police officers, but this paper has little to say about the canonical "faceless" bureaucrat whose performance is never directly observed by consumers.

2 The Basic Model

An allocation A must be made to a consumer, where A can take on a value of 0 or 1. An allocation of 1 is beneficial to the consumer: for concrete purposes, think of $a = 0$ as arresting a suspect and $a = 1$ as freeing him. Society, which consists of M (interested) individuals, values the allocation based on a characteristic of the consumer. *Average* per capita social surplus from the allocation depends on a parameter α and is given by

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

Thus social surplus is positive only if A is properly matched to the underlying environment, α: so for example a suspect should only be arrested if there is a reasonable belief that the person is guilty. The true value of α is unknown and can take two values $\alpha = 1$, or $\alpha = 0$. I assume that each state occurs with equal probability.

Information There are three actors in this model: society (the principal), an agent (or bureaucrat), and a consumer. The agent is used to propose an allocation to the consumer. To do so, she collects information on α; she observes α_a which is correct with probability $q \geq \frac{1}{2}$, where with probability $1 - q$, she observes $\alpha_a \neq \alpha$. The precision of the agent's estimate depends on an unobserved effort decision by her, where she chooses effort $e \geq 0$ at a disutility of e. Let $q(e)$ be the precision of the agent's estimate, where $q(0) \geq \frac{1}{2}$, $q'(e) > 0$, $q''(e) \leq 0$, and $q(\infty) < 1$.

Second, the consumer observes α. Let $V(A, \alpha)$ be the utility obtained by the consumer if his type is truly α and the allocation is A. I assume that the consumer has different preferences to society as

$$V(1, \alpha) > V(0, \alpha). \tag{2}$$

Thus, the consumer wishes to receive the benefit ($a = 1$) independent of his type, so his preferences are not aligned with those of society.

Somebody is charged with overseeing the agent: this choice is made by the principal. The overseer is uninformed unless he carries out an investigation. To model a role for investigations, I assume that the overseer chooses a probability of observing α at some cost.[8] Specifically, the overseer chooses a probability ρ with which he observes α, at a *per capita* cost $\kappa(\rho)$, where $\kappa'(\rho) > 0, \kappa''(\rho) > 0, \kappa'(0) = 0, \kappa(0) = 0$, and $\kappa'(1) \geq 1$. For simplicity, throughout the paper I assume that monitoring costs are quadratic: $\kappa(\rho) = \frac{\gamma\rho^2}{2}$, where $\gamma \geq 1$. At this stage, the identity of the overseer remains vague, and depends on whether monitoring is internal or external, as described below.

Heterogeneity Average social surplus in (1) is made up of the different preferences of individuals in society. The idiosynchratic preference of person i is given by v_i, and measures the extent to which he values $A = 1$ over $A = 0$. The value that person i places on the allocation is

$$S_i(A; \alpha) = \frac{S(A; \alpha)}{M} + v_i A. \tag{3}$$

As mentioned in the introduction, my interest is in how bureaucrats weigh one error (where $A \neq \alpha$) relative to another: for example, how bothered are they by a qualified applicant being denied compared to an unqualified applicant being accepted? This is measured by v_i. To understand the import of these preferences, consider the cost of misallocating A to α. For person i, the cost of misallocating the benefit to the agent ($A = 1$) is $1 - v_i$, while the cost of incorrectly failing to allocate it ($A = 0$) is $1 + v_i$. Thus, v_i measures the extent to which a person values one error over another. Thus, in the law enforcement example, lower v_i types are less worried about the innocent being convicted than are higher v_i individuals individuals. For simplicity, I assume that v_i can take two values: $\overline{v}$ with probability z and $\underline{v}$ with probability $1 - z$, where $z\overline{v} + (1 - z)\underline{v} = 0$. Reservation utilities are independent of v.

Actions and Contracts The bureaucrat has two actions: (i) an effort level, $e(v)$, and (ii) an allocation to give the consumer, $a(v)$. Based on this allocation the consumer sends a message $m \in \{n, c\}$, where the message n means that no complaint is made and message c implies that a complaint has been made.

The principal (society) has three choices. First, he employs N employees, where he chooses a

[8] I assume that the signal received by him is correct with probability 1. This assumption is used to rule out "nuisance complaints", where a customer complains even when he knows that the bureaucrat made the right decision in the hope that the principal will come to the wrong conclusion and overturn the bureaucrat's decision in the consumer's interest. I am largely interested here in cases where complaints are informative of bureaucratic error and so ignore this possibility by assuming that the principal never makes errors after an investigation.

fraction μ of $\overline{v}$ types and $1-\mu$ of $\underline{v}$ types. The choice of μ is unobserved by the agents.[9] N is assumed small relative to the population. Second, he chooses oversight of the bureaucrat, where oversight is either internal or external. If monitoring is external, the principal himself oversees. If monitoring is internal, the employees choose monitoring levels to maximize their expected utility.[10] The party charged with overseeing the agent chooses to investigate the performance of the agent with probability $\rho(a,m)$, where monitoring potentially depends both on whether the customer complains and on the allocation made by the bureaucrat. The overseer cannot commit to oversight probabilities.

I assume that insiders have a comparative advantage in observing the details of cases over external parties. Specifically,

- Internal parties can observe both the allocation proposed by the bureaucrat a and the response by the consumers m.
- External parties only observe whether a consumer complains. As a result, if the consumer does not complain they do not observe the allocation proposed by the agent.

This assumption is meant to reflect the fact that outsider typically do not observe much about routine cases, but instead have their attention piqued by consumers. For example, the media typically knows little about the details of efficient drug arrests. I discuss this assumption in Section 6.

If an investigation turns up evidence that the agent made a mistake, the overseer overturns the agent's decision and allocates the correct one. If the investigation concurs with the agent's findings, or there is no investigation, the decision remains that suggested by the agent.

Preferences The objective of society is to maximize per capita surplus: $\frac{S}{M}$ from the allocation, minus investigation costs $\kappa(\rho(.,.))$ minus the effort costs of the agent ($\frac{e}{M}$). No contracts can be offered to the consumer, whose objective is to maximize V minus a cost of complaints.

The preferences of the agent are a little more complex. Her preferences differ from the population's in two ways. First, as is standard, she is effort averse. Second, she has idiosynchratic preferences given by v_i. I further assume that there are monetary consequences to her actions, where she is rewarded for making correct decisions and penalized for poor decisions. Incentives

[9] It seems unlikely that agents would observe the preferences of all members of the bureaucracy. I discuss the effect of relaxing this assumption in Appendix 3.

[10] The results are identical if one employee is randomly selected to oversee.

are provided by career concerns.[11] Assume that bureaucrats differ in their ability, where ability is unknown to all parties. Bureaucrats are of two types; those who are always right, and those who are always wrong.[12] The fraction of agents who are always right is given by $q(e)$.[13] The bureaucrat's wage depends on perceptions of her ability held by the overseeing agency, and is given by $E[q(e(v))|\Omega]\Delta$, where Ω is the information available to the overseer. The bureaucrat's payoff in any state is

$$\frac{S}{M} - \kappa(\rho(.,.)) + v_i A - e + E[q(e)|\Omega]\Delta, \tag{4}$$

The first two terms are simply the average valuation of the population for the allocation and investigation costs. The v_i term reflects the personal preferences of the agents, and e is the (fully internalized) cost of the agent.

Δ is exogenous, and reflects the marginal return to ability. In keeping with much of the literature on bureaucrats, such as Wilson, 1989, I assume that such incentives are weak, in that effort is below its first best optimal level. As a result, improvements in monitoring which increase incentives to exert effort will also increase surplus.[14] This will be the case when M is large enough (see footnote 20).

The timing of the game is as follows. First, society hires agents and decides who should oversee them. Second, nature assigns α to the consumer and the agent exerts effort. Third, the customer and the agent privately observe their signals. Next, the agent proposes an allocation a. Following this, the customer send a message m, i.e., he complains or not. The overseer observes the relevant information and then monitors with probability $\rho(a, m)$ to maximize his utility at that point. If he observes α, he allocates $A = \alpha$. Otherwise, the agent's recommendation is implemented. Finally, the agent is paid.

A Role For Complaints I am interested in cases where complaints are informative of bureaucrat error and, hence, increase the likelihood of oversight. To do this, I assume that the consumer has

[11] It is generally accepted that pay-for-performance contracts play a very limited role with public sector bureaucrats (see Dixit, 2000, Dewatripont et al., 1999, and Prendergast, 1999, for relevant examples), and that instead incentives for public sector employees are provided by career advancement induced by impressions of their ability. Consequently, incentives in this model are provided through career concerns.

[12] The model easily generalizes to less extreme types.

[13] Thus, the return to effort is to increase the fraction of agents who make correct decisions. This assumption is a little different from the standard career concerns model, such as Holmstrom, 1982, where ability is independent of effort. For example, it would be possible to specify the probability of being correct as, say, $q(ability, e)$ where the overseer infers ability from whether the agent was correct or not. In the usual way, agents have incentives to exert effort even though in equilibrium, such efforts are futile. It would be straightforward to use this form of career concerns without changing results: it is purely for simplicity that I use this form.

[14] This would typically be the case whenever there is a marginal cost to using wages as the mechanism of providing incentives, where that marginal cost of incentive provision is traded off against its benefits. This assumption is made merely to illustrate the increased surplus that is likely to arise from improving private benefits to exert effort.

an arbitrarily small cost of complaining: as the principal observes α upon investigation, consumers never complain when the bureaucrat makes the correct decision. The only issue is whether they complain when an error has been made. Although the model is not (quite) cheap talk, there is a multiplicy of equilibria based on how language (the message) is interpreted. To eliminate this, I assume that messages are meant literally, in the sense that a complaint strictly increases the probability of investigation.

Limited Private Benefits I show below that there is a role for hiring biased agents, as this improves their incentives for exerting effort. But if agents are very biased - v too large - they will make inefficient allocations simply because the private benefits from doing so are so large. For example, if a police officer values arrests so much that he arrests people he believes to be innocent, the advantages described below disappear. As a result, for the bias to play an efficient role, it needs to be of limited size that private benefits do not exceed the value of their beliefs. A sufficient condition for this is given by (5) which implies that v is low relative to the value of their information. This assumption is discussed in detail in Section 6.

$$|v| \leq \frac{(1-q(0))(1-\rho(1,n)) - q(0)(1-\rho(0,c))}{q(0)(1-\rho(0,c)) + (1-q(0))(1-\rho(1,n))} \tag{5}$$

Finally, the model is only interesting if the agents affect the allocation. This arises when agents recommend α_a. But for poorly informed agents, a better allocation is simply to always deny the consumer, and to allow only the consumer's complaints to affect the allocation. For low enough q, this is a better allocation mechanism than giving the bureaucrat the power to allocate the good. Throughout the paper, I assume that $q(0)$ is high enough that the agent is given the power to allocate the good.[15]

I consider pure strategy Bayesian Nash equilibria of the model where (i) the agent allocates $\alpha_a = a$ (if such an equilibrium exists), (ii) $\rho(.,c) > \rho(.,n)$[16], (iii) (5) holds, and one technical assumption, (iv) that the agent does not update beliefs on μ from the decision of the principal to hire her.[17] To find such equilibria, there are four important incentives to consider.

[15]The surplus from denying the consumer the good unless he complains is given by $\frac{1}{2} + \frac{\rho(0,c)-\gamma\rho(0,c)^2}{2}$. For $q(0)$ large enough, this is always dominated by allowing the agent to report truthfully. By contrast, consider the case where $q = \frac{1}{2}$ for all e. Then straightforward calculations show that the surplus from always denying the consumer is $\frac{3}{4}$, while allowing the agent to make an allocation of $a = \alpha_a$ yields surplus of $\frac{5}{8}$. As a result, I need to assume that there is a substantive role for the agent.

[16]This rules out the usual babbling equilibrium.

[17]I assume that the agent does not update on μ from the decision of the principal to hire her. Instead, the agent infers μ only by using Bayes Law from the equilibrium strategies below. This assumption rules out corner solutions where the principal in equilibrium hires none of a particular type of worker, which is enforced by the out-of-equilibrium

2.1 Incentives

The Incentive to Complain If a complaint increases the probability of an investigation, the consumer complains only if an error has been made *and* he has been denied the good $a = 0$. Consider the outcome when a bureaucrat recommends allocation j but the consumer knows that allocation i is socially optimal. If the agent complains, this increases the likelihood of i being ultimately implemented. The agent will complain if the bureaucrat errs by allocating j if $V(i,i) \geq V(j,i)$, where the small cost of complaint is ignored. To see this, note that a complaint is only relevant if the principal monitors and corrects the decision to $A = i$. This occurs with probability $\rho(a,c)$ and changes utility by $V(i,i) - V(i,j)$. If no complaint is made, the monitoring probability is $\rho(a,n)$. Thus the consumer can credibly reveal information on bureaucratic error only if $(\rho(a,c) - \rho(a,n))V(i,i) > (\rho(a,c) - \rho(a,n))V(i,j)$ for all i and j. But from (2), this only holds for $a = 0$ if $\rho(a,c) > \rho(a,n)$.[18]

Truth-Telling The agent proposes a based on her information, α_a. Truth-telling requires that the agent has an incentive to allocate $a = \alpha_a$.

Effort The agent cares about both about the allocation and the monetary consequences of her actions. She chooses effort to maximize (4).

Oversight Finally, consider the incentive to monitor. This depends on who oversees the firm.

1. If oversight is external, society chooses monitoring probabilities to maximize ex post surplus. Let $\tau(a,m;v)$ be the probability (as perceived by the overseer) that the agent of type v is incorrect based on an allocation a and a message m sent by the consumer. Then the external monitor sets $\rho(a,m;v)$ to maximize

$$\tau(a,m;v)\rho(a,m;v) - \frac{\gamma\rho^2(a,m;v)}{2} \tag{6}$$

   so that the optimal investigation probability is $\rho^*(a,m;v) = \frac{\tau(a,m;v)}{\gamma}$.[19] [For notational convenience, I do not condition the choices below on v, but the assumption is implicit.]

belief that if any of the other type is hired, those types infer *from that action* that the principal is hiring in such a suboptimal way that the principal prefers to hire none of that type. I rule this out by assuming that the decision to hire a given worker does not change that worker's beliefs about μ.

[18] To avoid the uninteresting problem of the out-of-equilibrium action where the customer complains when allocated $a = 1$, assume that the principal always monitors in that case.

[19] I assume that γ is large enough to guarantee an interior solution: allowing corner solutions would not change the results.

2. If oversight is internal, the preferences of the employees determine the probability of monitoring. Let μ be the fraction of $\overline{v}$ types in the firm. Then the internal monitor sets oversight probabilities to maximize the average utility of the employees. For example, assume that the agent recommends $a = 0$ and there is a probability of $\tau(0, m)$ that the allocation is incorrect. Then the overseer chooses $\rho^*(0, m)$ to maximize

$$\tau(0,m)\rho(0,m)[\mu\overline{v} + (1-\mu)\underline{v} + 1] - \frac{\gamma\rho(0,m)^2}{2}, \tag{7}$$

so that

$$\rho^*(0,m) = \frac{\tau(0,m)[\mu\overline{v} + (1-\mu)\underline{v} + 1]}{\gamma} \tag{8}$$

Only in the case where $\mu = z$, where bureaucrats are chosen to be the population average, does internal monitoring coincide with the preferences of society.

3 Internal Oversight

Begin by considering the case where the principal allows insiders to choose oversight probabilities. Again consider each incentive in turn.

Oversight Simple calculations show that

$$\rho^*(0,c) = \frac{1 + (\mu - z)(\overline{v} - \underline{v})}{\gamma}, \tag{9}$$

$$\rho^*(0,n) = 0, \tag{10}$$

and

$$\rho^*(1,n) = \frac{(1-q)[1 + (z-\mu)(\overline{v} - \underline{v})]}{\gamma}. \tag{11}$$

Note how the preferences of the employees determine the incentive to investigate: more $\overline{v}$ employees result in increased monitoring when the consumer is denied the benefit (and complains), but reduce its likelihood when allocated the asset.

Effort The agent chooses e to maximize the sum of her private benefits and her expected monetary returns. Her expected monetary return potentially depends on the allocation made, whether a complaint occurs, and on the outcome of the investigation. In any outcome where the agent reveals the state truthfully, a complaint illustrates that she was incorrect (as the consumer only complains when an error has been made). As a result, if the consumer complains, the agent earns 0. On the

other hand, if there is no complaint when the agent allocates $a = 0$, the agent is correct and is paid Δ. The expected value of then reporting 0 is $q(e)\Delta$, the prior.

On the other hand, if the agent allocates $a = 1$ the consumer never complains. If the overseer investigates, she earns an expected wage of $q(e)\Delta$. But no information is revealed on the agent without an investigation, and so the agent earns $q(Ee(v))\Delta$, where $Ee(v)$ is the expected effort of type v. Let $\hat{\rho}$ refer to the expected level of oversight.[20] No investigation occurs with probability $\frac{1-\hat{\rho}^*(1,n)}{2}$. Thus, if truth-telling holds the ex ante expected monetary reward is

$$q(e)\Delta\frac{1+\hat{\rho}^*(1,n)}{2} + q(Ee(v))\Delta\frac{1-\hat{\rho}^*(1,n)}{2}$$

The agent then exerts effort to maximize

$$U_a(v) = \frac{[q(e)+(1-q(e))\hat{\rho}^*(0,c)](v+1)}{2} - (1-q(e))\frac{\gamma\hat{\rho}^*(0,c)^2}{2} + \frac{[q(e)+(1-q(e))\hat{\rho}^*(1,n;v)](1-v)}{2} - \frac{\gamma\hat{\rho}^*(1,n;v)^2}{2} + q(e)\Delta\frac{1+\hat{\rho}^*(1,n;v)}{2} + q(Ee(v))\Delta\frac{1-\hat{\rho}^*(1,n)}{2} - e. \quad (12)$$

Note that the principal cannot directly affect the effort of any given agent: the effect of the principal's actions on effort exerted derives solely through who they hire (i.e., through μ). But μ is not directly observed, so that effort levels are based on expected oversight. This limits the ability of the principal to affect effort choices.

Truth-telling The calculation above also implies that the truth-telling constraint is satisfied with internal monitoring. This is because the overseer can condition wages on a. As a result, for any allocation, learning is simply a mean preserving spread on expected ability if the agent reports truthfully. If the agent reports truthfully, her expected wage in equilibrium is then $q(e)\Delta$ (because in equilibrium $e(v) = E(e(v))$). By contrast, if she offers an allocation different from her true belief, she earns $[q(e) - (2q(e) - 1)\rho^*(1,n)]\Delta < q(e)\Delta$. Thus, there is a monetary cost to misallocating. But there are also private benefits to be considered. However, (5) guarantees that these are not large enough to overcome the value of allocating correctly.

Lemma 1 *Assume that (5) holds. Then truth-telling is always satisfied with internal monitoring.*

[20] Expectations are required here as the agent must form a belief about μ.

3.1 Incentives and Selection

To understand the tradeoffs with internal monitoring, consider how v affects incentives. Straightforward differentiation of (12) yields that effort is characterized by

$$\left(\frac{q'(e)(1-\hat{\rho}^*(0,c))(v+1)}{2}+\frac{q'(e)\gamma\hat{\rho}^*(0,c)^2}{2}+\frac{q'(e)(1-\hat{\rho}^*(1,n)(1-v)}{2}\right) + \frac{\Delta q'(e)(1+\hat{\rho}^*(1,n))}{2}=1, \quad (13)$$

where $q''(.) < 0$ guarantees that the second order condition holds. The final term on the LHS provides the monetary effects of exerting effort, and depends on Δ. The remaining terms reflect the private returns to the agent from making the correct allocation. In order to understand how selection affects these incentives, note the effect of v on effort (holding $\hat{\rho}$ constant):

$$\frac{d^2U_a}{dedv}=\frac{q'(e)(\hat{\rho}^*(1,n)-\hat{\rho}^*(0,c))}{2}. \quad (14)$$

Evaluating (14) at $\mu = z$

$$\frac{d^2U_a}{dedv}=\frac{-q'(e)q}{2\gamma}<0. \quad (15)$$

Put simply, hiring agents who are more sceptical of the value of allocating the asset to the agent increases their incentives.[21] The reason for this is that their mistakes of denial ($a = 0$) are corrected with probability $\rho^*(0, c)$, which at $\mu = z$ is given by $\frac{1}{\gamma}$. When γ is low (close to 1) this implies that there is little equilibrium likelihood of the consumer being denied this important benefit. As a result, bureaucrats have little incentive to exert effort as they know that these errors will ultimately be corrected. By contrast, at $\mu = z$, $\rho^*(1, n)$ equals $\frac{(1-q)}{\gamma} = (1 - q)\rho^*(0, c)$. Therefore there is a higher likelihood of the agent being inefficiently given the asset ($[1 - q(e)](1 - \rho^*(1, n))$), as the consumer does not point out these errors. This higher likelihood is gives the agent an incentive to exert effort. Therefore, bureaucrats exert effort more to correct errors of giving the agent the benefit than to correct those of denial.

The reason for biased selection is then a short step. Since some agents care more about errors of acceptance more than others (i.e., those who do not value the claims of consumers as much as the population average), those individuals exert more effort than people who are more concerned with errors of denial. This induces the principal to bias the selection of bureaucrats. While $\rho(1, n; v)$ depends on v and can be used to partially offset reductions in effort, the costliness of this

[21] See Aghion and Tirole, 1997, for other work emphasizing how private benefits can induce effort exertion.

intervention implies that surplus is increasing in e for effort below the first best level.[22] Therefore, those who are more sceptical of the demands of consumers realise that without their contribution, consumers are quite likely to inefficiently end up with the allocation they prefer: as a result, they exert more effort in the (efficient) task of denying them these benefits.

3.2 The Principal's Objective

The principal's objective is to maximize social surplus subject to (i) the consumer's incentive to complain, (ii) the agent's effort incentives, (iii) the incentives to oversee given by (9), (10), and (11), and (iv) the agent's truth-telling constraint.[23] Expected social surplus is given by the weighted average of the productivity of the proportion μ who are type $\overline{v}$ and the complementary probability that are type $\underline{v}$. Let $\overline{e}$ be the effort decision of type $\overline{v}$ and $\underline{e}$ be the effort decision of type $\underline{v}$. Then expected surplus is[24]

$$\begin{aligned}\frac{\mu[q(\overline{e})+(1-q(\overline{e}))\rho^*(0,c)]}{2}+\frac{(1-\mu)[q(\underline{e})+(1-q(\underline{e}))\rho^*(0,c))]}{2}\\ -\frac{[1-\mu q(\overline{e})-(1-\mu)q(\underline{e})]\gamma\rho^*(0,c)^2}{4}+\frac{\mu[q(\overline{e})+(1-q(\overline{e}))\rho^*(1,n;\overline{v})]}{2}+\\ \frac{(1-\mu)[q(\underline{e})+(1-q(\underline{e}))\rho^*(1,n;\underline{v}))]}{2}-\mu\frac{\gamma\rho^*(1,n;\overline{v})^2}{4}-(1-\mu)\frac{\gamma\rho^*(1,n;\underline{v})^2}{4}\\ -\mu\overline{e}-(1-\mu)\underline{e}. \quad (16)\end{aligned}$$

Surplus has a number of components. First, the surplus from the allocation arises both when the agent makes the correct decision (the $q(e)$ terms) and when her errors are corrected (the

[22]This result is based on the assumption that increasing effort is beneficial for the principal, i.e., that effort is below first best. The first best chooses effort levels and oversight to maximize expected surplus given by

$$\frac{[q(e)+(1-q(e))\rho^*(0,c)]}{2}-\frac{[1-q(e)]\gamma\rho^*(0,c)^2}{4}+\frac{[q(e)+(1-q(e))\rho^*(1,n)]}{2}-\frac{\gamma\rho^*(1,n)^2}{4}-e$$

Maximizing this yields the first best levels of effort and oversight given by

$$\rho^*(0,c)=\frac{1}{\gamma},\rho^*(1,n)=\frac{1-q(e^*)}{\gamma},$$

and

$$q'(e^*)=\frac{1}{M\left(1-\frac{1}{4\gamma}-\frac{1-q(e^*)}{2\gamma}\right)}.$$

For M large enough, first best effort always exceeds that privately chosen by the agent.

[23]Note that no individual rationality constraint is included here. This is already incorporated in the program of maximizing surplus, as the reservation utilities of the two types of agents are identical. It is worth noting that the required wage to employ a worker of type v is decreasing in the fraction of workers of the same type, as investigations are more closely matched to their desires as v rise. But this benefit is offset by the loss to the other type of agents, including those who do not work in the firm. As a result, the surplus maximization program includes this effect.

[24]It is always optimal to choose $\rho(0,n)=0$ so this is ignored here.

$(1-q(e))\rho^*(.,.)$ terms). This applies to both types of agents ($\overline{v}$ and $\underline{v}$) and to the cases where $a=0$ and $a=1$ arise. Note that monitoring in the case where $a=1$ is proposed are not focused on the cases where mistakes happened, unlike the case where $a=0$. Maximizing surplus subject to (9), (10), (11), and the agent's incentive constraint (13), yields Proposition 1. All proofs are in the Appendix.

Proposition 1 *Assume that at $\mu=z$, the agents' efforts e are below the first best level. Then if the principal chooses internal monitoring, the optimal level of μ is less than z, i.e., bureaucrats are on average biased against consumers.*

Proposition 1 reflects the two conflicting effects in operation when choosing bureaucracts. The objective is to maximize societal surplus when overseeing. This leads to choosing $\mu=z$ to align with the preferences of society. But agents who are sceptical of the consumer's demands exert more effort. This effects alone would lead to only choosing $\underline{v}$ agents. At $\mu=z$, oversight losses are second order, but the benefit from replacing a $\overline{v}$ agent with a $\underline{v}$ agent is first order, and so the principal chooses the fraction of $\overline{v}$ types below z.

To summarize, this section offers a theory of bureaucracies where bureaucrats are optimally under-responsive to consumer complaints. This lack of (ex post) efficient response arises because they both exert effort and play an oversight role. But these objectives are in conflict, because of the key assumption that internal monitoring depends on the preferences of insiders. I find it realistic that most internal monitoring mechanisms reflect the preferences of the individuals in that setting, so that (for example) a police force disproportionately made up of officers who are sceptical of claims of innocence is unlikely to intervene in the way that society believes to be appropriate. From (9), it can be readily seen that for $\mu<z$, $\rho^*(0,c)$ falls short of the level desired ex post by society (i.e., $\frac{1}{\gamma}$), while from (11) it follows that officers who release suspects will come under excessive scrutiny ($\rho^*(1,n)>\frac{1-q}{\gamma}$). As a result, this section offers a model of endogenously biased internal monitoring caused by the poor signals received by superiors about the performance of agents, which themselves derive from the unreliability of information from the consumers of benefits.

4 External Monitoring

An alternative way of overseeing is to allow society to allow external parties to intervene, perhaps through public opinion, congressional hearings, the courts, the media, and so on. These channels can play an important role in affecting policy and in disciplining many public sector employees. For example, media attention focused on the Los Angeles Police Department was a key factor in

its overhaul, so that now all consumer complaints must now be investigated by Internal Affairs (see Figure 1). Similarly, congressional oversight focused attention on the practices of the IRS, which resulted in the Internal Revenue Service Restructuring and Reform Act of 1998, requiring the creation of an independent Oversight Board to create a climate of accountability within the IRS.[25]

External monitoring has an advantage over internal mechanisms in that it can intervene when society sees fit rather than relying on when bureaucrats see it as appropriate. But a cost of external monitoring[26], as in Prendergast, 2001, is that external parties are more reliant on signals of wrongdoing before they intervene. Specifically, I show that external oversight exacerbates the truth-telling condition above so that bureaucrats who face external monitoring can have the overarching incentive to avoid investigation. This can lead to consumers being allocated benefits when agents believe that they do not deserve them. In effect, the agent realises that by simply giving the good to the consumer, he will not complain and so the agent is less likely to be penalized.

It is worth considering the difference between internal and external monitoring in the context of this model. In the previous section, the internal monitor could observe the recommendations of bureaucrats, and could (in equilibrium) condition wages on that recommendation. This implied that, for example, a police officer who correctly arrested a suspect could be rewarded more than a police officer who did not arrest a suspect but was not investigated. This distinction was important for giving officers an incentive to arrest, and lead to the irrelevance of the truth-telling condition above. But external monitors observe nothing about the details of cases unless the consumer complains, and only by investigating do they get more details on the case. This implies that bureaucrats are similarly rewarded if (without complaint or investigation) (i) they correctly arrested a suspect, and (ii) they did not arrest the suspect. I show that this inability to distinguish may imply that the consumer is given the benefit, even when the agent believes that the consumer does not deserve it. For example, consider the Internal Revenue Service. It has been argued that the increased congressional oversight has resulted in "a sharp roll-off in tax investigations as auditors, fearing for their bureaucratic lives, proceed timidly..[as]..tax collectors are too worried about their jobs to be aggressive" (Star Tribune, 2000).

The objective of the principal is again to hire to maximize expected surplus subject to the same constraints as above. Consider the relevant incentive constraints.

[25]For more details on the effect of the Act on taxpayer rights, see www.henderco.com/1998act.html.

[26]One obvious drawback could be that the cost of investigation is higher for external parties. This will lead to reduced incentives for bureaucrats for the obvious reason that it will reduce the likelihood of an investigation. I do not address this straightforward cost of external oversight by assuming identical investigation costs.

Oversight Here the investigation probabilities are constrained to depend only on whether the agent complains or not. If the consumer complains, an error of denial must have been made (as the consumer only complains if denied) and so the overseer chooses

$$\rho^{**}(0,c) = \frac{1}{\gamma}. \tag{17}$$

If the consumer does not complain, this is either because (i) the agent recommended $a = 0$ and was right or (ii) she recommended $a = 1$. The decision is overturned only if the consumer was given the benefit, and erred. The likelihood of this conditional on no complaint is[27]

$$\frac{1 - q(e)}{1 + q(e)}, \tag{18}$$

and so monitoring is given by

$$\rho^{**}(.,n) = \frac{1 - q(e)}{\gamma(1 + q(e))}. \tag{19}$$

These are the ex post optimal levels of intervention that maximize social welfare. Note that oversight is less precise here than with internal monitoring, as the external parties cannot distinguish within the cases where no complaint was filed. Thus, $\rho^{**}(.,n) < \rho^{*}(1,n)$ above.

Effort If the truth-telling condition is satisfied, the incentive of the agent is to maximize

$$U_a(v) = \frac{[q(e) + (1 - q(e))\hat{\rho}^{**}(0,c)](v+1)}{2} - (1 - q(e))\frac{\gamma\hat{\rho}^{**}(0,c)^2}{2} + \\ \frac{[q(e) + (1 - q(e))\hat{\rho}^{**}(.,n)](1-v)}{2} - \frac{\gamma\hat{\rho}^{**}(.,n)^2}{2} + q(e)\Delta\frac{1 + \hat{\rho}^{**}(.,n)}{2} + q(Ee(v))\Delta\frac{1 - \hat{\rho}^{**}(.,n)}{2} - e \tag{20}$$

Maximizing this expression yields the first order condition for the agent's effort

$$\frac{q'(e)(1 - \hat{\rho}^{**}(0,c))(v+1)}{2} + \frac{q'(e)\gamma\hat{\rho}^{**}(0,c)^2}{2} + \\ \frac{q'(e)(1 - \hat{\rho}^{**}(.,n))(1-v)}{2} + \frac{\Delta q'(e))(1 + \hat{\rho}^{**}(.,n))}{2} = 1 \tag{21}$$

For exactly the same reason as in Section 3, $q(\overline{e}) < q(\underline{e})$, so the principal has an incentive to choose $\mu < z$ to improve incentives. Lemma 2 immediately follows.

[27] From an ex ante perspective, there is a probability $\frac{1}{2}$ that either state occurred. Then there is a probability $\frac{q(e)}{2}$ that the agent correctly allocated $a = 0$, and a probability $\frac{1}{2}$ of $a = 1$, which is corrected with ex ante probaility $\frac{1-q(e)}{2}$. Summing these yields (18) as the probability of correction.

Lemma 2 *Assume that (5) holds, and that effort is below the first best level. Consider any equilibrium where truth-telling holds. Then* $\mu = 0$.

This immediately follows from the fact that effort is higher for $\underline{v}$ types than for $\overline{v}$ types. Since with external oversight, monitoring probabities are independent of make-up of the sector, there is no reason to hire $\overline{v}$ workers. This is one of the advantages of external oversight, namely, that principal chooses only those agents who have better incentives as oversight decisions are made with the objectives of society in mind and no longer depend on the make-up of the burueacracy.

4.1 Truth-telling

Yet these effects only hold if truth-telling is not a problem. When the agent believes that the consumer should be allocated the good, truth-telling is not a binding constraint (assuming (5) holds) as the agent is both carrying out the optimal action and is minimizing the likelihood of a complaint. Consider an equilibrium where $\alpha_a = a$. The problem arises when $\alpha_a = 0$ and the agent should deny the benefit to the consumer. If she denies the benefit to the consumer, with probability $1 - q(e)$ she is incorrect and so will be the subject of a complaint. The complaint implies that the agent was incorrect and so the expected wage of the agent is given by $q(e)\Delta + (1 - q(e))0 = q(e)\Delta$. By contrast, suppose that the agent now offers the out-of-equilibrium allocation $a = 1$ to the consumer. The consumer never complains, and investigations occur with probability $\rho^{**}(., n)$ above. If an investigation occurs when the agent allocates $a = 1$, the agent is penalized if she was correct and rewarded if her guess was incorrect, which occurs with probability $1 - q(e)$. Finally, the agent earns a wage w_0 if not investigated. In any equilibrium where truth-telling occurs, this arises either because no complaint was made when allocating $a = 0$ or because the agent allocated $a = 1$. The conditional probability of the agent being correct is $\frac{2q(e)}{1+q(e)}$ which exceeds $q(e)$. In other words, not being investigated increases the wages of the agent from the prior to

$$w_0 = \Delta \frac{2q(e)}{1 + q(e)} > q(e)\Delta. \tag{22}$$

Then the expected monetary return to offering $a = 0$ is $q(e)\Delta$, while the monetary return to misallocating $a = 1$ is $\rho^{**}(1, n)[1 - q(e)]\Delta + (1 - \rho^{**}(1, n))w_0$. Adding in the private benefits for a worker of type v yields that truth-telling only then arises if

$$\begin{aligned} q(e)(\Delta + S_i(0,0)) + (1 - q(e))[1 - \rho^{**}(0,c)]S_i(0,1) + (1 - q(e))\rho^{**}(0,c)S_i(1,1) \geq \\ (w_0 + (1 - q(e))S_i(0,1))[1 - \rho^{**}(.,n)] + (1 - q)\hat{\rho}^{**}(.,n)\Delta + (1 - q(e))S_i(1,1). \end{aligned} \tag{23}$$

This constraint is more easily satisfied for $\underline{v}$ employees, for two reasons. First, their private benefits from allocating $a = 0$ are higher than for the $\overline{v}$ workers. Thus the net effect of private benefits for $\underline{v}$ agents is to relax the constraint. Second, these agents exert more effort, and the truth-telling constraint is relaxed as q increase. Since Lemma 2 illustrates that the principal would want $\underline{v}$ workers anyway, I only need to consider the truth-telling condition for those workers.

Intuitively, truth-telling requires that the cost of a complaint to the agent not be large. If the agent recommends that the consumer be denied, the agent's ability is revealed only if the news is bad. Alternatively, the agent can incorrectly allocate $a = 1$ in which case they face the prospect of not being investigated (and getting a wage exceeding $q(e)\Delta$) or being investigated, in which case they are likely penalized: only if they were wrong are they mistakenly perceived to be talented. There is no necessary reason to expect the return from truth-telling to be higher. After some manipulation, (23) simplifies to

$$q - \frac{2q}{1+q} + \rho^{**}(1,n)[\frac{2q}{1+q} - (1-q)] \geq \frac{(1-\rho^{**}(0,c))(1-q(e))(1+\underline{v}) - (1-\rho^{**}(.,n))q(e)(1-\underline{v})}{\Delta}. \tag{24}$$

There are two parts to the truth-telling condition. The term on the right hand side gives the private benefits from allocating $a = 1$ rather than $a = 0$. This is negative as $\underline{v} < 0$. My interest is in the left hand side, the monetary effects of denying the benefit to the agent. Proposition 2 illustrates the importance of this condition.

Proposition 2 *Assume that the truth-telling condition for the $\underline{v}$ agent is violated. Then the surplus from internal monitoring (weakly) exceeds that of external monitoring.*

But truth-telling need not hold. For example, consider the case where Δ is large so that the monetary consequences of actions are most important. Then as the cost of investigation (γ) gets large, the left hand side of (24) is negative as $\rho^{**}(1,n) \to 0$ as $\gamma \to \infty$. Then for truth-telling to hold, it must be that the likelihood of investigation when the consumer allocates the good to the consumer is sufficiently high. If the principal never oversees these cases, the agent always offers $a = 1$ as this would for sure increase his pay above the prior. It follows that there is a critical value of γ above which truth-telling condition is a problem, and for Δ large enough, internal monitoring is used.

The incentives of police officers seem particularly apposite here. To guarantee that they deny benefits to suspects, it must be the case that their decisions are overseen reasonably frequently. But in many cases, this is extremely difficult. Consider this calculus for a police officer who is charged with, say, arresting drug suspects. This involves identifying suspects and may involve searching,

detention, and use of force. Sometimes these procedures backfire and a "suspect" is treated in a way that can seem incompetent, malicious, or racist (New York Times, 2000a, 2000b). This can land the officer in question in trouble. An alternative is simply to let suspects pass, thereby avoiding the possible problems that arise with a false arrest or an accusation of assault. However, it may be extremely hard to monitor those cases where the person was not arrested. Simply randomly monitoring citizens will hardly be efficient, as they are unlikely candidates. In terms of the model above, γ is likely to be high, and so it may be very difficult to give officers incentives to reveal information credibly.

5 The Los Angeles Police Department

One way to test the ideas of this paper is by tracking how a change in oversight affects the behavior of a group of public officials. In this section, I consider the effect of recent changes that have increased the ability of consumers to complain against Los Angeles Police Department officers. The increased oversight of officers was partly caused by well-publicized media accounts of infractions by anti-gang squad officers during the mid-1990s.[28] These concerns lead to the Rampart Board of Inquiry, whose recommendations affected the policies of the Department. One of these policies was that the Department became more responsive after January 1998 to perceived infractions by increasing the penalities that officers face for malfeasance. For example, all complaints against officers are now investigated by the Internal Affairs division of the police force.[29] The result of these investigations has often had disastous effects on the careers of police officers. An implication of the model above is that this is likely to result in officers avoiding actions that could lead to an investigation, even when those actions are efficient. (In the formal terms of the model, both $\rho(0, c)$ and Δ have increased.)

In this section, I provide some exploratory evidence[30] that officers are indeed avoiding such confrontations. In summary, the changes appear to have caused less aggressive policing, which is correlated with an enormous increase in the homicide rate over the last two years. Thus, as in the model, increased external oversight may have backfired, by reducing the incentives for officers to

[28] Among these claimed infractions were the use of excessive force, framing suspects, and cocaine dealing.

[29] The Internal Affairs division, rather unfortunately named, acts like external monitors in my model, as their interest is largely induced by consumer complaints, and they tend to be poorly informed about cases unless they respond to consumer complaints to investigate.

[30] One particular group in the police department, the anti-gang squad in Rampart, was the particular focus of the department's investigations, and resulted in a Board of Inquiry. In their report, they provide data on the procedures and practices in Rampart and five other bureaus. The five other bureaus are Hollywood, Southeast, Van Nuys, Wilshire and the 77th Street Area. The data in Figures 2 and 3 come from these six areas. Some of the data used here are available at www.lapdonline.com.

do their primary job, namely, confronting and arresting criminals.

The premise of the paper is that a primary mechanism for overseeing bureaucrats is through its responsiveness to complaints ($\rho(0, c)$ in the terminology of the model). Among the first changes implemented after the scandal were procedures that allowed the public to fill out complaints on line,[31] and an order by the Chief of Police that every complaint against a police officer be forwarded to the Internal Affairs Division. As can be seen from Figure 1, the Department increased the fraction of complaints that led to an investigation after 1998.

Figure 1: Percentage of Complaints Investigated

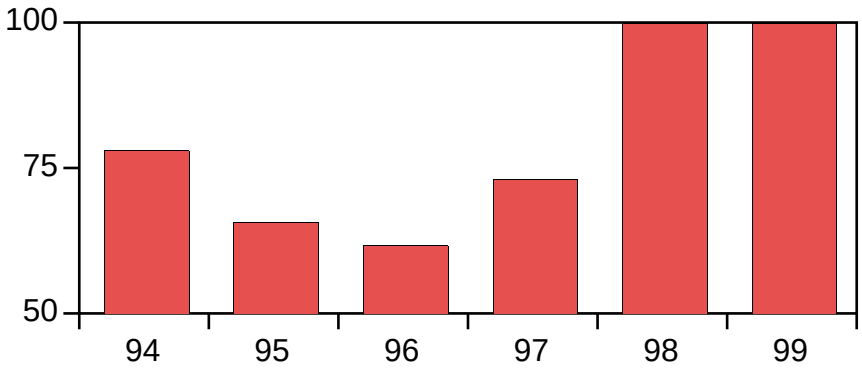

Not surprisingly, this resulted in a huge increase in complaints.[32] Data are not provided for all years for the entire force, but the department notes that 5339 complaints were made against officers (and civilian staff) in 1998 and 5280 in 1999, up from 1912 in 1997, which itself was down from a previous peak of 2359 in 1992.[33] (There are approximately 10,000 officers in the L.A.P.D.) One might imagine that this increase in complaints after 1998 would involve more marginal infractions being brought forward than in prior years, where only the most serious problems would be worth pursuing. If this was so, a lower fraction of these new complaints should ultimately be sustained against the officer after the changes. Despite this, the fraction of complaints sustained (against

[31] See www.policeabuse.org, a website devoted to Police Complaints, for obstacles placed in the way of individuals wishing to complain about an officer in other juristictions.

[32] In the model, the consumer always complains if he is denied the asset and the bureaucrat makes an error: thus, the propensity to complain is independent of $\rho(0, c)$. But this is an artifact of the assumption that complaints are costless. In Prendergast, forthcoming, I allow consumers to have a non-trivial cost of complaint, where higher responsiveness of the bureaucracy makes complaints more likely.

[33] Data are provided from 1994 to 1998 for the six areas covered by the Board of Inquiry. These data show a gradual decline in complaints from 1994 to 1997, with a 108% increase in 1998.

the officer) increased from 50% in 1997 to 53% in 1998, which suggests not only that consumer complaints were now easier to file and investigate, but also that they were more likely to harm the officer involved.

These changes have not been well received by the officers, who now perceive themselves more at risk of suspension and losing their jobs. For example, between 1992 and 1997, an average of 13 officers per annum were removed from the force for malfeasance. In 1998, 55 officers were fired, with 44 fired in 1999. Cannon (2000) argued recently that in total these changes have resulted in over 800 officers being disciplined, 113 terminated, and many who have left the force rather than be investigated. According to Tharp (2000), "more than 300 officers have either been forced out or quit in the face of complaints". Furthermore, in both 1998 and 1999, almost 5% of all officers were suspended. Thus, it seems that the instruments used by this bureaucracy correspond closely to those in the model: they can change their responsiveness to complaints, and the propensity to fire and suspend conditional on infractions after a case has been investigated. These are typical career concern incentives.

The premise of the paper is that increased responsiveness to consumer complaints changes the behavior of officers in ways that avoid consumer complaints (or worse, media attention). To use the conclusion of a recent New York Times article, the upshot may be police officers who "too often cruise down the street in the patrol cars - a practice known within the L.A.P.D. as drive and wave - instead of engaging in aggressive policing" (Cannon, 2000, p. 62). There are a variety of ways in which police behavior could change. First, officers may become more circumspect about the use of force, where officers are now more likely to be disciplined for going beyond acceptable levels of restraint. Yet such restraint is often necessary. Second, they may avoid confrontations with individuals which could result in accusations of brutality. Third, they are likely to become less likely to discharge their weapons, given accusations of officers shooting indiscriminately. (For example, witness the recent riots in Cincinnati caused by the shooting of a black teenager, as in New York Times, 2001.) Fourth, they may simply cut down on arrests, as suspects do not complain when set free.[34]

I provide four pieces of information to suggest that police have become less aggressive in their behavior: use-of-force data, officer-involved shootings, attacks *on* officers, and arrest rates.[35] It is

[34]It is worth bearing in mind that in general there is no reason that increased oversight will result in less aggressive policing. Increased oversight could, for instance, penalise an officer for failing to use force to subdue a suspect in the same way as it could penalise him for using excessive force. But the premise of this paper is that because it is consumers who raise flags to point out malfeasance, officers will become more reticent in their dealings with suspects.

[35]It is important to note that these responses are not necessarily inefficient. It could be that the objective of the changes is to reduce "aggressive" policing: all that I argue here is that it appears that officers respond to the increased oversight, and that there may be a price from this reaction.

worth emphasizing at the outset that these data are time series averages so it is difficult to devise a precise test of incentives on behavior, especially as it took time for officers to realise the gravity of the changes in oversight. However, I believe that the range of data offered suggests a change in the behavior of police officers towards suspects in the last three years. First, consider evidence on formal use of force by officers in Figure 2.

Figure 2: Use-of-Force Per Arrest (%)

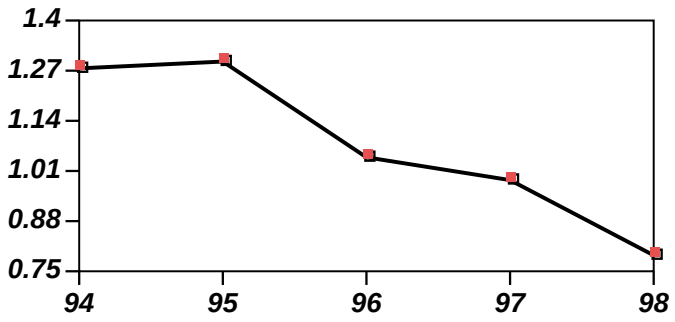

As can be seen above, the use of force per arrest fell by almost 20% in 1998, the year of the change in policy.[36] However, it is difficult to claim that this was necessarily caused by the change in oversight as there was a negative trend in the propensity to use force since 1994. While it may be a response, the claim seem unproven.[37] (Unfortunately, no data on use-of-force are available after 1998.) Another measure provided in Figure 3 is the frequency of officer-involved shootings: again, the paper predicts that increased oversight of officers will result in less shootings.[38]

[36]These data are for the six areas covered by the Rampart Board of Inquiry. For the department as a whole, the reduction is lower, at 13%.

[37]It is not clear how to normalize use-of-force data. Here I normalize by incidents per arrest, but since arrests are endogenous to the oversight process, this may be inappropriate. Another method is to use the number of incidents of use-of-force per officer. This shows an 4% increase in incidents in 1997, and a 15% decline in 1998, which is more suggestive of a response to the policy by officers.

[38]I have excluded shooting at animals and accidental discharges here.

Figure 3:Officer-Involved Shootings per Officer (%)

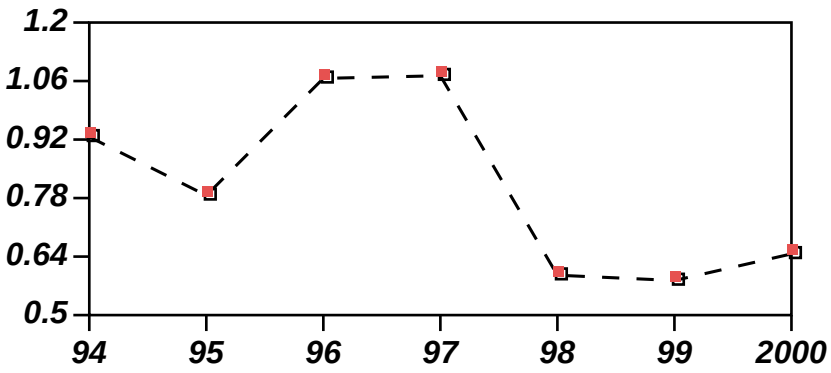

As can be seen from Figure 3, shootings declined in 1998 by 45%, which has remained in the succeeding two years.[39] These data suggest some response by police officers, but say little about whether there is a price to pay for such reactions. Perhaps the most interesting and relevant evidence comes from changes in gang activity in Los Angeles. After many years of decline, gang-related violence in Los Angeles increased significantly over the last two years: for instance, by December, 2000, the homicide rate was 143% higher than at a comparable time the previous year, and drive-by shootings up 69%.

The premise of the paper is that increased oversight will result in officers being less likely to confront suspects. One measure which reflects the aggressiveness of policing is the number of assaults *on* officers. If officers are more likely to engage in "drive and wave", where they disengage from troublesome situations, we would expect to see that assaults on officers have declined recently, as less aggressive policing results in fewer confrontations. Notably, by December 2000, assaults on police officers were down 35% from 1997. This may result in a higher crime rate if such aggressive policing has a payoff.[40] Figure 4 provides recent data on this, where I compare homicide rates to attacks on officers, from 1992 through the end of 2000. Both are normalized to 100 in the year

[39]In 1996, there was a change in department policy towards what is called bean-bag deployment, where bean-bag rounds were fired rather than live bullets. This factor increases the shooting numbers from 1996 to 1998: before 1996, these rounds were not used. When bean-bag incidents are excluded from the calculations, there was a decrease of 15% in shootings from 1996 to 1997, and a reduction of 40% from 1997 to 1998, where the number of shootings in 1998 were only half its 1994 level.

[40]According to the New York Times (2000), the result of this policy is that "violent crime is up 9 percent this year in poor neighborhoods after years of decline. Gang-related homicides are up 116 percent. Arrests and field interviews of suspected criminals are down" (p.62). To phrase this in the terminology of this paper, giving the customer what they want looks like an increasingly popular strategy to officers faced with this increased responsiveness to complaints.

preceding the change, 1997.

Figure 4: Gang-Related Homicides and Officer Attacks (1997 = 100)

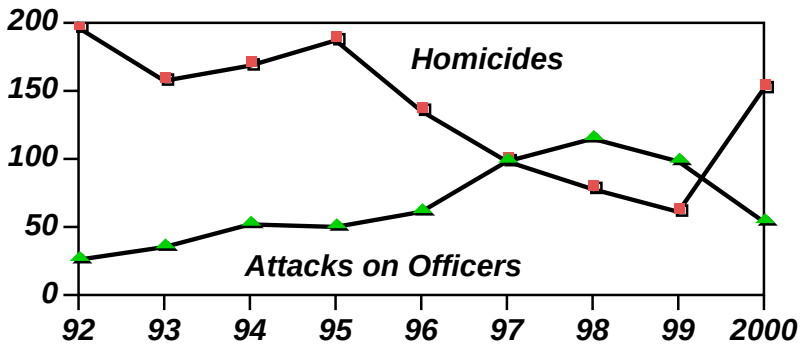

These data suggest a negative relationship between attacks on officers and gang-related homicide rates. Until 1998, gang related homicides fell as attacks on officers rose. However, after the change in oversight in 1998, assaults on officers have declined while the homicide rate has soared to a rate not seen since the mid-1990s. These trends have continued through the first 4 months of 2001, where gang homicide rates are a *further* 100% higher than for the same period in 2000, while attacks on officers are *down* a further 2.6% from 2000. One interpretation of this is that since 1998 officers have been responding to increased oversight by actions which, although keeping them out of trouble, also result in higher crime.[41]

In the interests of completeness, I also provide econometric estimates of the time series correlation between gang-related homicides and attacks on officers, using monthly gang data from March 1998 to April 2001. (The relevant gang related statistics are not available on a monthly basis before March 1998.) As there is likely some lag between the policies of officers and the response of gang members, I estimate the homicide rate in a given month with (i) assaults on officers in the contemporaneous month, (ii) assaults on officers in the previous two months, and (iii) month dummies, where the final term picks up seasonality effects in the homicide rate. The estimated OLS regression is given in Table 1.

[41]There is obviously the possibility of unobserved heterogeneity generating this correlation. Yet most unobserved heterogeneity that I can think of is likely to cause positive correlation between the two measures. For instance, an unobserved shift in the returns to crime will both likely increase the crime rate and the confrontations that officers have with criminals, thus generating a positive relationship. There is also an important issue of causality here: it could be that officers simply avoid confronting gang members when homicide rates are high.

Table 1: Estimating the Monthly Homicide Rate.

H(t)	Constant	Dummies	OA(t)	OA(t-1)	OA(t-2)
Coefficient	39.90	-	-0.876	-0.417	-0.886
T-statistics	(6.18)	-	(-3.14)	(-1.35)	(-2.8)

where H refers to the gang homicide rate, OA is the number of attacks on officers, and t is the month. The t-statistics are provided below the estimates. The R-squared is 0.59 and $N = 35$.[42] The significant estimates illustrate that the negative correlation evident in Figure 4 continue to hold in the monthly data for the post 1997 period.

It is also worth pointing out that in general, the available measures of gang crime are positively correlated with each other. For instance, the gang-related aggravated assault rate has a correlation coefficient of +0.64 with the homicide rate. This should not be surprising as there are likely factors that shift the demand for many types of crime. But this correlation is not true for these measures of crime and the level of attacks on officers; both the assault rate and the homicide rate are negatively correlated with the contemporaneous level of Attacks on Officers (the contemporaneous correlation coefficients are -0.04 and -0.37 respectively). Thus the factors causing attacks on officers seems to be different than for other measures of gang activity.

Finally consider the arrest record of officers. The central premise of the paper is that suspects are more likely to be released when oversight increases. Figure 5 plots the ratio of arrests to crimes for the more serious Part 1[43] crimes from 1996 through September 8, 2001.[44]

[42] None of the month dummies are significantly different from zero: there is little seasonality in homicides.

[43] Part 1 crimes consist of homicide, rape, robbery, aggravated assault, burglary, larceny, and auto theft.

[44] An alternative is simply to plot the total number of Part 1 arrests. These data are 42,800 (1996), 41,900 (1997), 40,700 (1998), 37,300 (1999), 33,700 (2000) and an annualized estimate of 31,300 for 2001 (the number of arrests until September 8 was 17949, compared to 18379 at a comparable point in 2000). These raw data would suggest a faster response (in 1998) to the increased oversight than in Figure 5.

Figure 5: Arrest Rate Per Part 1 Crime

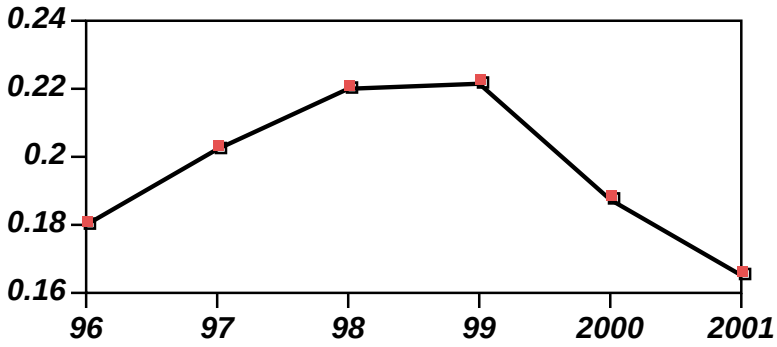

As can be seen, the steadily rising arrest rate of the Los Angeles Police department was abruptly halted in the last two years, with a 15.3% decline in arrest rates between 1999 and 2000, and a further decline in 2001. These changes coincide with the point when the changed disciplinary measures began to have effect.[45] These data concern arrests for all Part 1 offences, which are considered the most serious offences. However, the majority of these are made up of larceny, auto theft, and burglary, which are less serious than homicide, rape, or aggravated assault. In Appendix 2, I decompose these crimes into their constituent parts, where I show that there has been a decline in arrests per crime after 1998 for all of these categories except auto theft, where there appears to be little change. Thus, there is no evidence that the change in officer behavior is focused on the less serious offences.

If police officers change their behavior to avoid attention, they are especially likely to do so in cases where the victims of crime are diffuse, and are unlikely to complain about a particular officer's actions. One instance where this is likely to arise is failure to solve narcotics crimes, where the victims are often unaware of the actions of the officer. The time series of narcotics arrests is given in Figure 6, where I obviously do not normalize by total crimes here, as this is unknown. (It

[45]There is two potentially important problems with the 2001 data. First, many of the most recent crimes have not been solved, as not enough time has elapsed to do so. This reduces the estimate for 2001 for reasons other than increased oversight. To address this, I compared the data up to September 8, 2001, with equivalent data up to September 8, 2000. At an equivalent point in time in 2000, the arrest per crime rate was 0.178 (compared to the end of year arrest rate of 0.188) suggesting that the data for 2001 are underpresented by approxiately 0.01 or thereabouts. This adjustment would not change the outcome in any qualitative way. Second, there are seasonality concerns to worry about, as some crimes may be more common at some points of the year. There are two reasons to suggest that this is not a major concern. First, there is little seasonality in Part 1 crimes in Los Angeles. See the 1999 and 1998 crime digests at www.lapdonline.com for more details. Second, this would only be a concern if the arrest/crime ratio varies considerably across the year, which seems less likely than the crime rate itself.

is precisely the fact that these crimes often go unreported that causes the agency problem here.) The number of narcotics arrests shows a significant fall in arrests over the last two to three years, once again suggesting a change in how policing is carried out after the changes in oversight.

Figure 6: Narcotics Arrests

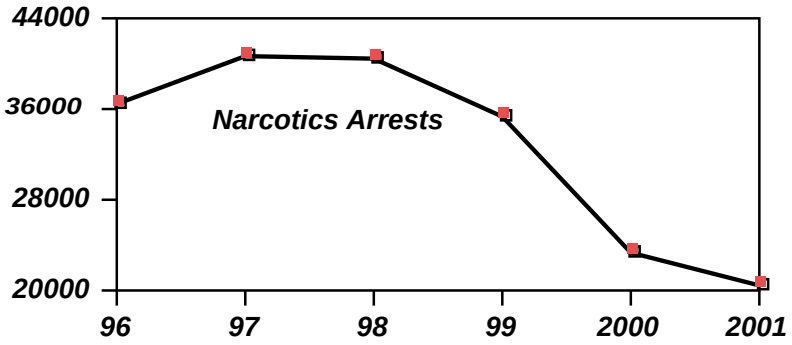

To conclude, it is important to point out that these data cannot be used to claim that the change in oversight was necessarily bad. Indeed, as there is no measure of police performance that all can agree upon, some likely see the changes as beneficial and others see it as a regressive step. The purpose of this section has simply been to point out that the changes in oversight have had a cost, in this case through reductions in the kind of aggressive policing that is often credited with reducing crime in Los Angeles during the 1990s. While most probably realised that these changes would have some restraining effect on officers, the data here suggest that the behavioral responses are large. For example, narcotics arrests were down almost 40% between 1998 and 2000, and down further this year. As such, this section sounds a warning note to those who believe that the way to resolve public sector inefficiencies is simply through more external oversight.

6 Extensions

6.1 Prices

So far, I have alluded to the agents above as public sector employees but have not described in any detail why the kind of activities carried out in the private sector is likely to differ. One important distinction that often arises between the public and private sectors is the more limited role for

purchasing benefits in the public sector. For instance, citizenship is not for sale, but the services of a (private sector) optician are. More generally, the typical transaction in the private sector involves payments from consumers. In this section, I briefly describe how this affects the results above.

Remember that the consumer's utility *from the allocation* is given by $V(A, \alpha)$ where $V(1, \alpha) > V(0, \alpha)$. But suppose that the utility function of a consumer is $V(A, \alpha) - pA$, so that the allocation of the good costs the consumer p and is separable. Furthermore, consider the case where there exists a price p such that

$$V(1, 1) - p > V(1, 0) \tag{25}$$

and

$$V(1, 0) - p < V(0, 0). \tag{26}$$

In the private sector example above, a patient goes to an optician, unsure if he needs spectacles. If the optician gives him glasses that he does not need, he is out of pocket by p for little benefit. But if he truly needs the glasses, his utility increases. If these conditions holds, then the principal can use prices to attain a better allocation than is achieved above.[46] To see this, assume that the principal choose a price such that (25) and (26) hold. This has one important change from the previous model, in that now the consumer complains if given the wrong allocation both when $a = 1$ and when $a = 0$.[47] Thus, the patient complains about the optician if either (i) he receives glasses he does not need, or (ii) he does not receive glasses that he needs. In the previous section, not paying for the good meant that no complaints were registered when he received the good. This changes the outcome in the following way.

Proposition 3 *Assume that at $\mu = z$, the agent's efforts e are below the first best level. Then if the principal can charge consumers a price p that satisfies (25) and (26), there is no reason to bias hiring. Furthermore, $\overline{e} = \underline{e}$.*

Proposition 3 implies that when consumers can be charged for goods in such a way as to make their responses more informative of error, there is no longer any need to bias the make-up of the workforce to induce effort exertion.[48] The idea that prices can be used to induce credible feedback

[46] Ex ante transfers can additionally guarantee that consumers pay 0 on average if this is desired.

[47] For example, consider the difference between a police force and a private security firm, who are deciding on manning for a project. If a police force assigns too many officers for the security detail, the consumer does not voice any concerns, as he does not pay for this. On the other hand, a private security firm charges for manpower at competitive rates, and is likely to upset the customer if he suggests a large number of officers.

[48] The starkness of the result that bias disappears when prices can be charged arises from the assumption that states 0 and 1 are equally likely. Changing this assumption will cause bias even when consumer feedback is credible. This is because the principal cares more about one state than another (as it is more likely to occur) and will select workers

has been previously considered in Prendergast, 2001: the issue of interest here is how that affects the selection of agents. As the model is specified, the principal would be indifferent between internal and external monitoring (as they have identical cost functions) but nevertheless it is clear that there is no need for bias in the selection of workers to induce effort exertion. As a result, it is the fact that consumers do not pay for assets, a central characteristic of the public sector, that generates this result.

Again, it is worth emphasizing here that this paper says nothing about whether an activity should be carried out in the public or private sector. Instead, all that this section says is that if prices can be charged for an activity, it may be possible to alleviate the agency concerns. It should be clear that there are many private sector examples which find it difficult to induce credible consumer responses. Those involving insurance are one case where it may not be efficient to require consumers to pay for benefits, as the cost of the monetary transfers are not separable from the allocation. As a result, it may not be efficient to require patients to pay to be approved for medical procedures if it violates the insurance purpose of these programs. There are also many private sector examples involving incentives where payments are unlikely to be a plausible way of inducing consumers to self-select. For example, it should matter to our students whether they pass our exams: it is highly unlikely that transfers can be used to induce students to self-select their grades, or workers to self-select their promotions.

Ultimately, the issue addressed here is whether there exists some known allocation which is inefficient, yet which reduces the number of complaints made by consumers? If agents know of such allocations, the problems outlined here exist. While there are private sector examples that satisfy this condition, the fact that consumers do not pay for most public sector benefits makes this the most obvious place to observe this kind of behavior.

6.2 Informational Assumptions

It is simplistic to assume that, without formal investigation, internal monitors always observe proposed treatments and that external monitors never do. Insiders sometimes know little. For example, a police chief often will have no idea about unarrested suspects, or a superior of a custom's official may have little idea about potential smugglers. On the other hand, there are cases where external monitors observe more than is allowed here, for example when the police fail to arrest a suspect in a well known case. As such, I make no claim that all internal and external oversight

who are most likely to find errors in these states costly. Therefore, if it is more common to turn down applicants for unemployment benefits more than it is to accept them, the principal will find this an additional reason to hire agents sceptical of the consumer's claims, and vice versa.

takes the form above. Instead, the objective of the modeling sections is to illustrate that in many instances, insiders are better informed about the cases where benefits were given to consumers than when institutions like the media or congress are used to investigate. The importance of the paper's distinction between internal and external monitoring then depends on the empirical differences in their ability to observe details of cases.

Another important assumption is not only that internal overseers have the ability to observe cases where consumers were denied benefits, but also that they have the instruments to reward agents for correctly doing so. But many overseers simply do not have the discretion to do so, whether for legal reasons or otherwise. For example, in many public agencies the only way to offer incentives is through the threat of firing, which is only carried out for gross violations of acceptable performance. In effect, there are few carrots for unusually good behavior. In those instances, truth-telling is likely to be a serious problem, even when overseers can observe cases where the agent correctly made the difficult decision to deny benefits: while the overseer can see that the agent has made the correct decision, there is no way to reward her for doing so. Thus, the outcomes here rely not only on the ability to observe behavior but also to reward bureaucrats for acting appropriately.

It is also unrealistic to claim that internal monitoring of public agencies reflects only the preferences of the rank and file. There is surely some ability for superiors to impose oversight which sometimes upsets the bureaucrats who are making the relevant decisions. As such, this paper offers an insight into why internal overseers of public bureaucracies are likely to have a very difficult job; their subordinates likely have preferences which place them at odds with external consituencies, whose objectives these superiors are employed to implement. But equally it is naive to think that internal monitoring does not reflect the preferences of the employees. First, many bosses are promoted internally, and are likely to reflect the preferences of insiders. Second, much of the relevant information on malfeasance is not spotted by superiors, but by co-workers, who must be relied upon to pass on relevant details. As such, the biases of bureaucrats is likely to be reflected in the way that internal oversight operates.

6.3 Limited Private Benefits

Bias by bureaucrats played two roles in the model above. First, it changed the incentives of the bureaucrats to exert effort. Second, it harmed oversight by intervening in ways different to the preferences of society. There is a final way effect of bias in that it can cause agents to allocate benefits based on private rather than social surplus. For example, an immigration official who gets private benefits from denying applicants entry may do so even when the applicants are qualified.

This effect has been ignored in the model by the assumption of limited private benefits: they were assumed small enough that they never overcome the desire to allocate honestly.

A more general model would relax this, perhaps by allowing outcomes to be continuous, rather than the binary outcomes here. In this way, there may be some marginal cases where society would like a consumer to be treated in one way, but the agents instead treats the person according to the agent's private benefits. For example, some arrests would be made that society would set free. Yet this generalization is unlikely to change the message of the paper, namely, that some bias in the composition of bureaucrats is likely to be beneficial. The idea behind the results above is simply that distortions are traded off against one another when choosing the make-up of the bureaucracy. While this extension will add one more distortion, the qualitative effects of biasing hiring will likely remain. If so, problems of internal oversight once again arise. While the degree to which agents will be biased will be tempered by this additional cost, there will remain the tendency to hire bureaucrats sceptical of the merits of the consumers case, and so the difficulty of using insider oversight is likely to remain.

7 Conclusion

Agency theory is concerned with understanding the constraints that limit the efficiency of institutions. One might imagine that agency theory would then focus on those institutions which seem least efficient, government bureaucracies being a prime candidate. Yet ironically the vast majority of agency work has been on occupations and institutions where these problems are least likely to be observed, such as executives in Fortune 500 companies, salesforce workers, and sharecropping farmers. See Prendergast, 1999, Malcomson, 1999, and Gibbons and Waldman, 1999, for details. The upshot of this is that by now we have a relatively good understanding of the role of pay-for-performance contracts and monitoring for workers whose objectives are easy to specify and monitor. Yet we have learned little about what motivates public sector bureaucrats such as police officers, benefit officials, doctors (in many countries), army officers, teachers, tax collectors, immigration officials, and so on. Much of what goes on in these occupations remains murky, despite the fact that in many countries they are a majority of the workforce. The purpose of this paper has been to formally address the environment in which bureaucrats operate to generate a better understanding of (i) how bureaucrats are selected, and (ii) how their performance is overseen.

The ultimate message of the paper is that it is futile to expect the solutions typically used in the literature to be as effective for the allocation of benefits, because bureaucrats are inherently harder to monitor. Agency concerns largely arise because it is difficult to monitor workers, and so

consumers are used to point out errors. But consumers are less reliable for many of the activities of the public sector, because so many of the goods given out in the public sector are free. As a result, more costly mechanisms of both providing incentives and overseeing performance are needed. In this paper, this took the form of hiring bureaucrats who are dubious about the merits of consumer claims, and of being left with the difficult choice of relying on them to oversee (though they are biased) or using less informed outsiders, where bureaucrats can change their actions to stay out of the limelight, as illustrated by the police data.

Appendix 1: Proofs of Results

Proof of Lemma 1:

A sufficient condition for Lemma 1 is that the private benefits from misallocation are negative. The relative private benefit from allocating $a = 0$ rather than $a = 1$ when $\alpha_a = 0$ is given by

$$q(e)S_i(0,0) + (1-q(e))\rho^*(0,c)S_i(1,1) \geq \\ (1-q(e))S_i(1,1) + q(e)[1-\rho^*(1,n)]S_i(1,0) + q(e))\rho^*(1,n)S_i(1,1). \quad (27)$$

Substituting yields

$$q(e) + (1-q(e))\rho^{**}(0,c)(1+v) \geq (1-q(e))(1+v) + q(e)[1-\rho^*(1,n)]v + q(e)\rho^*(1,n)(1+v). \quad (28)$$

For an allocation of $a = 0$, this constraint can only be violated for type $\overline{v}$. However, simple manipulations shows that this condition is always satisfied if (5) holds.

The relative private benefit from allocating $a = 1$ rather than $a = 0$ when $\alpha_a = 1$ is given by

$$q(e)S_i(1,1) + (1-q(e))\rho^*(1,n)S_i(0,0) + (1-q(e))[1-\rho^*(1,n)]S_i(1,0) \geq \\ (1-q(e))S_i(0,0) + q(e))\rho^*(0,c)S_i(1,1). \quad (29)$$

Substituting yields

$$q(e)(1+v) + (1-q(e))\rho^{**}(1,n) + (1-q(e))[1-\rho^*(1,n)]v \geq (1-q(e)) + q(e)\rho^*(0,c)(1+v). \quad (30)$$

For an allocation of $a = 1$, this constraint can only be violated for type $\underline{v}$. However, simple manipulations shows that this condition is always satisfied if (5) holds.

Proof of Proposition 1:

First, if $\rho(.,c) > \rho(.,n)$, the consumer only complains if $a = 0$ and the agent is incorrect in any Bayesian Nash equilibrium. Second, given this, Bayes Law requires that oversight is given by (9), (10), and (11). With these oversight probabilities, and the assumption that a given worker does not update on μ from the fact that he was hired, effort decisions must then be given by (12). By the assumption in footnote 17, their beliefs are tied down by the equilibrium choice of the principal on the assumption that the principal's action does not change the perception of μ. The objective of the principal is that to maximize (16) subject to (9), (10), (11), and (12). Straightforward

differentiation of (16) with respect to μ yields

$$\frac{[q(\overline{e})+(1-q(\overline{e}))\rho^*(0,c)]-\frac{\gamma\rho^*(0,c)^2}{2}}{2}-\frac{[q(\underline{e})+(1-q(\underline{e}))\rho^*(0,c)]-\frac{\gamma\rho^*(0,c)^2}{2}}{2}$$
$$+\frac{[q(\overline{e})+(1-q(\overline{e}))\rho^*(1,n;\overline{v})]-\frac{\gamma\rho^*(1,n;\overline{v})^2}{2}}{2}-\frac{[q(\underline{e})+(1-q(\underline{e}))\rho^*(1,n;\underline{v})]-\frac{\gamma\rho^*(1,n;\underline{v})^2}{2}}{2}$$
$$-\mu(1-q(\overline{e}))(\mu-z)\frac{(\overline{v}-\underline{v})^2}{\gamma}-(1-\mu)(1-q(\underline{e}))(\mu-z)\frac{(\overline{v}-\underline{v})^2}{\gamma}$$
$$+\mu(z-\mu)\frac{(\overline{v}-\underline{v})^2}{\gamma}+(1-\mu)(z-\mu)\frac{(\overline{v}-\underline{v})^2}{\gamma}=0 \quad (31)$$

From (13), (14), and (15), it should be clear that $q(\overline{e})<q(\underline{e})$. For $q(e)$ below the first best level specified in footnote 22, this implies that surplus is higher from the effort of the $\underline{v}$ agent. Therefore,

$$\frac{[q(\overline{e})+(1-q(\overline{e}))\rho^*(0,c)]-\frac{\gamma\rho^*(0,c)^2}{2}}{2}-\frac{[q(\underline{e})+(1-q(\underline{e}))\rho^*(0,c)]-\frac{\gamma\rho^*(0,c)^2}{2}}{2}$$
$$+\frac{[q(\overline{e})+(1-q(\overline{e}))\rho^*(1,n;\overline{v})]-\frac{\gamma\rho^*(1,n;\overline{v})^2}{2}}{2}-\frac{[q(\underline{e})+(1-q(\underline{e}))\rho^*(1,n;\underline{v})]-\frac{\gamma\rho^*(1,n;\underline{v})^2}{2}}{2}<0. \quad (32)$$

Consequently, the first order condition for equilibrium requires that

$$\mu(1-q(\overline{e}))(\mu-z)\frac{(\overline{v}-\underline{v})^2}{\gamma}+(1-\mu)(1-q(\underline{e}))(\mu-z)\frac{(\overline{v}-\underline{v})^2}{\gamma}-(z-\mu)\frac{(\overline{v}-\underline{v})^2}{\gamma}<0, \quad (33)$$

or $\mu<z$.

Proof of Proposition 2:
If truth-telling does not holds, there exists no equilibrium where the agent can be induced to reveal $\alpha_a=a$. Instead, there are two possible pure strategy Bayesian Nash equilibria of the model: (i) where $a=0$ is always recommended, and where $a=1$ is always recommended. (These equilibria can be supported by offering $w=\Delta q(E(e))$ if the agent deviates.) In these equilibria, the agent is rewarded with a wage $w=\Delta q(E(e))$, where $E(e)$ is the expected level of effort exerted by the agent. But since $w=\Delta q(E(e))$ is independent of true effort, this implies that $e=E(e)=0$.

But an equilibrium where the agent recommends a independent of his information and exerts $e=0$ is at least weakly dominated by internal monitoring where $\mu=z$. In that case, there is an equilibrium where the agent allocates $a=\alpha_a$, and exerts effort $e>0$. If that equiibrium has a payoff which dominates offering $a=0$ or $a=1$ with probability 1, then internal monitoring is strictly better. Otherwise, there exists an equilibrium with internal monitoring where the same strategies are followed as in the external case posited here, and hence the payoffs are equal. Thus, if truth-telling is violated, internal monitoring at least weakly dominates external monitoring.

Proof of Proposition 3:
Consider first the role of internal monitoring.

The Incentive to Complain If the principal chooses p to satisfy (25) and (26), the agent complains both when an error is made by recommending $a=0$ and $a=1$, if complaints increase

oversight, as assumed.

Oversight If the agent always complains when an error has been made, the overseer chooses

$$\rho^{***}(0,c) = \frac{1+(\mu-z)(\overline{v}-\underline{v})}{\gamma}, \tag{34}$$

$$\rho^{***}(1,c) = \frac{1-(\mu-z)(\overline{v}-\underline{v})}{\gamma}, \tag{35}$$

and

$$\rho^{***}(.,n) = 0. \tag{36}$$

Truth-Telling Assumption (5) guarantees truth-telling as in Proposition 1.

Effort Effort is chosen to maximize

$$U_a(v) = \frac{[q(e)+(1-q(e))\rho^{***}(0,c)](v+1)}{2} - (1-q(e))\frac{\gamma\rho^{***}(0,c)^2}{2} + \frac{[q(e)+(1-q(e))\rho^{***}(1,c)](1-v)}{2} - \frac{\gamma\rho^{***}(1,c)^2}{2} + q(e)\Delta - e \tag{37}$$

Maximizing this expression yields the first order condition for the agent's effort

$$\frac{q'(e)(1-\rho^{***}(0,c))(v+1)}{2} + \frac{q'(e)\gamma\rho^{***}(0,c)^2}{2} + \frac{q'(e)(1-\rho^{**}(1,c)(1-v)}{2} + \Delta q'(e) = 1 \tag{38}$$

In order to understand how selection affects incentives, note the effect of v on incentives:

$$\frac{d^2U_a}{dedv} = 0. \tag{39}$$

Thus, unlike the previous model, selection has no effect on incentives. As a result, $\overline{e} = \underline{e}$. But if this is the case, the optimum is where $\mu = z$ as this maximizes oversight propensities. This implies that the best internal monitoring outcome involves $\mu = z$. But this is also the same optimal outcome as with external monitoring (as $\overline{e} = \underline{e}$), and so the best outcome involves $\mu = z$.

Appendix 2: Arrest Rates by Crime

Figure A2.1: Arrests Per Crime: Homicide

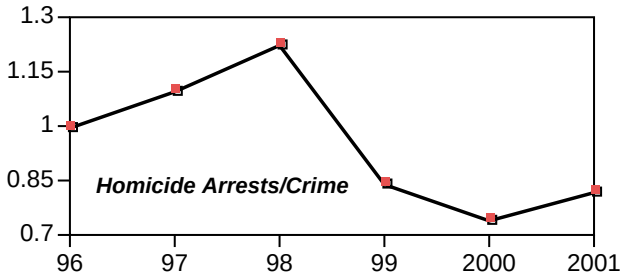

Figure A2.2: Arrests Per Crime: Rape

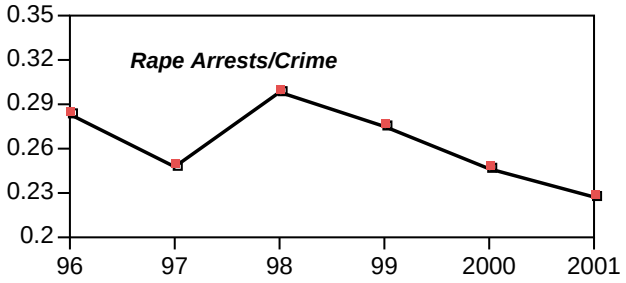

Figure A2.3: Arrests Per Crime: Aggravated Assault

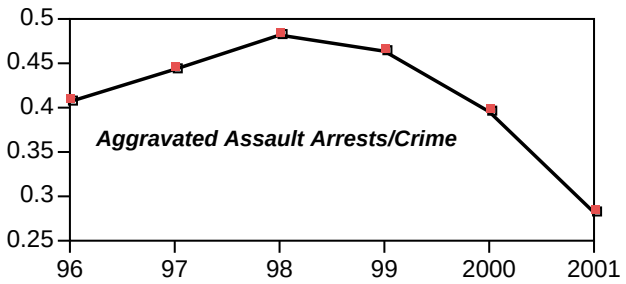

Figure A2.4: Arrests Per Crime: Robbery

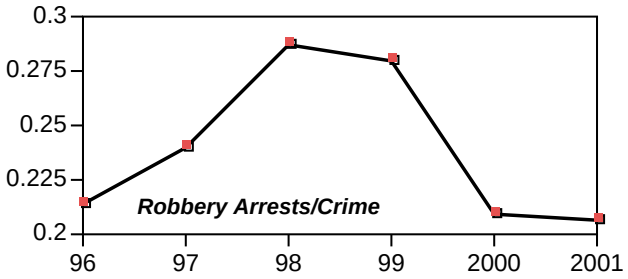

Figure A2.5: Arrests Per Crime: Larceny

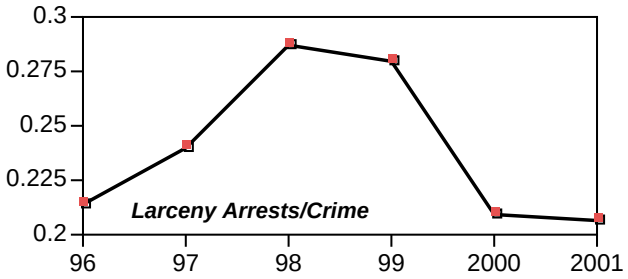

Figure A2.6: Arrests Per Crime: Grand Theft Auto

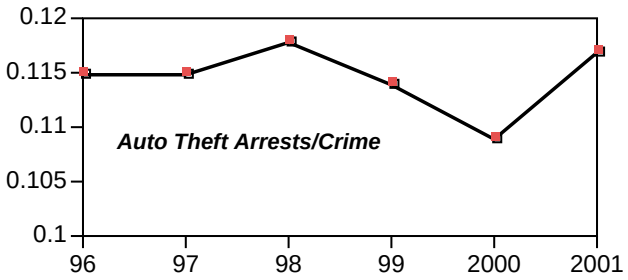

Appendix 3: The Assumption of Unobserved μ

Throughout the paper, I assumed that the agents could infer μ from the actions of the principal, but did not observe it directly. The impact of this assumption is that the principal could not affect effort *for a given person* through the choice of whom to hire. Instead, the effect of hiring policies were simply to overselect those who work harder. But allowing the agent to observe the fraction of $\overline{v}$ workers to hire adds two additional effects.

First, remember that incentives are muted by the fact that in some states (namely, those where the agent gives the consumer the benefit and there is no investigation), there is no information revealed on the agent. This, not surprisingly, reduces incentives. But the principal can change the likelihood of this state, by its hiring practices. In particular, increasing the fraction of sceptical agents (the $\underline{v}$ types) increases oversight propensities in this state (see (11)) and therefore increases effort. Consequently, another reason to increase the fraction of sceptical agents is to make investigations more likely when the consumer does not complain, and so increase incentives.

Second, the premise of the paper is that the prospect of making errors (partly) drives the incentives of the agents. But agents realise that the principal will step in to correct some errors. This then gives the principal an incentive to commit not to correct errors, in the interest of inducing more effort, as the agent's decisions then become more pivotal. But the principal cannot directly commit: instead, oversight is carried out in a subgame perfect fashion. Yet he can commit in an indirect fashion through hiring practices. For example, suppose that the principal chooses to hire only $\overline{v}$ agents. Then the agent realises that oversight is likely very high when the agent is denied the benefit (as $\rho^*(0,c) = \frac{1+(1-z)(\overline{v}-\underline{v})}{\gamma}$), but very low when the agent is given the benefit (as $\rho^*(1,n) = \frac{(1-q)[1+(z-1)(\overline{v}-\underline{v})]}{\gamma}$). This reallocation of oversight propensities has an ambiguous effect on incentives, but could increase them. This implies that the selection of agents can have an commitment-like effect on oversight in such as way as to increase agent incentives. This is the second effect of observing μ on incentives.[49]

[49] I offer these extensions not because of their empirical importance, where agents can observe the preferences of their co-workers, *and* the principal can use this to his advantage in the ways outlines above. Instead, this section simply addresses other effects which would arise with different modeling choices.

References

Aghion, Philippe and Jean Tirole, "Formal and Real Authority in Organizations", *Journal of Political Economy*, 1997, 105, 1-29.

Cannon, Lou, "One Bad Cop", New York Times, October 2, 2000, New York.

Dixit, Avinash, "Incentives and Organizations in the Public Sector: An Interpretive Review", mimeo, 2001.

Gibbons, Robert and Michael Waldman, "Careers in Organizations: Theory and Evidence", forthcoming, *Handbook of Labor Economics*, 1999.

Hall, Sarah, "Protests as Prague Airport Asylum Screening Resumes", *The Guardian*, August 28, 2001.

New York Times, "Besmirched Deportland Wrestles With the INS", August 31, 2000, New York.

New York Times, "Police in Cincinnati Pull Back in Wake of Riots", July 19, 2001, New York.

Malcomson, James, "New Developments in the Study of Contracts in Labor Markets", forthcoming in *Handbook of Labor Economics*, 1999.

Prendergast, Canice, "The Provision of Incentives in Firms", *Journal of Economic Literature*, 1999, 37, 7-63.

Prendergast, Canice, "The Limits of Bureaucratic Efficiency", mimeo, 2001.

Prendergast, Canice, "Consumers and Agency Problems", Dennis Sargan Lecture, Royal Economic Society, 2001: forthcoming, *Economic Journal*.

Star Tribune, "Congressional Cures for the IRS Worse than Agency's Ills", August 30, 2000.

Tharp, Mike, "L.A.'s Go-to Guy Gets Down To Business", *U.S. News and World Report*, 2000, September 11.

Wilson, James, *Bureaucracy*, Basic Books, New York, 1989.