Professionalism and Contracts in Organizations

Canice Prendergast

December 11, 2013

Abstract

Employees in public agencies rarely have pay for performance: instead their incentives are often guided by a sense of professionalism. This paper concerns how organizations should monitor professionals. The primary outcome of the paper is that weak incentives lead public agencies to exhibit bias in their oversight, by rewarding the interests of their employees to the detriment of other constituencies' concerns. In some instances, this bias is complete by entirely ignoring other interests.

University of Chicago Booth School of Business. Email: canice.prendergast@chicagobooth.edu.
A previous version of this paper circulated under the title "Contracts and Conflict in Organizations". Thanks to Mike Waldman, Wouter Dessein, Bob Gibbons, and seminar participants at the University of Chicago, Columbia University, Northwestern University, Harvard/MIT, UCLA, UNC, ESRI, and the NBER for helpful suggestions.

Public agencies typically have two characteristics. First, they have competing interests, where their employees are called upon to satisfy multiple objectives. Second, their employees rarely have pay for performance.[1] Motivating such employees is a central concern of the public administration literature, but has received relatively little attention by economists.[2] A commonly cited motivation of public officials is their sense of professionalism. This paper concerns how firms should monitor professionals. The primary outcome of the paper is that optimal oversight typically involves public agencies showing *bias* in how they trade off objectives, by excessively internalizing the interests of their employees to the detriment of other constituencies' concerns. In some instances, this bias is complete by entirely ignoring other interests. In this sense, agencies appear to be captured by their employees' interests.

There is a small academic literature on professionalism. Two issues are clear from this. First, professionals often exert considerable effort, often based on the importance of the perceptions held by "peers". For example, Wilson, 1989, describes professionals as those who "receive some significant portion of their incentives from organized groups of fellow practitioners" (p.60). Second, while professionals exert effort, their interests are often not aligned with those of their employers. The widely accepted definition of professionalism derives from Hall, 1968. Two of his defining characteristics of professionals are: (i) *self-regulation*: "the belief that the person best qualified to judge the work of a professional is a fellow professional", and (ii) *autonomy*: "the feeling that the professional ought to be able to make his own decisions without external pressure from clients and those who are not members of his profession" (p.92).

This suggests that a problem with professionals is that they only answer to each other rather than to all competing interests. So, for example, social workers value benefit provision to clients over cost control, police officers are easier to motivate to catch "bad guys" than to do community outreach, engineers working on auto safety seek technical solutions to problems rather than driver behavior modification, academics value research over administration, and so on. Accordingly, professionalism is characterized here as the combination of (i) agents having incentives for reasons other than pay for performance, and (ii) incentives biased in

[1]Despite this, many continue to exert considerable effort. Much of this literature is in the field of public administration or political science (such as Goodsell, 1998, and Brehm and Gates, 1997). As an extreme case, note that the US Post Office - hardly known for its extensive use of performance pay - has ontime delivery rates of mail in the region of 98%, and in less than 3% of cases do government officials fail to give enough benefits to welfare recipients (Goodsell, 1998).

[2]Dixit, 2002, offers a review of the existing economic literature and offers many examples of the competing interests of public sector employees.

favor of certain activities.

The baseline model assumes that the incentives of professionals revolve around the desire to "look good" in the eyes of others, akin to the career concerns literature. This arises in a setting where professionals are assumed to be predisposed to favor some activities over others, all else equal.[3] The focus of the paper is how firms respond to such incentives in their oversight: the issue is not why professionals want to focus on certain activities, but rather why their employers encourage them to do so.

Oversight here refers to instruments other than pay for performance that firms use to affect effort exertion. It is assumed throughout that firms can affect how workers allocate their efforts in other ways. First, organizations choose how to allocate rewards. Public sector institutions do have fixed monetary rewards (promotion for example), and how these decisions are made are partly at their discretion, by deciding what gets weighted when these decisions are taken.[4] Second, the choice of who to hire affects what workers do on the job. This is a central theme of Wilson's famous treatise on professionalism. As but one example, he notes that disproportionately hiring engineers at the National Transport Safety Association resulted in psychological approaches to safe driving being ignored. By choice of whom to hire, efforts are affected. Finally, firms allocate resources to tasks: not only dollars, but they also decide what manpower to allocate across tasks. Firms can then affect effort by selectively assigning assets that are complementary to efforts - they can fund certain projects, assign the best workers to particular tasks, and so on.

The second central assumption of the paper is that these alternative sources of incentives are often substitutes, in the sense that an increased focus on one activity implies less focus on another. This arises as many (non-monetary) resources are in limited supply, and those resources needs to be assigned across activities. As an example, consider a control rights example - who gets to decide on promotions. By offering one party more say over whether a candidate is promoted, another party's influence likely declines. The results below arise because of the rivalry between these non-monetary instruments.

In the model below, the mission of the agency involves production on two objectives

[3]It is easy to imagine reasons why professional want to focus on specific activities: often it is what they have been trained to do. For example, social workers are trained to care for their clients, not to save money by denying them benefits. Brehn, 1997, and Kadushin, 1985, describe such orientation for social workers. Other examples of such bias are provided below.

[4]One way of implementing this is through changing control rights over these decisions. If potential evaluators vary in their preferences, then the allocation of these control rights affects what workers are likely to do. See Goldin, 2000, for details of how changing decision makers affected evaluation during the Reagan administration.

or "outputs". These could be service provision and cost control in a welfare department, security and accountability in a police department, and so on. The agency or firm cares equally about each. However, the agent is predisposed to work more on one of these. Given this, how should firms oversee performance? A natural intuition would be that oversight by firms should be focused on the ignored activity, as marginal returns are higher there. So for example if social workers are predisposed to offering benefits to their clients, social welfare departments might respond by emphasizing cost control. However in many cases, it appears that oversight does not seem to counterbalance these interests. Instead, if anything it reinforces these biases. This is manifested in perceptions that many public agencies and not-for-profits have been captured by the interests of their employees, to the detriment of other constituencies. As examples, police departments are often accused of systematically ignoring complaints from the public about police behavior (see News and Record, 2013, as one recent example), academics are routinely believed to work on subjects so arcane that they are of interest only to a narrow set of specialists, and Social Welfare departments are believed to show no interest in cutting costs (Brehm and Gates, 1997). Why institutions may be biased in this way - professional bias - is the principal interest of the paper.

The central outcome of the paper is the efficient use of professional bias, so that downplaying relevant interests may be an efficient ex ante device for providing incentives. I show that in some instances, the optimal strategy is to completely ignore other interests such that professionals only focus on their own narrow interests. This arise because organizations are limited in the incentives that they can provide and so have to prioritize what they have agents focus on: when incentives are weak, they prioritize the activity on which the agent is predisposed to exert effort. So, for example, it may be efficient to ignore the public's complaints about public officials, to only reward academics for their contributions to their narrow fields rather than encourage inter-disciplinary contributions, or to allow social workers to not be responsible for the cost implications of their actions, and so on. This crowding out of other activities is done in order to induce more effort on their professional activities: this is the outcome of the fixed natures of the firm's resources to aid effort exertion.

This outcome ignores the intuitive possibility that oversight be more focused on the ignored activities. As incentives are poorer there, why is this not the focus? I show that professional bias relies on the agent having weak incentives. Specifically, professional bias always occurs when the principal is trading off inefficiently low efforts on both activities. In cases where the agent can be induced to exert excessively high output on one activity, professional bias need not arise, and in the limiting case where close to first best efforts

arise, the principal focuses excessively on the "ignored" activity. Hence weak incentives - in the sense of effort levels being below first best levels - are an important ingredient for professional bias.

Two extensions of the model are addressed. First, the baseline model considers career concerns as the motivation as incentives. However, all that is necessary is that the agent's incentives are biased: career concerns is simply one interpretation of such bias. Second, the basic model simply posits that agents have biased incentives: they are more willing to exert effort on one task rather than the other. This is endogenized below, by showing that there is a reason for firms to specialize in hiring due to increasing returns to scale on the monitoring side. Specifically, performance is easier to monitor when agents are similar. As a result the asymmetry described above can be easily induced from a more symmetric primitive.

The paper addresses the incentives of professionals. The motivating examples are typically public servants. Yet professionals work in many private sector settings where profits are the motive. This is addressed by considering how the ability to use performance pay affects oversight. This is shown to depend not on the efficiency of such contracting *per se* but rather on what can be contracted upon. Intuitively, oversight becomes less biased when aggregate output is contracted upon. This shows how biased oversight is a response to poor contracting. However this is only necessarily true when aggregate output is contracted upon rather than its individual components. Instead, the ability to contract on the components of output reduces bias only if the professional activity can be contracted upon better than the weaker activity. Hence oversight depends on the specifics of contracting.

Professionalism has received little attention in the economics literature. An important exception is Dewatripont, Jewitt, and Tirole (1999). They also address an environment with multiple constituencies where career concerns induce effort exertion. Their results derive from an assumption that career concerns incentives become diluted when agents are "spread too thin".[5] Because of this, agencies impose "missions", where the institution bans the agent from carrying out certain activities. By doing so, total effort rises. This paper also predicts an outcome of public agencies is excessive focus. This does not arise through multiple activities weakening the informativeness of incentive signals. Instead, here agencies simply do not have enough incentives ammunition to satisfy all constituencies and must prioritize which activities on which to focus the agent. The central outcome is that with weak incentives, agents excessively focus on their professional interests. This gets reversed as the willingness

[5]Specifically, when an agent engages in many activities, it is harder to infer ability from performance measures, and so effort falls. This is either because more tasks involve more aggregate noise, or because the speed of learning increases in task-specific effort.

to exert effort in aggregate gets larger. These insights continue to hold when the agent's motivation is other than career concerns. How one might empirically distinguish the insights of this work from those of Dewatripont et al. is discussed below.

Section 1 begins by describing the model. Section 2 shows how biased oversight arises as the optimal way of providing incentives and also shows the importance of weak incentives. Section 3 considers examples. Following this, Section 4 illustrates how contracting on aggregate performance mutes this tendency towards bias, while a more nuanced outcome arises when contracts can be written on individual performance. Extensions are considered in Section 5, which show that the insights are generally robust to other assumptions.

1 The Model

A risk neutral firm employs a risk neutral agent to produce two outputs, A and B. Each output depends on the agent's ability and effort on that task. Output in activity i is given by

$$y_i = m_i + e_i + \epsilon_i, \tag{1}$$

where the agent has ability m_i at activity i, exerts effort e_i, and ϵ_i is a noise term, $i = A, B$.[6] Outputs are observed but cannot be contracted upon. Effort is assumed to be both costly and unobserved by the principal. The cost of effort is given by $C(e_A, e_B)$, where $C_i > 0, C_{ii} > 0$, and $C_{ij} \geq 0$, for $i \neq j$.

There is symmetric uncertainty about the agent's abilities. The belief of both actors is that m_i is Normally distributed with mean μ, and variance σ_m^2. These are common across the two activities in order to focus on agency concerns. The distribution of ϵ_i is Normal with mean 0 and variance $\sigma_i^2, i = A, B$, and the noise terms are uncorrelated with each other and the m_i. In order to reflect the priority of a "profession", there is an asymmetry between the two activities. In keeping with Hall's definition above where peers are the best monitors, it is assumed that performance in one activity is more easily measured: it is assumed that activity A is the agent's profession, where $\sigma_A^2 < \sigma_B^2$. (This asymmetry is exogenous, but an endogenous interpretation is given below.) The interpretation here is that these variances refer to the quality of information garnered by fellow professionals, who also generate the rewards for the agent.

[6] This technology allows the possibility of output to be negative. This will also be true for wages below. This should not be taken literally, but rather these technologies represent local approximations to address small enough changes in effort or ability.

The objective of the firm is to maximize expected "profits", the sum of the two outputs minus wages. The agent cares about income and effort: her utility from income I and efforts e_i is given by $U(I,e) = I - C(e_A, e_B)$. She derives income from two sources. First, she earns a wage, w, which in the initial part of the paper is a salary. Second, she gains a benefit from perceptions of her expected abilities. This is couched in terms of career concerns below. The agent has a reservation utility normalized to 0 so accepts the position if $EU(I,e) \geq 0$. (The role of the market is addressed below.)

Career Concerns A by now large literature points to the willingness of agents to exert effort to affect others' perceptions of their abilities. While the agent begins with expected abilities μ, observations of output will change this perception, which affects future compensation. Accordingly, after observing the vector of outputs, $\mathbf{y}$, the agent's expected ability is updated to $E(m_i|\mathbf{y}), i = A, B$.

The agent has capacities at two activities, A and B, and after observing $\mathbf{y}$, these are updated to $E(m_A|\mathbf{y})$ and $E(m_B|\mathbf{y})$.[7] As is standard in career concerns settings, worker welfare is increasing in her perceived ability (Holmstrom, 1999, Gibbons and Murphy, 1992). The novel assumption of the model is that the firm chooses how these are weighted in assigning benefits to the worker. There are a number of ways this can occur, which are elaborated upon below. Here we provide an interpretation of how benefits are assigned, where the principal decides whether to retain the worker after observing her performance. It makes this decision based on her expected abilities. Retaining her job has an exogenous benefit to the worker normalized to \$1. However, the principal chooses weights between 0 and 1 on each ability in the retention decision. I assume that (locally) the probability of being retained is $pE(m_A|\mathbf{y}) + (1-p)E(m_B|\mathbf{y})$, where p is chosen by the firm.[8] Hence higher ability increases the probability of retention, but the principal chooses the importance of each.

The agent's expected utility is therefore given by

$$EU(I,e) = E_{\mathbf{y}}[w + pE(m_A|\mathbf{y}) + (1-p)E(m_B|\mathbf{y}) - C(e_A, e_B)]. \tag{2}$$

Note that p is the only instrument that the principal holds to affect effort exertion.

[7] As examples, a university professor has research and teaching abilities, a police officer is valued both for his ability to solve crime and to handle the public, a social worker has an aptitude for both service provision and cost control, and so on.

[8] These are interpreted here as probabilities, so we assume that this lies between 0 and 1 to ignore the case where career concerns are irrelevant. The appropriate interpretation is that these are local approximations. Below I consider the case where the weights need not sum to 1.

The timing of the model is as follows. First, the firm offers a contract to the agent, where the contract is (w, p). Hence the firm can commit to oversight. The agent either accepts or rejects the contract. If the contract is rejected, the game ends. If she accepts the contract, she exerts efforts e_A and e_B. After efforts are exerted, outputs are observed, and the agent receives $w+pE(m_A|\mathbf{y})+(1-p)E(m_B|\mathbf{y})-C(e_A, e_B)$, while the principal receives y_A+y_B-w.

2 Biased Oversight

The firm chooses oversight to maximize profits. By varying p appropriately, it can induce effort exertion on either activity. The objective of this section is to illustrate where the firm locates on this spectrum.

To begin, consider the agent's incentives. She exerts effort to affect (out-of-equilibrium) ability perceptions. Career incentives are standard here. As the error terms are independent, each output y_i is sufficient for updating expected ability m_i. Standard updating implies that

$$E[m_i|\mathbf{y}] = s_i[y_i - E(e_i)] + (1 - s_i)\mu \tag{3}$$

where $s_i = \frac{\sigma_m^2}{\sigma_m^2+\sigma_i^2}$, and $E(e_i)$ is the market's expectation of the agent's effort. The value of marginally increasing e_i on expected ability is given by s_i. As activity A can be more easily monitored, the agent's incentives are greater in her professional activity. This is the source of biased incentives in the model. To see this, note that the agent chooses e_i to maximize (2) subject to (3) and so

$$ps_A = C_A(e_A, e_B) \tag{4}$$

and

$$(1 - p)s_B = C_B(e_A, e_B). \tag{5}$$

Intuitively, effort exertion depends both on the informativeness of output and how much the firm weights each ability in its evaluation. Importantly, oversight and informativeness are complements. Note that as $s_i \leq 1$, the first best ($C_i = 1$) is not possible, and efforts are inefficiently low.

The focus of this section is the exercise of professional bias, by which it is meant that oversight favors activity A. The definition used here is the natural *ex post* benchmark - as the principal cares about both outputs equally, the "passive" benchmark is where both are treated equally.[9] As a result, I consider oversight to be biased in favor of the profession if

[9]So as an example, if a scandal broke in the newspaper where it appeared one activity was being ignored,

$p > \frac{1}{2}$.[10] By doing so, the principal will not simply passively react to the agent's incentives to exert more effort on A, nor set oversight to counteract it, but reinforces it to induce a further specialization of their efforts towards the preferred interest of the agent. As a further (more extreme) benchmark, I also consider the case where the principal sets oversight to completely ignore activity B by choosing $p = 1$.

The objective of the firm is to maximize profits subject to the incentives of the agent. As the agent's participation constraint binds in equilibrium, this is equivalent to the principal choosing p^* to maximize surplus $y_A + y_B - C(e_A, e_B)$ subject to the agent's incentives.[11]

Optimal oversight depends on two factors: (i) how different are effort responses to oversight for the two activities, and (ii) how different is surplus from any extra effort? To see this, it is useful to consider the principal's first order condition when setting oversight. Differentiating surplus above with respect to p yields the first order condition

$$\frac{de_A}{dp}(1 - C_A(e_A, e_B)) \geq -\frac{de_B}{dp}(1 - C_B(e_A, e_B)). \tag{6}$$

where $\frac{de_i}{dp}$ are determined by (4) and (5), with the inequality strict if p is strictly between 0 and 1. Professional bias arises if the solution to this implies an optimum $p^* > \frac{1}{2}$.

This condition illustrates the relevant tradeoff. For the career concerns reasons above, effort is more responsive on the margin for A than on B. This leads to a tendency for professional bias. However, as the agent is predisposed to exert more effort on A than B, residual surplus is higher on B, which would lead to oversight being biased in the opposite direction.[12] The results below are a tradeoff between these two conflicting influences.

In order to focus on the role of professionals having greater incentives at one activity, I initially consider the case where the only difference between the two activities is the ability to monitor performance, by assuming that effort costs are are independent, symmetric, and

how would the firm arbitrate this scandal? As they care about each equally, the natural benchmark is to not favor one constituency over the others.

[10]There is another potentially useful definition of bias which is marginal surplus is not equalized across the two activities. Marginal surplus on activity i is given by $1 - C_i$. One sense of bias is where oversight is chosen such that $1 - C_A < 1 - C_B$. This weak form of bias will always arise below.

[11]The participation constraint of the agent only depends on p through effort decisions as $Em_A = Em_B$.

[12]Two benchmarks are useful here. First, if both efforts respond equally to oversight ($\frac{de_A}{dp} = -\frac{de_B}{dp}$) marginal surplus is equated. Hence the only possibility for distorted effort is as a means of providing incentives when there is an asymmetry between the two activities, as seems appropriate for professional settings. Second, if the marginal costs of effort are identical, such as when $C(e_A, e_B) = C(e_A + e_B)$, the only thing that matters for oversight is which effort is most responsive on the margin. These extreme cases illustrate the two influences.

quadratic: $C(e_A, e_B) = \frac{\gamma e_A^2}{2} + \frac{\gamma e_B^2}{2}$. Unless otherwise stated, the Propositions below relate to this case. More general cost functions are considered below to illustrate the generality of the results.

Proposition 1 *Assume that effort costs are symmetric and quadratic:* $C(e_A, e_B) = \frac{\gamma e_A^2}{2} + \frac{\gamma e_B^2}{2}$. *Then the firm always exhibits professional bias* ($p^* > \frac{1}{2}$)*:*

- *If* $s_B > s_A(1 - s_A)$*, then* $p^* = \frac{s_A - s_B + s_B^2}{s_A^2 + s_B^2} > \frac{1}{2}$.
- *If* $s_B \leq s_A(1 - s_A)$*, then* $p^* = 1$ *and the agent exerts no effort on activity* B.

This is the central idea of the paper: when agents are predisposed to ignore an activity, the principal reinforces that bias through its oversight. Hence the actions of the firm serve to reinforce certain interests being ignored: not only are agent's professional biases allowed, they are encouraged. To see this, note that unbiased oversight would induce relative effort of $\frac{e_A}{e_B} = \frac{s_A}{s_B}$. From Proposition 1 equilibrium relative efforts are more distorted than this because of biased oversight.

If the ability to provide incentives on B is sufficiently poor (s_B low enough), then effort is only exerted at activity A so the agent is efficiently induced to ignore certain constituencies. So for example, a police office is shielded from being accountable to the public in order to induce him to focus on catching "bad guys", or a social worker is allowed to ignore cost control in order to better induce effort on service provision to clients. This arises because if constituency A is given more weight, by necessity it reduces the influence of constituency B - this is where the rivalry of incentives becomes key.

Proposition 1 also offers a useful comparative static. Activities in public sector settings are often amorphous and hard to observe. As a result, it is useful to consider how outcomes change as it becomes more difficult to monitor performance. Consider the case where activity B is harder to monitor. As s_B declines, bias increases, so bias is reinforced with poor monitoring of the other activity.[13] Furthermore, when these activities become sufficiently hard to monitor, the firm simply gives up trying to induce the agent to exert any effort on task B. In this sense, designing oversight to make agents unaccountable to certain constituencies may be a natural outcome of public agencies.

[13] $\frac{dp^*}{ds_B} = \frac{-s_A^2 + s_B^2 - 2s_A s_B(1 - s_A)}{(s_A^2 + s_B^2)^2} < 0$

2.1 More General Cost Functions

So far, I have considered the case where the cost functions for the two activities are independent with similar quadratic costs. As a result, effort responses differ between the two activities only because of different s_i. Now consider more general cost functions, $C(e_A, e_B)$. Note first that by substituting (4) and (5) into (6) and evaluating at $p = \frac{1}{2}$, the principal exhibits bias for any cost function if

$$\frac{de_A}{dp}(2 - s_A) > \frac{de_B}{dp}(2 - s_B). \tag{7}$$

As $s_A > s_B$, this implies that a sufficient condition for professional bias is that effort on A is more responsive to oversight changes than is effort on B. The results on more general cost functions below simply make that more precise, by showing conditions under which this is so.

Proposition 2 *Assume that the agent has a cost function $C(e_A, e_B)$. The firm chooses $p^* > \frac{1}{2}$ if*

$$\frac{s_A(2 - s_A)}{C_{AA} + C_{AB}} > \frac{s_B(2 - s_B)}{C_{BB} + C_{AB}}. \tag{8}$$

Activity B is ignored if $\frac{s_A(1-s_A)}{C_{AA}+C_{AB}} > \frac{s_B}{C_{BB}+C_{AB}}$.

In the previous section, effort on A is more responsive if $s_A > s_B$. With more general cost functions, this is true if $\frac{s_A}{C_{AA}+C_{AB}} > \frac{s_B}{C_{BB}+C_{AB}}$. The economic logic is as above: bias arises if the responsiveness of effort to oversight is greater for activity A. All that differs from the model above is that it now also depends on the marginal cost of efforts. This formulation allows us to consider natural extensions to the cost function used above:

Cost Interactions: Consider the quadratic interaction case: $C(e_A, e_B) = \frac{\gamma e_A^2}{2} + \frac{\gamma e_B^2}{2} + \delta e_A e_B$, where $\delta < \gamma$. Then (8) simplifies to $s_A(2 - s_A) > s_B(2 - s_B)$, which always holds.

Aggregate Effort Costs: Consider the case where $C(e_A, e_B) = C(e_A + e_B)$. Then $p^* = 1$ as surplus is identical for the two activities.

Task Specific Costs: Consider the quadratic case where the marginal cost curve is activity specific: $C(e_A, e_B) = \frac{\gamma_A e_A^2}{2} + \frac{\gamma_B e_B^2}{2}$, where $\gamma_A \neq \gamma_B$.[14] In this case, activity A is more responsive to oversight than activity B if $\frac{s_A}{\gamma_A} > \frac{s_B}{\gamma_B}$, which is a sufficient condition for bias.[15] The logic

[14] A natural assumption here would be that marginal costs are lower for the agent's profession: $\gamma_A < \gamma_B$.

[15] Bias occurs if $\frac{s_A(2-s_A)}{\gamma_A} > \frac{s_B(2-s_B)}{\gamma_B}$.

is identical to above - oversight is focused on the activity whose marginal effort response is highest.

Symmetric Cost Functions: Consider the case the cost function is symmetric so $C(a,b) = C(b,a)$. Then while (8) can fail when $s_A > s_B$, this relies on $\frac{s_A}{C_{AA}} < \frac{s_B}{C_{BB}}$, which can only occur if third derivatives of the cost function are sufficiently large as the relative C_{ii} only depend on the level of efforts.

As a result of these cases, the insights above have more general validity than the special case of Proposition 1: when the worker's objective is given by (2), the firm exhibits bias towards the activity that is most responsive to effort. For the cost functions considered above, this will be the activity in which they are predisposed to exerting effort, activity A.

To summarize, public agencies typically cannot use incentive pay to orient actions. Instead, they rely on workers' sense of professionalism, yet this sense of professionalism typically results in distorted incentives across competing interests. While the agency could allocate resources and decision rights to counterbalance that incentive, they choose not to do so. Instead, they reinforce the distorted incentives of their employees, with the implication that some competing interests are downplayed, and in some cases, completely ignored.

2.2 Weak Incentives

An important assumption above is that incentives are weak relative to first best efforts. Remember that the agent maximizes $w + pE(m_A|\mathbf{y}) + (1-p)E(m_B|\mathbf{y}) - C(e_A, e_B)$. This formulation rules out the possibility that the agent exerts more than efficient effort, which would arise if marginal incentives exceed unity. Hence effort is interior to the optimum and the principal is trading off inefficiently low efforts on both activities. To understand the importance of this assumption, we now allow for career concerns to be stronger, by no longer assuming that weights on ability need to sum to 1. Specifically, assume now that the agent's utility is

$$w + \kappa(pE(m_A|\mathbf{y}) + (1-p)E(m_B|\mathbf{y})) - C(e_A, e_B) \tag{9}$$

where $\kappa > 1$. Hence career concerns have been increased in importance. (This could be for the standard reason offered in Holmstrom, 1999, where workers garner a good reputation long into the future, or it could simply be that agents value greatly the impressions that others have of their skills.)

This shift in the agent's preferences is non-trivial for the oversight decisions of the firm because as the agent's choices get closer to the efficient levels, optimal oversight changes.

First, for κ large enough, it should be clear that the principal can induce first best effort levels.[16] Assume that first best efforts are not possible. In this case, optimal bias is given by Proposition 3.

Proposition 3 *Assume that the agent maximizes (9) and the first best is not feasible. Then the firm exhibits professional bias only if* $\kappa < \frac{2}{s_A+s_B}$. *If* $\kappa > \frac{2}{s_A+s_B}$, *oversight is biased on favor of* B.

This proposition reflects an additional intuitive possibility for oversight, namely that as the agent has enough incentives to carry out one activity, oversight is focused on emphasizing the other one. With κ large enough, the principal changes his focus from the professional activity to the activity with less effort. But note that this relies on the possibility of the agent exerting excessively high effort levels for some oversight.

More generally, this section illustrates an important implication of the model: professional bias is a response to poor incentives. In this case $\frac{dp^*}{d\kappa} = \frac{s_B - s_A}{\kappa^2} < 0$, so bias is a direct response to weak incentives. Remember from above that optimal oversight trades off two factors - differences in responsiveness of effort and differences in marginal surplus. When incentives are stronger, the value of additional effort on the professional activity becomes small (as effort is closer to the first best) and so oversight moves towards supporting the weaker activity. Hence these insights apply best to cases where incentive provision is difficult.

2.3 A Broader Interpretation

The asymmetry between the agent's willingness to exert effort is characterized by $\frac{s_i}{C_i}$ from (4) and (5). The source of differences in the s_i has been interpreted though through the lens of the career concerns logic of the informativeness of signals. Yet it does not rely on this interpretation. Instead, *any* asymmetry in willingness to exert effort between tasks generates professional bias. Specifically, the results arise when $\frac{de_A}{dp} > -\frac{de_B}{dp}$ with weak incentives.

Consider two other plausible alternative sources of bias in professionals. First, agents may have intrinsic preferences to exert effort, where those preferences differ between activities. For example, following Brehm et al, 1997, social workers likely feel an innate willingness to exert effort to help clients that exceeds their inherent desire to control costs. From this

[16] In the previous sections, it always is optimal to maximize oversight in aggregate by choosing weights that sum to 1. With more incentives, this is not longer necessarily true. The principal now chooses oversight p_A and p_B subject to $p_A + p_B \leq 1$. If the agent has enough incentives to exert effort, the first best is possible, by choosing $p_i = \frac{1}{\kappa s_i}$. The first best is then possible if $\kappa \geq \frac{1}{s_A} + \frac{1}{s_B}$.

perspective, reinterpret $\frac{s_A}{C_A}$ as the agent's private marginal return to service provision and $\frac{s_B}{C_B}$ their marginal interest in controlling costs. Consider the case where the agent cares equally about both outputs, with weight s, but where one activity is more unpleasant to exert (marginal) effort on than the other, where $C_{BB} > C_{AA}$. Then $\frac{de_A}{dp} = \frac{s}{C_{AA}+C_{AB}} > -\frac{de_B}{dp} = \frac{s}{C_{BB}+C_{AB}}$ and the results above follow identicallyNote also that from this perspective, the outside market plays no role other than the principal guaranteeing the agent her reservation utility.

The case of intrinsic preferences raises another relevant instrument reminiscent of Holmstrom and Tirole (1991). When there are interactions in the cost function $C_{AB} > 0$, increased incentives on say activity A will reduce effort on B. With intrinsic preferences to exert effort, it may be optimal to shut down the responsiveness of pay to output on A to induce more effort on B. A good example of this (in settings outside formal pay for performance) is Carmichael (1988) who argues for the role of tenure among academics as a way of inducing them to not worry so much about their own individual performance ("A"). In Carmichael, this is so as to induce them to select better colleagues (who otherwise might threaten their jobs) but the logic would easily extend to other "B" activities.

This reinterpretation of the model does not rely on the preferences for exerting effort being intrinsic. Instead they could be weak monetary incentives for example, say where the agent is attempting to be promoted to a higher paying position in the organization. Let s_i now be the marginal monetary return to exerting effort on activity i in terms of promotion returns, as in Lazear and Rosen. So, for example, s_A could exceed s_B as it is more easily monitored. Then the agent will be biased towards A if $\frac{s_A}{C_{AA}+C_{AB}} > \frac{s_B}{C_{BB}+C_{AB}}$ in exactly the same way as above. As a result, the insights depend only on the agent having some reason to prefer exerting effort on one activity over another.

3 Instruments and Examples

The model relies on the assumption that firms have instruments to affect incentives aside from the usual measure considered, incentive pay. A number of such instruments are potentially feasible:

Control Rights In the model above, it was assumed that the principal placed weights p and $1-p$ on a retention decision. A natural way in which this holds in firms is through control rights. Many decisions made in firms are based on control rights. Somebody or some

group gets to choose outcomes in the absence of contractibility. An important decision in many settings is who gets to choose. This underlies much of the literature on incomplete contracting following Grossman and Hart, 1986. Yet control rights are often zero sum: if I get to choose what to do, you don't. A natural interpretation of p above is that certain constituencies get to decide how to treat the agent. This is modeled here in a continuous way rather like a marginal voter model, where oversight biased in favor of A makes it likely that A is the marginal voter over the agent's prospects.[17]

Such an extension could be formally introduced to the model. Assume that there is a benefit that the agent wishes to attain, which has a value of \$1 if achieved. Oversight for that decision is randomly assigned: with probability p those who value her talent at A control the decision over the receipt of that benefit, and with probability $1-p$ those who value her ability at B choose. The agent perceives the probability of receiving that prize as linear in her expected ability at the relevant task, where relevance is determined by who chooses, and her probability in that state is $\tau E(m_i|y)$. Then her incentives are identical to above (modulo the factor τ) and the qualitative results continue to hold.

Hiring An alternative way to interpret p is through hiring decisions. Agencies choose who to hire for positions. Potential candidates vary in how they are likely to spend their time as they vary in how they trade off various objectives. There is considerable evidence that staffing decisions directly affect how jobs are done. Take the example elaborated on in Pruitt, 1979, on the National Transportation Safety Agency, whose mandate was to encourage safe driving. He notes that because the NTSA chose to almost exclusively hire engineers (rather than psychologists), the agency focused on technical solutions to car safety (such as airbags) to the detriment of the kind of behavioral approaches to safe driving that psychologists would more likely have focused on.

Another clear example of this is Wilson's description of the contrast between the Forestry Service and the Park Service. Both have similar mission - involving both the need to preserve the wilderness but also to make it open to the public - but they carry them out very differently. The Forestry Service typically hires those with degrees in Forestry and is often accused of focusing on the preservation aspects of their job over making land accessible to

[17] A natural example arises from our own profession. Universities vary in how they evaluate candidates for tenure. Beyond the obvious tradeoff between research and teaching, some universities have more intervention at the levels of Deans and Provosts who have been know to seek such characteristics as "real world applications", or "interdisciplinary contributions". By allocating control rights to different parties, junior faculty interested in attaining tenure at that university will likely change their behavior.

the public. By contrast, the Park Service typically hires engineers and law enforcement officers, and is more focused on its public service component. While Wilson is more interested in how such staffing decisions make change difficult, it is clear that staffing decisions affect orientation in the job. Wilson (p.62) also argues that a focus on hiring engineers at NASA may have contributed to the Challenger disaster through an excessive focus on numbers.

As a final example, social workers are trained to provide services to their clients, and their sense of professionalism typically revolves around the provision of high quality services. Yet social workers are additionally charged with the task of controlling costs, by denying services to those deemed (often by others) to not be sufficiently needy. Their professional incentives do not rewards them for such cost control activities, as much of their training is associated with helping others and Brehm and Gates, 1997, document this resistance. In a similar vein, Derthick, 1979, also provides some evidence on such conflicts for social workers when they were asked by the Social Service Administration to be instrumental in denying coverage to applicants. Their widespread refusal to act on this mandate illustrates the type of bias described here. In a survey on the preferences of social workers, Robert Peabody, 1964, notes that "by far the most dominant organizational goal perceived as important..is service to clientele" (p.66), where 83 percent of survey respondents view such service as important, compared to only 9 percent who see "obligation to taxpayers" or "assistance to the public in general" as important concerns affecting their decisions.[18]

Complementary Investments Firms assign assets to aid workers in their tasks - other employees to help, capital equipment, research budgets, and so on. To the extent that these assignments are activity specific (such as funding particular research) and these budgets are fixed, the results above naturally arise. Below, I consider a case where the firm has a budget of \$1 and has to choose how to allocate that dollar to each task to affect its productivity. The assets are complementary in that they make output more likely. (Think of a project that might never come to fruition without sufficient resources.) Take a very a simple case - an allocation of z cents to an activity implies that output in that activity arises with probability z and zero output is observed with probability $1 - z$. The firms then has to choose how to allocate its dollar to maximize productivity, and the issue is whether the firm biases its

[18]From Section 2.3, all that is necessary is that different hires have different motivations. However, there are some examples where career concerns appear to lie behind such motivations. Take Wilson's description of the Federal Trade Commission (p.60), which hires both lawyers and economists to pursue anti-trust cases. He describes how the actions of lawyers are explicitly oriented towards achieving a high paying law job in the private sector, while the objectives of economists reflect their own professional orientation.

complementary investments towards the agent's profession. This extension is discussed in more detail below.

The central point of the paper is that when the firm holds few instruments to induce incentives, it has to choose where to focus, and how it focuses depends on the weakness of those instruments. This final interpretation relies on limited supply of these complementary assets: ideally one would simply assign more assets to each activity. However, to the extent that public agencies are rationed in their ability to assign such assets, the same influences as above arise.

4 Contracting

This paper addresses the behavior of professionals. The examples above are mostly from the public sector. Yet professionals pervade most industries. For example, almost all Fortune 500 companies likely hire engineers, yet all the examples above are for cases where the profit motive is absent. In this section, I extend the analysis to consider a case where outputs can be contracted upon. While formal contracting may be feasible in some parts of the public sector (as in say the large recent increase in incentive pay for public school teachers), many of these cases arise in the private sector. As such, this section addresses incentives for professionals more generally.

Two possible issues of interest arise here. First, how does contracting affect the relative efforts of the two activities? First consider the case where contracts are based on a measure of aggregate performance. Not surprisingly, the ability to contract moves relative efforts towards the weaker activity as contracting improves. Second, how does the ability to contract affects governance on the non-contractual issues? For example, how are promotion decisions made for engineers in companies that can also contract on performance? It is shown that the ability to contract reduces the demand for professional bias when aggregate performance is contracted upon. When individual performance measures are feasible, it depends on which measure can be better contracted upon. In the case where each activity can be equally well contracted upon, bias remains unchanged. Hence how governance is affected by the ability to contract depends on precisely what can be contracted upon.

4.1 Contracting on Aggregate Output

Suppose that the principal can now contract on an unbiased measure of total output. However, the measure that can be contracted upon potentially differs from true output. Following

Baker, 1992, it is assumed that the firm has a distorted measure of aggregate output to contract on. Specifically if $y = y_A + y_B$ is true aggregate output, assume that the measure on which contracts are feasible is given by $\tilde{y} = \omega y$, where ω can be observed by all but cannot be contracted upon. ω is assumed to be Normally distributed with mean 1 and variance σ_ω^2. Higher values of σ_ω^2 imply more distorted performance measures.

The agent is no longer paid a salary but instead is offered a wage $w = \beta_0 + \beta\tilde{y}$, so the contract offered to the agent is now (β_0, β, p). Assume that the principal does not observe ω until after outputs are observed so that contracts cannot be conditioned even implicitly on their realization. The firm has two instruments that affect efforts, pay for performance and oversight. Let β^{**} and p^{**} be the optimal contract.

Proposition 4 *The ability to contract on aggregate output reduces professional bias, $p^{**} < p^*$. Optimal bias is given by*

$$p^{**} = min\{\frac{[s_A - s_B](1-\beta^{**}) + s_B^2}{s_A^2 + s_B^2}, 1\}, \tag{10}$$

where $\beta^{**} = \frac{2 - p^{**}s_A - (1-p^{**})s_B}{2(1+\sigma_\omega^2)}$.

The focus of this paper is the idea that favoring professional interests may be optimal. The examples given above all arise from the not-for-profit world. The result above provides a simple reason for this - when output can be contracted upon, the value of professional bias declines.

4.2 Contracting on Individual Outputs

Now assume that the firm can contract on individual outputs. This matters here as the firm can now target incentives on particular activities. The measure on which contracts for output i are feasible is given by $\tilde{y}_i = \omega_i y$, where ω_i can be observed but cannot be contracted upon. ω_i is now assumed to be Normally distributed with mean 1 and variance $\sigma_{\omega i}^2$. The agent is no longer paid a salary but instead is offered a contract $w = \beta_0 + \beta_A\tilde{y_A} + \beta_B\tilde{y_B}$. The contract offered is now $(\beta_0, \beta_A, \beta_B, p)$. Let p^{***} be the optimal contract. Bias is given by Proposition 5.

Proposition 5 *The ability to contract on individual outputs only reduces professional bias if activity A can be more efficiently contracted upon ($\sigma_{\omega A}^2 < \sigma_{\omega B}^2$). If each can be contracted on equally well, bias is unchanged. Optimal bias is given by*

$$p^{***} = min\{\frac{Rs_A - s_B + s_B^2}{Rs_A^2 + s_B^2}, 1\}, \tag{11}$$

where $R = \frac{\rho_1}{\rho_2}$ *and* $\rho_i = \frac{r\sigma^2_{\omega i}}{1+r\sigma^2_{\omega i}}, i = A, B.$

This offers a caveat to the idea that the favoritism shown to professional interests necessarily declines with the ability to contract. A natural benchmark is where the two can be contracted upon equally well: here bias is unchanged by any ability to contract. The reason is simply that the *relative* demand for more effort on B rather than A is unchanged by contracting (each declines by the same proportion ρ_i above) and so oversight is unchanged. Hence how professional oversight changes depends on the specifics of what can be contracted upon.

One contrast is notable between the two contracting cases above: when individual outputs can be contracted upon, the first best is possible as contracting converges to perfect efficiency, while bias still arises in the limit when aggregate output is contracted upon. The reason for this can be seen by noting the asymmetry between the two career incentives: without any contractual incentives, the agent exerts more effort on A than B. Now consider the case where contracting is possible on each component of output. Then to get to the first best (where total incentives are 1) more contractual incentives are necessary on B than A. This is possible when each output can be separately contracted upon. But this is not possible when aggregate output is contracted upon, as incentives in A are $s_A + \beta$ and incentives in B are $s_B + \beta$. As $s_A > s_B$ these cannot both be simultaneously 1 and so distortions arise even in the limit.

5 Extensions

5.1 Markets

So far, the market has played little role in the analysis. In cases where the source of the bias is not related to career concerns, this is not an issue, as the firm simply chooses wages to meet reservation utilities in Section 2.3. However, when perceptions of ability are the source of incentives, the market may play a larger role on the margin. One interpretation of these results is that either the information or the abilities are firm specific: while they affect rewards in the firm, they play little role in how other employers value workers. One case where information is firm specific is when only peer professionals in the firm observe performance. If outsiders see no implications of this information, there is little reasons for outside options to change. Alternatively, it may be that whatever skills workers have in firm has little correlation for their prospects outside. As an example, social workers may gain

prestige among their co-workers and superiors by handling their case loads well, it is not clear it has much effect on their outside options, either because that information is not well known or because social work skills do not translate well elsewhere. In these settings, the results above continue to hold.

Yet there are many cases where ability may be general. In this section, I embed the results above in a more standard market setting. The model above has a worker with two abilities. Furthermore, there is no market operating in the background. A simple reinterpretation shows that both of these assumptions can be relaxed. Consider the case where there is a single ability, $m_A = m_B = m$, which adapts the agent's objective function in (2) to

$$EU(I, e) = E_{\mathbf{y}}[w + E(m|\mathbf{y}) - C(e_A, e_B)]. \tag{12}$$

This is analogous to the objective function in Holmstrom, 1982, where future compensation depends linearly on the market's perception of ability. Assume that this is the agent's market based compensation.

The action that the principal takes is the allocation of complementary assets. These assets affect whether the activity is a success or not. Specifically, output in each activity is either a success, in which case y_i is as in (1), or a failure, in which case output is 0. Assets are necessary for a success: by allocating z of the asset to an activity, output is a success with probability z. The total supply of this asset is 1, and they are assigned and observed before the effort decision. Here the firm orients actions through the probability that each activity will be a success, and reveals information on ability.

Consider the agent's incentives with assets of p assigned to activity A. There is a probability p^2 that "only A" is a success, $(1-p)^2$ that "only B" is a success, and $p(1-p)$ that "both A and B" are a success. The agent's incentives are then to maximize

$$EU(I, e) = E_{\mathbf{y}}[w + p^2 E(m|y_A) + (1-p)^2 E(m|y_B) + p(1-p)E(m|y_A, y_B) - C(e_A, e_B)]. \tag{13}$$

The impact of oversight on incentives is qualitatively identical to the baseline one above. Incentives in the model above were determined by s_i. The only complication here is that if both activities are a success, expected ability is given by $E(m|y_A, y_B) = s_0^*\mu + s_A^*(y_A - Ee_A) + s_B^*(y_B - Ee_B)$, where $s_i^* = \frac{\sigma_m^2\sigma_j^2}{\sigma_m^2\sigma_A^2+\sigma_m^2\sigma_B^2+\sigma_A^2\sigma_B^2}$, $j \neq i$, and $s_0^* = \frac{\sigma_A^2\sigma_B^2}{\sigma_m^2\sigma_A^2+\sigma_m^2\sigma_B^2+\sigma_A^2\sigma_B^2}$. Then efforts are given by

$$p^2 s_A + p(1-p)s_A^* = C_A(e_A, e_B) \tag{14}$$

and

$$(1-p)^2 s_B + p(1-p)s_B^* = C_B(e_A, e_B), \tag{15}$$

The same influences arise to distort oversight from $p = \frac{1}{2}$. At $p = \frac{1}{2}$, increasing p increases effort on A by $\frac{s_A}{C_{AA}+C_{AB}}$ and decreases it on B by $\frac{s_B}{C_{BB}+C_{AB}}$, and the logic for professional bias is as above. This reinterpretation allows us to adapt the model to one with a market determined wage and one-dimensional ability, as is more standard in the literature.

5.2 Endogenous Asymmetry

So far, the asymmetry between the two tasks has been exogenously assumed. In this section, the asymmetry is endogenized through increasing returns to scale in monitoring. The economic intuition is simple: the performance of agents can be more easily measured when "similar" agents are employed by using relative performance evaluation. As a result, an institution that values two endeavors equally may choose to specialize in one, in order to improve incentives, even though it causes the other activity to be ignored. The extension in this section formalizes this.

I make two changes to the model above. First, the firms now hires two agents, 1 and 2, where the agents are either trained in profession A or profession B. Let j refer to the agent's identity, $j = 1, 2$, and i to his profession, $i = A, B$. Output produced is as in (1) but I now assume that $\sigma_A^2 = \sigma_B^2$, so the two activities are ex ante identical.

The only change to the technology above is that the noise in the agent's performance in his given profession is composed of a profession specific component and a person specific component. Training typically involves a particular way of carrying out tasks. Accordingly, assume that if the agent has profession i his performance in activity i is given by

$$\epsilon_{ij} = \zeta_i + \hat{\epsilon}_{ij}. \tag{16}$$

The distribution of ζ_i is assumed to be Normal with mean 0 and variance $\sigma_\zeta^2, i = A, B,$, the distribution of $\hat{\epsilon}_{ij}$ is Normal with mean 0 and variance $\sigma_{\hat{\epsilon}}^2, i = A, B$, and the noise terms are uncorrelated with all others. The error term for the other activity is unchanged from above.[19] To be consistent with the previous section, $\sigma_\epsilon^2 = \sigma_\zeta^2 + \sigma_{\hat{\epsilon}}^2$. As a result, there is no difference in the primitives of each activity, so we are beginning from a case where both activities are symmetric.

As the activities are ex ante identical the only relevant choice for the principal is whether to hire both of the same type (say type A) or hire one worker of each type.

[19] Allowing for a similarity in how they do task B would strengthen the tendency to specialize.

Proposition 6 *The firm hires two agents of profession A, as effort on task A is higher than hiring one agent of each profession, and effort on task B is unchanged. Effort on task A increases by* $\frac{p(\overline{s}_A - s_A)}{\gamma} > 0$*, where*

$$\overline{s}_A = \frac{\sigma_m^2}{\sigma_m^2 + \sigma_\epsilon^2 + \hat{\sigma}_\zeta^2} \tag{17}$$

and $\hat{\sigma}_\zeta^2 = \frac{\sigma_\zeta^2(\sigma_m^2+\sigma_{\hat{\epsilon}}^2)}{\sigma_\zeta^2+\sigma_m^2+\sigma_{\hat{\epsilon}}^2} < \sigma_\zeta^2$.

Hence it may be efficient to specialize when hiring public agents as a means of improving monitoring. This arises because another agent who is similar allows more noise to be filtered out, and so incentives improve. By contrast, without any such benchmark, it is harder to attribute performance. But specialization induces an asymmetry in monitoring. Of course, the influence can be overturned if there is decreasing returns to scale from two agents involved in the same activity. However, for small enough diminishing returns, the optimal strategy for the firm is to specialize in one profession.

As an example, a central theme of Wilson's description of government agency is how hiring decisions affect their actions. NASA and the NTSA chooses to hire engineers, the Forestry Service focuses on those with Forestry degrees, and so on. What is striking in these examples is the extent to which these agencies chose not to diversify their expertise, but instead become focused on a narrow set of experts. As a result, even agencies with very similar missions (the Forestry Service versus the Park Service) can often have very different ways of satisfying their missions. While there are many ways in which this can come about, this section interprets this in terms of agency concerns, where the ability to monitor exhibits increasing returns to scale for a profession.

5.3 Specialization

So far, one source of asymmetry has been emphasized, namely that one activity is easier to monitor than the other. However, an alternative natural asymmetry is where one task is more commonly done than the other. Here the interpretation of bias is towards the common activity. To see this, consider an alternative where outputs are given by y_A and y_B as before but the agent exerts unobservable efforts on two tasks, 1 and 2. Let effort on task i be given by e_i, $i = 1, 2$, with cost $C(e_1, e_2) = \frac{\gamma e_1^2}{2} + \frac{\gamma e_2^2}{2}$. All effort by the agent on task 1 increases the returns solely of output A. By contrast, effort on task 2 has a shared benefit. A fraction x of effort on task 2 benefits output A, while the remaining $(1 - x)$ benefits output B, where

$0 \leq x < \frac{1}{2}$.[20] Hence $y_A = e_1 + xe_2 + m_A + \epsilon_A$ and $y_B = (1-x)e_2 + m_B + \epsilon_B$. Tasks are therefore partially specialized, as reflected by the parameter x. To focus solely on the role of specialization, assume that $\sigma_A^2 = \sigma_B^2$.

Proposition 7 *Consider the case where tasks are specialized ($x > 0$), but $s_A = s_B$. Then (i) bias ($p^* > \frac{1}{2}$) arises for all $x > 0$, and (ii) bias is (weakly) increasing in x.*

Hence oversight is biased if tasks are specialized. Once again, the outcomes here do not rely on the specifics of performance being better monitored in A: what matters is that the agent has *some* reason to focus on one activity over the other.

6 Conclusion

The focus of agency theory has largely been on pay for performance as a motivator. Yet most workers do not have any formal pay for performance (Lemieux, McLeod, and Parent, 2009). The assumption motivating this work is that firms continue to have instruments to motivate workers even without formal pay for performance. These can occur in the ex ante way described here by assigning resources or control rights. Alternatively, this could occur ex post, such as where infractions by public official come to light in the media and the agency has to decide whether to defend the official or not. This paper would suggest that it may be optimal to come to their defense even when other constituencies are clearly harmed.

The reason for biased oversight here is to induce effort exertion - effort on one task is sacrificed to increase effort on the other. This is reminiscent of two previous contributions. First, in Holmstrom and Milgrom's (1991) work on multitasking, firms scale back pay for performance on some activities in order to increase effort on others. This arises because of interactions on the cost side for workers. Here the focus is on substitutability on the oversight side. Perhaps more fundamentally, pay for performance is not relevant in many settings where professionalism acts as an incentive, and so is silent on the issues of interest here. This paper offers an interpretation of why some incentives are "crowded out" in settings with no incentive pay.

Second, as described above, Dewatripont et al also provide an economic theory of professionalism. The environments that they study - those with the kind of multiple principals described in Dixit, 2002 - are similar to those considered here. What matters in each paper is multiple activities in a setting where formal incentive provision is largely absent. In both

[20] It is assumed that $x < \frac{1}{2}$ so that effort on task 2 primarily benefit task B.

of these cases, the optimal outcome is a form of focus, where efforts are concentrated on a subset of the possible activities. While the setting and outcomes are somewhat similar, the reasons are very different. Outcomes here arise because organizations are limited in the incentives that they can provide and so have to prioritize what they have agents focus on. By contrast, Dewatripont, et al. provide a theory of "missions" caused by it being harder to infer ability from performance signals when agents do more tasks.

Although the outcomes are similar in both papers, three possible ways of empirically distinguishing between the two papers come to mind:

1. This paper relies on an asymmetry between the two activities. When it is equally easy to induce effort on activities, this paper does not predict bias. By contrast, Dewatripont generate focus with symmetric tasks. If focus empirically arises in symmetric settings, this would argue for Dewatripont et al over this work.

2. This paper does not require career concerns to be the source of professional incentives for bias to arise. This was shown in Section 2.3. By contrast, the insights of Dewatripont et al., rely on the informativeness of signals as the central ingredient. If focus arises in settings where motivation arises from other sources, this would argue in favor of this work.

3. A final source of differentiation may arise through the strength of incentives. My interest is largely on *which activities* organizations choose to focus upon. The implication of this model (from Section 2.2) is that when incentives are weak, the firm induces focus on those activities that are easier to monitor, while when incentives become stronger, they reverse and excessively focus on the activities that are harder to monitor. This could conceivably be used as a way of empirical differentiation.

The idea of biased oversight is also similar in outcome to Prendergast, 2003. In that work, oversight of bureaucrats is biased against consumers. The reason for this is that if consumer complaints are responded to with sufficient frequency, bureaucrats acquiesce to their demands, even in settings where they should be denied. Once again, although the outcomes are reminiscent, the mechanisms are very different.[21]

[21] Also loosely related are Che and Kartik, 2009, and Van Den Steen, 2007, 2010, who study the value of biased agents. Similarly, Prendergast, 2007, offers another reason for the benefits of bias, based on agents' altruism towards their clients. There is also a literature showing how monetary contracts that are biased towards some activities can improve efficiency, such as Itoh, 1992 , Dessein, Garicano, and Gertner, 2008, and Rotemberg and Saloner, 1995. See also MacLeod, 2003, for other work on the costs of conflict in settings with subjective performance measures.

To summarize, this paper argues that even without formal monetary incentives, organizations can still affect incentives in a variety of ways. While this may seem uncontroversial, there is little in the literature that guides how these should be used, other than Dewatripont et al's contribution on banning tasks. Here the choices are more continuous, and point to the role of emphasizing dominant activities when incentives are weak. This takes the form here of - either partially or completely - ignoring activities which are weaker to provide incentives. Yet this is a response to weak incentives, as stronger incentives can result in oversight focusing more on those ignored activities. Finally, by considering the role of contracting, this work also places the role of such bias firmly in the realm of public agencies, where such contracting is difficult.

References

Brehm, John, and Scott Gates. 1997. Working, shirking, and sabotage: bureaucratic responses to a democratic public. University of Michigan Press, Ann Arbor.

Che, Yeon-Koo and Navin Kartik. 2009. Opinions as incentives. *Journal of Political Economy* 117, no.5: 815-860.

Dessein, Wouter, Luis Garicano and Robert Gertner. 2008. Organizing for synergies: a theory of hybrids. Unpublished manuscript. Chicago GSB.

Derthick, Martha. 1979. *Policymaking for social security*, Brookings Institute, Washington.

Dewatripont, Matthias, Ian Jewitt, and Jean Tirole. 1999. Missions and the accountability of government agencies, *Review of Economic Studies* 66, no. 1: 199-217.

Dixit, Avinah. 2002. Incentives and organizations in the public sector: an interpretative review, *Journal of Human Resources* 37, no 4: 696-727.

Golden, Marissa. 2000. *What motivates bureaucrats?*. Columbia University Press, New York.

Goodsell, Charles. 2004. *The case for bureaucracy*, CQ Press, Washington.

Gibbons, Robert and Kevin M. Murphy. 1992. Optimal incentive contracts in the presence of career concerns. *Journal of Political Economy* 100, No. 3: 468-505 .

Grossman, Sanford and Oliver Hart . 1986.The costs and benefits of ownership: a theory of vertical and lateral organization. *Journal of Political Economy* 94: 691-719.

Hall, Richard. 1968. Professionalization and bureaucratization. *American Sociological Review* 70, no. 2: 92-104.

Holmstrom, Bengt. 1999. Managerial incentive problems: a dynamic perspective. *Review of Economic Studies* 66: 169-82.

Holmstrom, Bengt and Paul Milgrom 1992. Multi-task principal agent

analyses: linear contracts, asset ownership and job design. *Journal of Law, Economics, and Organization* 7:24-52.

Kadushin, Alfred. 1985. *Supervision in social work*. Columbia University Press, New York.

Itoh, Hideshi. 1992. Cooperation in hierarchical organizations: an incentives perspective. *Journal of Law, Economics, and Organizations* 8, no.2: 321-45,

Lazear, Edward and Sherwin Rosen. 1981. Rank order tournaments as optimal labor contracts. *Journal of Political Economy*, 89: 841-64.

Lemieux, T., MacLeod, Bentley, W. and Parent, D. 2009. Performance pay and wage inequality, *Quarterly Journal of Economics* 124, no. 1:149.

Milgrom, Paul, and John Roberts. 1988. Employment contracts, influence activities, and efficient organizational design, *Journal of Political Economy* 96, no. 1: 42-60.

MacLeod, Bentley. 2003. Optimal contracting with subjective evaluation. *American Economic Review* 93, no. 1: 216-40.

News and Record 2013. Some criticize police response to complaints, June 13.

Peabody, Robert. 1964. *Organizational authority*, Atherton Press, New York.

Prendergast, Canice. 2003. The limits of bureaucratic efficiency. *Journal of Political Economy* 111, no. 5: 929-958.

Prendergast, Canice, 2007. The motivation and bias of bureaucrats. *American Economic Review* 97, no. 1: 180-96.

Pruitt, Charles. 1979. People doing what they do best: the professional engineers and NHTSA, *Public Administration Review* 39, no. 1: 363-71.

Rotemberg, Julio and Garth Saloner, 1995. Overt interfunctional conflict. *RAND Journal of Economics* 26, no.4: 630-53.

Van Den Steen, Eric. 2007. The cost of incentives under disagreement: too motivated. Unpublished manuscript. Harvard business school.

Van Den Steen, Eric. 2010. Disagreement and the allocation of control. *Journal of Law, Economics, and Organization* 26, no.2: 385-426.

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

Proof of Proposition 1: First note that the principal chooses w to satisfy the agent's participation constraint and so $w + pE(\mu_A) + (1-p)E\mu_B - C(e_A, e_B) = 0$, and so $w = -\mu + C(e_A, e_B)$, as $E(\mu_i) = \mu$. As a result, the principal maximizes $y_A + y_B - w = y_A + y_B + \mu - C(e_A, e_B)$, subject to the agent's incentives. As $E[m_i|\mathbf{y}] = s_i[y_i - E(e_i)] + (1-s_i)\mu$, this implies that $ps_A = C_A(e_A, e_B)$ and $(1-p)s_B = C_B(e_A, e_B)$. Differentiating surplus with respect to p yields

$$\frac{de_A}{dp}(1 - C_A(e_A, e_B)) + \frac{de_B}{dp}(1 - C_B(e_A, e_B)) \geq 0. \tag{18}$$

Let the optimal choice be p^*. From (4) and (5), this simplifies to

$$\frac{de_A}{dp}(1 - p^* s_A) \geq -\frac{de_B}{dp}(1 - (1-p^*)s_B). \tag{19}$$

with the equality strict if $0 < p^* < 1$. Let ρ be the relative responsiveness of effort on A to effort on B when oversight is changed, $\rho = -\frac{\frac{de_A}{dp}}{\frac{de_B}{dp}}$. Then the first order condition is given by $\rho \geq \frac{1-(1-p^*)s_B}{1-p^* s_A}$. Then consider two natural cases - where bias arises ($p^* \geq \frac{1}{2}$) and where the weaker activity is ignored ($p^* = 1$). Bias arises if

$$\rho > \frac{2 - s_B}{2 - s_A}. \tag{20}$$

Furthermore, activity B is ignored if $\rho > \frac{1}{1-s_A}$.

With general cost functions, $\frac{de_A}{dp} = \frac{C_{BB}s_A - C_{AB}s_B}{C_{AA}C_{BB} - C_{AB}^2}$ and $\frac{de_A}{dp} = -\frac{C_{AA}s_B - C_{AB}s_A}{C_{AA}C_{BB} - C_{AB}^2}$. Then consider the case where $C(e_A, e_B) = \frac{\gamma e_A^2}{2} + \frac{\gamma e_B^2}{2}$. Then $\rho = \frac{s_A}{s_B}$ and bias always holds if $\frac{s_A}{s_B} > \frac{2-s_B}{2-s_A}$, which is always true. Furthermore, substituting for $\rho = \frac{s_A}{s_B}$ in the first order condition and solving for p^* yields $p^* = min\{\frac{s_A - s_B + s_B^2}{s_A^2 + s_B^2}, 1\} > \frac{1}{2}$ as required. Finally, $p^* = 1$ if $s_B < s_A(1 - s_A)$.

Proof of Proposition2: With more general cost functions, $\frac{de_A}{dp} = \frac{C_{BB}s_A - C_{AB}s_B}{C_{AA}C_{BB} - C_{AB}^2}$ and $\frac{de_A}{dp} = -\frac{C_{AA}s_B - C_{AB}s_A}{C_{AA}C_{BB} - C_{AB}^2}$. Substitution for $p = \frac{1}{2}$ yields (8) while substituting for $p = 1$ yields $\frac{s_A}{C_{AA} + C_{AB}} > \frac{s_B(1-s_B)}{C_{BB} + C_{AB}}$.

Cost Interactions: When $C(e_A, e_B) = \frac{\gamma e_A^2}{2} + \frac{\gamma e_B^2}{2} + \delta e_A e_B$, where $\delta < \gamma$, the agent's choice of efforts are given by $ps_A = \gamma e_A + \delta e_B$ and $(1-p)s_B = \gamma e_B + \delta e_A$. Then $e_A = \frac{ps_A - \frac{\delta}{\gamma}(1-p)s_B}{1 + (\frac{\delta}{\gamma})^2}$ and $e_B = \frac{(1-p)s_B - \frac{\delta}{\gamma}ps_A}{1 + (\frac{\delta}{\gamma})^2}$. Then $\rho = \frac{s_A - \frac{\delta}{\gamma}s_B}{s_B - \frac{\delta}{\gamma}s_A}$ and substitution yields $p^* = min\{\frac{(s_A - s_B)(1 + \frac{\delta}{\gamma}) + (s_B - \frac{\delta}{\gamma}s_A)s_B}{s_A^2 + s_B^2 - 2\frac{\delta}{\gamma}s_A s_B}, 1\} > \frac{1}{2}$.

Task Specific Marginal Costs: When $C(e_A, e_B) = \frac{\gamma_A e_A^2}{2} + \frac{\gamma_B e_B^2}{2}$, the agent's choice of efforts are given by $ps_A = \gamma_A e_A$ and $(1-p)s_B = \gamma_B e_B + \delta e_A$. Then $\rho = \frac{\frac{s_A}{\gamma_A}}{\frac{s_B}{\gamma_B}}$ and substitution yields $p^* = min\{\frac{\frac{s_A}{\gamma_A} - \frac{s_B}{\gamma_B} + \frac{s_B^2}{\gamma_B}}{\frac{s_A^2}{\gamma_A} + \frac{s_B^2}{\gamma_B}}, 1\} > \frac{1}{2}$.

Proof of Proposition 4: The firm observes ω when making inferences of ability and so $E[m_i|\mathbf{y}] = s_i[\frac{y_i - E(e_i)}{\omega}] + (1 - s_i)\mu$, and so career incentives are unchanged. The firm now chooses p and β to maximize expected surplus $E[y_A + y_B - \frac{\gamma e_A^2}{2} - \frac{\gamma e_B^2}{2}]$, subject to $ps_A + \omega\beta = \gamma e_A$ and $(1-p)s_B + \omega\beta = \gamma e_B$. Let the optimal choice be p^*. This is equivalent to maximizing

$$E[2\omega\beta + s_A p + s_B(1-p) - \frac{(\omega\beta + s_A p)^2}{2} - \frac{(\omega\beta + s_B(1-p))^2}{2}]. \tag{21}$$

Taking expectations, and noting that $E[(\omega\beta)^2] = E[(\omega\beta)]^2 + Var(\omega\beta)$, this simplifies to maximizing

$$2\beta + s_A p + s_B(1-p) - \frac{(\beta + s_A p)^2}{2} - \frac{(\beta + s_B(1-p))^2}{2} - \beta^2\sigma_\omega^2. \tag{22}$$

Differentiating this with respect to β yields

$$\beta^{**} = \frac{2 - p^{**}s_A - (1-p^{**})s_B}{2(1+\sigma_\omega^2)}, \tag{23}$$

and

$$p^{**} = min\{\frac{[s_A - s_B](1-\beta^{**}) + s_B^2}{s_A^2 + s_B^2}, 1\}. \tag{24}$$

Substituting for β^{**} in (24) yields

$$p^{**} = min\{\frac{[s_A - s_B](1 - \frac{2 - p^{**}s_A - (1-p^{**})s_B}{2(1+\sigma_\omega^2)}) + s_B^2}{s_A^2 + s_B^2}, 1\}. \tag{25}$$

In the case where p^{**} is not at the boundary, totally differentiating yields

$$\frac{dp^{**}}{d\sigma_\omega^2} = \frac{\frac{[s_A - s_B](2 - p^{**}s_A - (1-p^{**})s_B)}{(1+\sigma_\omega^2)}}{2(1+\sigma_\omega^2)(s_A^2 + s_B^2) - (s_A - s_B)^2} > 0, \tag{26}$$

as required.

Proof of Proposition 5: The firm observes ω_i when making inferences of ability and so $E[m_i|\mathbf{y}] = s_i[\frac{y_i - E(e_i)}{\omega_i}] + (1 - s_i)\mu$, and so career incentives are unchanged. In a natural extension of the previous section, the agent exerts effort $\gamma e_A = ps_A + \omega_A\beta_A$, and $\gamma e_B =$

$(1-p)s_B + \omega_B\beta_B$. The firm now chooses p and β_i to maximize expected surplus $E[y_A + y_B - \frac{\gamma e_A^2}{2} - \frac{\gamma e_B^2}{2}]$, subject to $ps_A + \omega_A\beta_A = \gamma e_A$ and $(1-p)s_B + \omega_B\beta_B = \gamma e_B$. Let the optimal choice be p^{***}. This is equivalent to maximizing

$$E[\omega_A\beta_A + \omega_B\beta_B + s_Ap + s_B(1-p) - \frac{(\omega_A\beta_A + s_Ap)^2}{2} - \frac{(\omega_B\beta_B + s_B(1-p))^2}{2}]. \quad (27)$$

Taking expectations, this simplifies to maximizing

$$2\beta + s_Ap + s_B(1-p) - \frac{(\beta + s_Ap)^2}{2} - \frac{(\beta + s_B(1-p))^2}{2} - \frac{\beta_A^2\sigma_{\omega_A}^2}{2} - \frac{\beta_B^2\sigma_{\omega_B}^2}{2}. \quad (28)$$

Differentiating this with respect to β_i yields $\beta_A^{***} = \frac{1-p^{***}s_A}{1+\sigma_{\omega_A}^2}$, and $\beta_B^{***} = \frac{1-(1-p^{***})s_B}{1+\sigma_{\omega_B}^2}$, while differentiating with respect to bias yields

$$p^{***} = min\{\frac{\frac{\sigma_{\omega_A}^2}{1+\sigma_{\omega_A}^2}s_A - \frac{\sigma_{\omega_B}^2}{1+\sigma_{\omega_B}^2}s_B + \frac{\sigma_{\omega_B}^2}{1+\sigma_{\omega_B}^2}s_B^2}{\frac{\sigma_{\omega_A}^2}{1+\sigma_{\omega_A}^2}s_A^2 + \frac{\sigma_{\omega_B}^2}{1+\sigma_{\omega_B}^2}s_B^2}, 1\}. \quad (29)$$

Then if $R = \frac{\frac{\sigma_{\omega_A}^2}{1+\sigma_{\omega_A}^2}}{\frac{\sigma_{\omega_B}^2}{1+\sigma_{\omega_B}^2}}$, this simplifies to $p^{***} = min\{\frac{Rs_A - s_B + s_B^2}{Rs_A^2 + s_B^2}, 1\}$, as required. Contracting on individual output then only affects bias through R. If $R > (<1)1$ bias increases (decreases) with contracting, while if each activity can be equally well contracted upon $R = 1$ and the outcome is unchanged.

Proof of Proposition 6: With with no agents or a single agent, the best estimate of an agent's performance has signal to noise ratio in her profession $\underline{s}_A = \frac{\sigma_m^2}{\sigma_m^2 + \sigma_\epsilon^2 + \sigma_\zeta^2}$. But as $\sigma_\epsilon^2 = \sigma_\zeta^2 + \sigma_{\hat{\epsilon}}^2$, this implies that $s_A = \underline{s}_A$ and so $ps_A = e_A$. As result, the agent chooses $\frac{ps_A}{\gamma} = e_A$ and $\frac{(1-p)s_B}{\gamma} = e_B$. When two agents are in the same profession, the firm can obtain a more precise estimate of ζ_i by using the other agent's performance (this is analogous to relative performance evaluation). Using one signal (say agent 1's output) allows the prior on ζ to be updated to $\hat{\zeta} = \frac{\sigma_m^2}{\sigma_\zeta^2 + \sigma_m^2 + \sigma_{\hat{\epsilon}}^2}[y_{1A} - \mu - E(e_{1A})]$, where $E(e_{1A})$ is expected effort by agent 1 on activity A, with variance $\hat{\sigma}_\zeta^2 = \frac{\sigma_\zeta^2(\sigma_m^2 + \sigma_{\hat{\epsilon}}^2)}{\sigma_\zeta^2 + \sigma_m^2 + \sigma_{\hat{\epsilon}}^2}$. Using this, the signal to noise ratio for profession A that profession becomes $\overline{s}_A = \frac{\sigma_m^2}{\sigma_m^2 + \sigma_\epsilon^2 + \hat{\sigma}_\zeta^2}$ where $\hat{\sigma}_\zeta^2 < \sigma_\zeta^2$. Note that $\overline{s} > \underline{s}$ for $\sigma_\zeta^2 < \infty$. As a result, incentives are greater with two agents in the same profession, with $p\overline{s}_A = \gamma e_A$, with e_B unchanged. The same logic is true for agents so both agent's efforts rise by $\frac{p(\overline{s}_A - s_A)}{\gamma}$,

Proof of Proposition 7: The agent now chooses $\gamma e_1 = ps$ and $\gamma e_2 = pxs + (1-p)(1-x)s$. The firm chooses p and to maximize $E[e_1 + e_2 - \frac{\gamma e_1^2}{2} - \frac{\gamma e_2^2}{2}]$, subject to $\gamma e_1 = ps$ and $\gamma e_2 = pxs + (1-p)(1-x)s$. Differentiation yields $p^* = min\{\frac{2x-(2x-1)(1-x)s}{s[1+(2x-1)^2]}, 1\}$. If $x = 0$, then $p = \frac{1}{2}$. As p^* is increasing in x and decreasing in s, the Proposition follows.