6 Incentives, Risk, and Accountability in Organizations

Timothy Besley and Maitreesh Ghatak

This chapter discusses incentives and risk management in organizations and how economists have thought about these issues. Risk is not merely the possibility of 'bad things happening', such as damage, loss, and accidents. Any deviation from expected outcomes constitutes risk, whether it is positive or negative. Clearly no organization can hope to achieve its objectives without putting its resources at some amount of risk. An excessive focus on risk-minimization leads to foregone opportunities. The question is striking the right balance between the risk that an organization is exposed to and achieving its objectives (e.g. shareholder value if it is a for-profit firm). A direct implication of this is that risk management and fulfilling the objectives of the firm are interrelated decision problems and cannot be studied in isolation from one another (see, for example, Poynter 2004).

An organization can mitigate risk by buying insurance, diversifying its portfolio, and maintaining sufficient solvency. However, it is impossible to eliminate all risk. If this was possible so that the earnings of the organization were invariant irrespective of its performance, this will also eliminate the drive to perform well. Therefore inevitably the problem of risk comes back to understanding how actors in an organization make key decisions that affect risk exposure. Variability of returns needs to be interpreted by managers and investors to determine whether their strategies are working. Here, therefore, we will focus on aspects of risk that can be altered by decisions made

within the organization even though the outcomes of such outcomes may in part be influenced by events outside of the organization.

Risk management in organizations whether public or private requires that incentives of those who work within an organization be appropriately aligned. The traditional economic approach to risk management stresses how explicit financial incentive schemes can play a role in making sure that members of an organization seek common goals. This is the cornerstone of the principal-agent approach to organization design. The key idea in such approaches is the importance of the fact that many actions taken in organizations are not observable. In this chapter, we review and assess this approach to risk management and its limitations. We then explore approaches that go beyond the standard principal-agent framework.

The standard economic approach to risk management has in mind a profit-maximizing firm with agents who are primarily motivated by financial concerns. The narrowness of this perspective in focusing only on money as a motivator has been questioned in the management and organizational behaviour literature (see for example, Herzberg 1987; Kerr 1995; Kohn 1993), and more recently, in economics literature (see Frey 1997). And for organizations that pursue non-pecuniary goals – such as government bureaucracies, public service providers and non-profits, the limitations of focusing only on financial incentive schemes is widely agreed upon (See Dixit 2002).

Starting with the same premise that the exclusive focus on financial incentives is misguided, we offer a different perspective that we characterize as the 'three M's' approach – missions, motivation and matching. We view individuals as being motivated not just by money, but by missions. However, only when matched with the right organization, and within the organization, the right task and project, is this intrinsic motivation fully realized. Therefore, matching

mitigates some of the potential principal-agent problems at the entry level by reducing the need to use explicit financial incentives, or supervision. This approach is able to meld the classical economic approach with non-economic approaches to organizations. It has rich implications for changing accountability structures and regulatory regimes.[1]

The chapter is organized as follows. We begin by laying out the classical principal agent model as well as discussing some more sophisticated variants of it. We then discuss the three M's approach followed by a discussion about how it can be applied to offer new perspectives on various aspects of risk management in organizations.

The classical principal-agent model

The economic view of an organization is that it is a network of principal-agent relationships. The shareholder-CEO relationship, the CEO-divisional manager relationship, the divisional manager-employee relationship are all examples of principal-agent relationships. An individual who is a principal in one relationship can be an agent in another relationship. The resulting network of relationships characterizes the accountability structure of an organization, i.e. the chain of command and control.

In the classical principal-agent model there are two parties, a principal and an agent. Each principal has some authority over the agent, in that he or she can determine what tasks the agent has to perform, and shape his or her compensation package. However the principal's power over the agent is limited by the presence of alternative employers who could hire the agent. In other words, the

[1] While our analysis is relevant to issues of risk management in an organization, the basic notion that organizations have individuals within them that make decisions that are not observable to their principals is of much wider applicability. It arises even when these decisions do not necessarily expose the organization to any significant amount of risk that needs to be managed.

overall payoff of the agent must be at least as high as what he or she could earn under a different principal (less the cost of moving and finding another job).

The agent works for the principal and undertakes some task that the principal cannot undertake, either because it requires some special expertise, or because the principal has a limited amount of time. The task could be in the nature of things that the agent has to do, such as working 'hard' or choosing the 'best' project. Alternatively, it could be in the nature of reporting to the principal some piece of information that the principal needs to make some decisions, such as the true cost or true value of a project, bad news or good news. These actions affect the payoff of the principal and so the principal has a direct stake in inducing the agent to take actions that are most desirable from his or her point of view. However, the agent is likely to have better information than the principal as to what actions he or she undertakes or what information he or she possesses. That is, there is

likely to be asymmetric information between the principal and the agent, with the agent having an informational advantage. Given this, the agent is likely to undertake the action that is most convenient from his or her private point of view. Similarly, he or she is likely to report information that is most favourable from his or her private point of view. These may not coincide with what is most desirable from the principal's point of view. In the literature the former possibility is described as the problem of 'hidden action' or 'moral hazard' and the latter possibility is described as 'hidden information' or 'adverse selection'. For the most part, our focus in this chapter will be on the former type of problem.[2]

[2] The terms 'moral hazard' and 'adverse selection' originated in the insurance industry. Purchasers of insurance are viewed as having a struggle with their conscience as to whether to do the right thing when it's tempting to do otherwise (e.g. not take enough precaution to prevent a loss). Similarly, those who select (i.e. choose) to buy insurance are the ones who are most likely to have an accident, and are therefore 'adverse' or undesirable from an insurance firm's point of view.

The principal-agent model assumes that the principal and the agent have objectives that are not fully aligned and that the actions undertaken by the agent cannot be perfectly monitored by the principal. If both these features are present, then the principal-agent problem has bite, i.e. there is an 'agency problem'. For example, the principal may care about the profits of the organization (e.g. if he is the owner) whereas the agent may care only about his or her own remuneration and how hard or easy the alternatives are from his or her private point of view (e.g. if he or she is the manager). Similarly, the principal may have only a 'noisy' measure of what the agent actually does or whether his or her reports are truthful or not. This 'noise' arises from the difficulty of perfectly monitoring or supervising the agent. Typically this is because the outcome of the relevant task is uncertain since it depends on both exogenous factors and the action of the agent.[3] For example, sales could be low because of some general factors affecting most firms in the industry, some specific factors affecting the firm that are beyond a manager's control, such as delays by a supplier, or because the manager was not putting in enough effort.

If the principal has access to a good monitoring or supervision technology, then he or she can either directly observe the agent's action or can infer it from the outcome by filtering out the effect of exogenous risk. In this case there is no agency problem and the agent can be induced to take the desired actions from the principal's point of view by making it part of the job description. Alternatively, if the agent is completely loyal to the principal, then too there is no agency problem – the agent can be trusted to do take the desired actions from the principal's point of view. The

---

[3] It is possible that the outcome of the agent's action is not subject to any uncertainty, but the measurement of it is noisy from the principal's point of view, say, because the supervisors can make mistakes.

combination of these two elements raises the possibility that the agent could be undertaking actions that are not in the best interests of the principal, and yet, there is no way the principal can directly catch him or her.

In the presence of moral hazard, the principal has to use indirect means to induce the agent to take the desired actions from the principal's point of view. In other words, the agent must be given incentives to perform.

Foundations of agency problems

Is it possible to design incentives such that the outcome is the same as when there are no differences in objectives between the principal and the agent, or when the principal can perfectly monitor the agent? In other words, can contractual methods achieve the 'first-best' situation, i.e.

as if there is either no asymmetry of information, or perfect alignment of objectives of the principal and the agent? The answer, as one would expect, is typically no. There are costs associated with providing incentives. The source of the problem lies in the difficulty of making the agent's objective closely aligned to that of the principal via an incentive scheme. The obvious way to do this is through rewards and punishments. However, if the agent is risk averse or there is a 'limited liability' constraint that limits how much the agent can be punished or fined in the event of the

performance measure being unsatisfactory, there will be a loss of efficiency due to the agency problem. Let us consider these by turn.

Suppose the agent is neutral towards risk. Technically, this means that if he or she faces an uncertain income stream, he or she cares only about the mean or expected income and not the potential variability of income around this mean. In this case, the principal can design incentives

that will achieve the first-best, i.e. it will be as if the principal could perfectly monitor the agent. The trick is to make the agent what economists call a 'full residual claimant'. For example, suppose the outcome can be either satisfactory or unsatisfactory. If the agent works hard, it will be satisfactory with probability 0.75 and unsatisfactory with probability 0.25. If the agent does not work hard, the probabilities are 50:50. Working hard is considered a more costly option to the agent, since he or she can presumably store his or her time and energy and devote them to more pleasurable pursuits. Because of exogenous uncertainty, just by observing unsatisfactory performance the principal cannot conclude that the agent did not work hard. The principal can however tell the agent that he or she will be paid a salary only if the outcome is satisfactory, otherwise he or she will not get paid. Faced with this reward scheme the agent is very likely to work hard.[4] Examples of such incentive schemes are piece rates and fixed-price contracts.

The trouble with the above incentive scheme is that it puts too much risk on the agent. Even if he or she works hard, due to exogenous uncertainty performance is sometimes going to be unsatisfactory and the agent will get 'punished' even though the fault is not his or hers. So long as the agent cares only about his or her expected salary, this does not matter. But almost everyone cares not just about his or her expected earnings, but also about the range of possible variation in the earnings. That is, they are risk averse. They evaluate a risky income stream not in terms of the expected value but the expected value less a discount the size of which depends on the variability of the income stream.

Because of the above reason, a very sharp incentive scheme may be too costly for the principal since the agent will discount it by the amount of risk he or she has to bear. Since the principal is

[4] For example, if the salary is $1000 a month if the outcome is satisfactory and 0 otherwise, then working hard entails an expected salary of $750 and not working hard yields an expected salary of $500. So long as the cost of working hard does not exceed $250 in monetary terms, the agent will work hard.

likely to be less risk-averse than the agent (e.g. because he or she is wealthier and better able to cope with risk), the solution lies in offering an incentive scheme that lies somewhere between a completely flat salary and a sharp incentive scheme where the agent bears all risk. Different forms of profit sharing and bonuses are examples of such incentive schemes. This solution comes with a cost. The agent still has to bear more risk than is optimal, since otherwise he or she will not put in extra effort. That means the principal has to pay him or her a higher expected salary to offset the discount due to risk aversion, which is costly from the principal's point of view. Also, it is possible that the principal will have to be reconciled to the agent putting in a low or moderate effort level, if the extra wage cost in order to give incentives does not justify the gain in expected profits to the principal from the high effort level.

Risk aversion is not the only possible reason why it is hard to implement a 'perfect' incentive scheme. Another reason is limited liability. Suppose the agent is not risk-averse but must be given some minimum amount of salary due to labour regulations or social norms. This puts a lower bound on how much the agent can be penalized if the outcome is unsatisfactory. It restricts the use of incentive mechanisms such as performance bonds. Given that the agent cannot be penalized very much if the outcome is unsatisfactory, the only way to induce the agent to put in effort is to reward the agent when the outcome is satisfactory. But from the principal's point of view, punishment is cheaper than rewards. As a result, the principal might want to induce a lower effort level than what would be possible under perfect monitoring to economize on rewards given out to agents.

Revisionist perspectives

The classical principal-agent model is well suited to analyse situations where there is a single principal dealing with a single agent as performance can be measured even though the measure is noisy, and the agent is entrusted with a single task. Clearly, these assumptions do not apply to many situations and accordingly the incentive theory has been extended to deal with these situations. Below we briefly discuss some of the major directions in which the basic principal-agent model has been extended (see Dixit 2002 for a detailed review).

Measurability

In many cases the measurement of performance is very difficult. For example, how does one measure the performance of a teacher? The task of a teacher is to provide 'good education', but this is much harder to measure than say, production in an assembly line, sales, or provision of banking services or even some public services such as garbage removal or power supply. This means that in these cases it would be hard to find good performance measures. If performance measures are noisy, then making rewards very sensitive to performance does not give effective incentives, and imposes unnecessary risk on the employee. In these situations, the solution is to offer a flat incentive scheme and introduce a subjective performance evaluation. If the agent's performance is considered satisfactory by the principal (but hard to quantify and explicitly introduced in the contract) then he or she will be given a raise in the future, or a promotion.

Multi-tasking

Most jobs involve several tasks. If some of these have good performance measures and not others, then making an agent's pay sensitive to the good performance measures will cause him or her to substitute effort away from the other tasks, and could result in a loss of efficiency. To continue with

the example of teaching, good education involves students being able to achieve high scores in standardized tests, but also encouraging a spirit of creativity, curiosity and the inculcation of good values. The former is easy to measure but if teachers are rewarded just on the basis of the performance of students in tests, this might lead to an excessive focus on examination skills at the expense of the other components of a good education. This makes provision of incentives hard when employees have to perform multiple tasks (Holmstrom and Milgrom 1991). Similarly, if hospitals are given incentives to cut costs, they are going to sacrifice quality by refusing to treat certain types of illnesses or being excessively selective in using expensive medical procedures.

Multiple agents

Sometimes several agents are involved in a task, and it is hard to separate the contribution of one agent from the other. This is referred to as the problem of moral hazard in teams (Holmstrom 1982). In this situation, in addition to the principal-agent problem between the principal and the team of agents taken together, there is also a free-rider problem among the agents. In other words, since individual contributions cannot be separated and so incentives of each individual agent will depend on the performance of the team, each agent will supply less effort than if performance depended only on his or her own effort.

Multiple principals

In some organizations an agent works for multiple principals who may not share the same objectives. In a profit-maximizing firm one would expect that in the ultimate analysis, only the net profits matter. But clearly members of the organization at the top of the hierarchy may have their

private agendas as well, and due to principal-agent problems between the owners and these individuals, the divergence of objectives may not be completely eliminated through incentive schemes. For example, the head of the engineering division may be driven by an objective of coming up with an innovation that will be his or hers, or the department's ticket to fame. But the head of the sales division may be concerned only with sales. Also, a doctor can be concerned with the success of a particular treatment method. This can be in conflict with the interests of the hospital management (or tax payers at large) who would like to minimize costs. This can also be in conflict with patients, who might not wish to be subjects of experimentation. Similarly, a teacher might want to give more emphasis on learning using expensive teaching aids, as opposed to imparting mechanical test-taking skills. This might make the (enlightened) parents happy, but the school principal or management might care more about the average test-record of their students, and tax-payers may be more concerned about the expenses. Since each principal would like to induce the agent to put more effort in activities that he or she cares about more, if the incentive schemes are not chosen to maximize the joint payoffs of the principals, there will be inefficiencies over and above the basic agency costs because of the lack of cooperation and coordination among the principals.

Implicit incentives

The classical principal-agent model takes the view that the agent works exclusively for the principal and has rewards set inside the organization. However, in practice, principal agent relationships are set in a market context. The market determines the outside option of the agent. Also agents may have one eye on the market when he or she chooses how much effort to put in – the rewards may come from the market rather than from the current principal. An agent who is

viewed as performing well will command a higher market price in future periods. Hence incentives become implicit (see Holmstrom 1999 and Dewatripont et al. 1999). In this instance, the principal loses some control over the incentive schemes that can be offered due to the operation of market forces.

Empirical evidence

Empirical evidence on the standard economic approach to principal agent problems is relatively thin especially in comparison to the extensive theoretical literature (see Prendergast 1999; Chiappori and Salanie 2003 for excellent surveys). There are two key issues that empirical research has focused on. Firstly, do incentive schemes affect performance? Secondly, are incentive schemes optimally chosen? Prendergast (1999) in his survey concludes that while the answer to the first question is, in general, yes, the answer to the second question is mixed and inconclusive. The main difficulty is what empirical economists call the problem of 'identification'. It is rare to find changes in incentives which occur for genuinely exogenous reasons. As a result, if one compares across two sets of incentive schemes and measures the difference in performance, it's not clear whether one is picking up the effects of the variation in the incentive schemes, or differences in the two environments or the characteristics of the agents. For example, if we observe one firm using high-powered incentive schemes and also having higher measures of productivity compared to another firm, we cannot conclude that greater incentive pay raises productivity. It could be that more productive agents self-select to work in firms that offer greater incentive pay. This problem of selection can contaminate evidence. Although, overall the evidence does point out a positive effect of incentive pay on performance, the selection effect can be important in specific cases.

For example, Lazear (2000) looks at how windshield fitters respond to the introduction of piece-rate incentives. He finds a large effect that can be explained in significant measure by the selection of different individuals in piece-rate jobs.

The three M's spproach

In this section we report on our ongoing work (Besley and Ghatak 2003; 2004a) which builds an approach to the principal-agent problem with some distinctive features. This approach can be applied to the design of incentive schemes in public and private organizations.

There are three key elements in our approach: missions, motivation, and matching. A mission consists of the attributes of a project that make some principals and agents value its success over and above any monetary income they receive in the process. This could be based on what the organization does (charitable versus commercial), how they do it (environment friendly or otherwise), who is the principal (kind and caring versus strict profit-maximizer) and so on. Motivation is any value in excess of the monetary rewards from doing something that a principal and agent may derive. This can be viewed as job satisfaction and is likely to be greater if things are done in the way an individual likes. For example, a sales representative in the Body Shop may put in extra effort since she is opposed to animal-testing and has strong pro-environment views. Matching is the process by which firms bring together like-minded principals and agents. Because they share similar views on missions, and therefore are motivated not just by the monetary rewards, the principal-agent problem is alleviated and the need to give explicit incentives (which are costly for reasons discussed above) is lessened. We discuss these concepts in greater detail below.

Missions

The mission of the organization, displaces the conventional notion of profit maximization. The idea that missions are important in public organizations is not a new. It is a central plank of James Q. Wilson's celebrated study of public bureaucracies (Wilson 1989: 95). He defines a mission as a culture 'that is widely shared and warmly endorsed by operators and managers alike'. It is an important and frequent theme in the literature on non-profit organizations (see, for example, Sheehan 1998). It is the nature of the activities in question and not whether the service is provided publicly or privately that unites mission-oriented organizations.

While the notion of missions is somewhat vague compared to notions like profit, we believe that it is an important departure when thinking about organizations that are not directly responsive to market forces. Missions can also be important in more standard private sector occupations. Firms frequently profess that their goal is to serve customers rather than to make their shareholders as rich as possible. One can question whether these firms are genuinely committed to these missions, or whether these are just a veil for some other underlying self-interested behaviour. However, there need not be any incompatibility between profit-maximization and commitment to specific missions. For example, firms that adhere to norms of corporate social responsibility and adopt business practices that seemingly sacrifice some profits (e.g. adopting environment-friendly technologies) may actually be better off doing so in net terms, since they would attract workers who share the same mission preference.[5] If principals and agents share a view of the mission, it is likely that an effective mission will economize on monetary incentives.

[5] Analogously, such firms may be able to attract customers who may be willing to pay a higher price to 'consume' the mission along with the product. In Besley and Ghatak (2004b) we analyse corporate social responsibility from this point of view.

We assume that the mission of the organization is determined by the principals in the organization. This can be a heterogeneous group with overlapping responsibilities. For example, in the case of a school, the principals are the parents, the government and the head teacher. Preferences over missions can be heterogeneous. For example, some parents may value high levels of discipline. There could also be disagreement on the right curriculum choices such as the weight to be attached to music teaching or languages. An important role of the management in a mission-oriented organization is to foster a congruent outlook. Thus as Miller (2002: 446-7) argues in the context of her case studies of twelve non-profit organizations, 'Non-profit board members do not expect conflict between the executive director and the purpose for which the organization was created. The board believes that the executive management will not act opportunistically and that what management actually does is ensure good alignment and convergence in its relationship with principals.'

Changing the mission of an organization in a way that is not favoured by the agents can reduce the efficiency of the organization. In that sense, the approach shows why mission oriented organizations are conservative and slow moving since there is a rigidity built in from the types of agents who are attracted to the organizations. Organizations without mission-oriented agents, such as private firms, are likely to be more flexible and adaptable.

Motivation

A key assumption in the principal-agent model is that performance benefits from the effort put in by agents and that this effort is costly and that the agents in question have to be motivated to put in effort. But rewards for putting in effort are not always purely monetary – agents may be motivated

because they care about the output that is being produced. However, such non-monetary rewards depend on the way in which the organization is structured. For example, teachers may care about teaching a curriculum that they think is most conducive to learning. Thus, the mission of the organization can affect the degree to which agents are willing to commit costly effort.

When goods are produced with external benefits (i.e. benefits that exceed the market value of the good or the service being produced) then individuals who work in the production of these goods may factor the value of the output that they produce in their decision to work in that sector and into the amount of effort that they put in. This is the labour market equivalent of the idea that individuals engage in private supply of public goods and those with the highest valuation of public goods may have the greatest interest in contributing. The model could also be one in which individuals are 'altruistically' motivated or that they get a 'warm glow' from doing social good.[6] In the former case, the level of the good being produced matters to the individual, but not who provides it. This can lead to free-riding. In the latter case, it's not the level of the goods, but how much the individual contributes to it that matters. It is clear that on either of these views the value of what they do should be attached to the job that they do and not the sector in which they do it. Thus, if a nurse believes that nursing is an important social service with external benefits, then it should not matter whether she is employed by the public or private sector except in so far as this affects the amount of the benefit that she can generate.

The general point here is that a system of organization and remuneration for the provision goods and services will have to take into account not only how on-the-job incentives affect those in the sector work, but also who is attracted to work there. This might alleviate the need to give high-

[6] These ideas are also related to the strong professional ethics that govern the behaviour of workers in the production of collective goods. Such ethical codes de-emphasize narrow self-interest.

powered incentives. Francois (2000) has shown that government bureaucrats are not residual claimants which implies that they can commit to a 'hands-off'' policy that elicits greater effort from workers who have 'public service motivation'. However, if individuals differ in terms of how motivated they are, and in addition have heterogeneous mission-preferences, it is important to examine the process by which agents are matched to an organization, a topic to which we turn to now.

Matching

Matching is the process by which principals and agents come together to create an organization. Matching serves an allocative role in bringing consumers to providers ('product market matching'') and of workers to providers ('labour market matching').

If consumers care about the missions adopted in organizations, then allowing them to choose between providers with different missions is a potentially important source of welfare improvements. This argument applies to both the private and the public sector. There is no reason why a consumer could not exercise choice between two competing hospitals or schools in much the same way that they choose a TV or a car. This application of private goods choice to public services underpins the standard argument for voucher provision of public services.

If workers care about missions in organization, then principals and agents can match with one another on the basis of the perceived mission of the organization. This is a natural consequence of organizations being mission oriented. This matching increases efficiency in the operation of organizations since the returns from putting in effort are higher when agents share the same goals as those espoused by the organization. This process of selection at the entry point alleviates the need to use explicit financial incentives on the job.

Applications

In this section we discuss how the above general approach can be applied to specific issues. We begin by discussing the role of incentives in different kinds of organizations. Secondly, we discuss differences between public and private organizations. Thirdly, we discuss different views of the role of competition and finally, incentives to innovate.

The role of incentives

The standard principal-agent model underpins the classic economic view of organization design. However, some of the business school literature on the firm (see, for example, Harvard Business School 2003) has tended to emphasize the importance of motivation in the firm without paying sufficient attention to the problem of getting incentives right. Moreover, this more hands-on literature also gives a much bigger role to selection of individuals whose motivations cohere with those of the organization. This view is therefore quite compatible with the three M's approach that we have outlined.

Akerlof and Kranton (2003) have recently developed a view of organizations based on the sociological notion of identity. In their view organizations develop identities and people conform to these identities in effective organizations. They emphasize the role of entry rituals in socializing individuals into the appropriate identity. Like the three M's approach, they put much less weight on the role of monetary incentives in organization design. Akerlof and Kranton cite Max Weber's notion of the importance of vocation. According to Weber, an office is a vocation: 'Entrance into an office is an acceptance of the fealty to the purpose of the office'. Akerlof and Kranton remark that if Weber's observation reflects the behaviour of most jobholders, 'the standard economic

theory of behavior in organizations (principal-agent theory) has missed most of what causes them to function'.

Which of these views is most appropriate to understanding any particular production process will vary a great deal by type of organization. We are not claiming that monetary incentives as modelled in the classical principal-agent problem are never relevant, even though economists have perhaps overplayed their significance. The three M's approach is able to bring a traditional economic view and a more sociological view of incentives closer together.

Public versus private organizations

Public sector reform is one of the key policy areas where issues of risk management and incentives in general are important. At the heart of many debates is an issue of whether the public and private sectors are fundamentally different in important ways. Under the auspices of the New Public Management (see, for example, Barzelay 2001; Hood et al. 1999) there has been a huge focus on bringing the practice of incentive pay into the public sector. A simplistic view of principal-agent relationships would support this as the *sine qua non* of effective organization. However, our analysis calls this into question.

One possibility is that ideas of mission and a culture of control based on limited monetary incentives are indeed peculiar to the public sector. This view was embodied in the idea of the Whitehall village. Thus, Heclo and Wildavsky (1974) characterize Whitehall in terms of:

> The traditional picture of a village world regulated in a relatively informal way through largely unwritten rules, a compliance culture and low relational distance between regulator and regulated still appeared to capture much of the style of regulation within Whitehall a quarter of a century after Heclo and Wildavsky's study.
>
> Quoted in Hood et al. 1999: 74

The debate between these competing visions of public management can be thought of in terms of the three M's framework. The Whitehall village model has strong use of missions, weak need for monitoring and emphasizes matching and socialization in controlling organizations. By contrast, the New Public Management model emphasizes weak mission preference and heavy use of monitoring and incentive pay.

The importance of mission in effective public organization is a dominant theme of the literature on public bureaucracies – see, for example, Wilson (1989). This suggests that a view of public organization with a strong emphasis in mission in aligning incentives makes a lot of sense in this context. This explains well why models of organization that de-emphasize the standard principal-agent concerns are found in this literature. The attack on this conception by public choice theorists like Niskanen (1971) fundamentally questions whether there are any key differences in motivation between public and private employees. While this is difficult to test systematically, there is evidence from the literature on non-profit organizations that motivation differs between individuals who work in for-profit and not-for-profit organizations.

What is clear is that there is no particular reason to believe that there is something special about publicly-owned organizations. Motivation seems much more likely to attach to what an organization tries to achieve – the extent to which it works in the public interest. While the ownership structure may be a correlation of this, it is not what matters per se. Thus, if there are differences in organizational structure between different organizational forms, we would expect them to be more broadly correlated with the way in which organizations operate. For example, we would not expect large differences between state and private schools and public and private hospitals.

The role of competition

We have not so far discussed the role of competition in affecting incentives and organization performance. However, the forces of competition are frequently appealed to in improving organizational effectiveness.

In traditional principal-agent models there are two effects that determine how competition affects incentives in organizations (see Schmidt 1994). Since the agent is rewarded out of any rents that the principal earns in the market place, competition that limits rents will tend to reduce incentive pay and hence organizational efficiency. However, competition may often reduce the probability that an inefficient organization will survive. If the agent earns a rent from working for the principal, this liquidation effect will tend to increase organizational efficiency.

In situations where there are many agents completing similar tasks, there is a potential role for competition based on benchmarking of performance – so-called 'yardstick competition'. A good example of this is the use of league tables in the regulation or schools and hospitals.

The three M's approach suggests that the role of competition in improving organizational efficiency can also work through improved matching in the labour market. Agents who are poorly matched will have higher incentive pay and lower levels of productivity. Matching reduces the misalignment in missions and hence increases efficiency.

Innovation in organizations

Organizations are frequently under pressure to seek out new ways to please customers and to improve efficiency. The forces that shape an organization's responsiveness to such opportunities are an important source of long-run advantage. In the standard principal-agent approach of agents

with preferences over pecuniary aspects of their jobs, innovation incentives are achieved by paying for adaptability pure and simple.

The three M's approach complicates matters. Since agents have an interest in the mission of the organization, they may also have a direct preference over innovations in so far as the latter interact with mission preferences. Thus, a new method of teaching in a school may appear attractive to one educationist but may be viewed by some teachers as interfering with their educational goals. Thus the three M's approach does lead to some understanding of organizational conservatism that does not appeal to emotions and innate conservatism. In other words, conservatism is quite consistent with a rationally based view of organization design. To the extent that the forces in the three M's approach are different across different types of organization, then we would expect organizational conservatism to differ too.

The approach also reveals that the efficiency of innovation needs to be considered in a broader framework – a narrow financial criterion may be a misleading basis for advocating efficiency improvements if there is de-motivation because agents are less in tune with organizational goals. These considerations need to be added to the purely financial criterion and brought into the criteria for efficient innovation.

Concluding comments

This chapter has discussed economic approaches to risk management and organization design. The three M's approach provides a means of thinking through some issues that fall outside the ambit of standard economic approach. Yet, it also maintains the spirit of many key ideas in principal-agent theory. The approach also helps to bridge the gap between thinking about incentives in economics and in other branches of the social sciences. Much remains to be done to further this agenda. But it

provides a better understanding of the limitations of focusing only on monetary incentives and an appreciation of the role of selection or matching in better aligning the objectives of employees of a firm with that of the owner or the manager.

References

George Akerlof and Rachel Kranton, 'Identity and the economics of organizations', typescript, University of California, Berkeley (2003).

Michael Barzelay, *The New Public Management – improving research and policy dialogue* (Berkeley & Los Angeles: University of California Press, 2001).

Timothy Besley and Maitreesh Ghatak, 'Incentives, choice and accountability in public service provision', *Oxford Review of Economic Policy* 19, no. 2 (2003), 235-49

Timothy Besley and Maitreesh Ghatak, 'Competition and incentives with motivated agents', typescript, London School of Economics and Political Science (2004a).

Timothy Besley and Maitreesh Ghatak, 'The economics of corporate social responsibility', typescript, London School of Economics and Political Science (2004b).

P. A. Chiappori and B. Salanie, 'Testing contract theory: a survey of some recent work', in M. Dewatripont, L. Hansen and S. Turnovsky (eds.), *Advances in economics and econometrics,* vol. 1 (Cambridge: Cambridge University Press, 2003).

Mathias Dewatripont, Ian Jewitt and Jean Tirole, 'The economics of career concerns', *Review of Economic Studies* 66, no. 1(1999), 189--217.

Avinash Dixit, 'Incentives and organizations in the public sector: an interpretive review', *Journal of Human Resources* 37, no. 4 (2002), 696-727.

Patrick Francois, 'Public service motivation as an argument for government provision', *Journal of Public Economics* 78 (2000), 275-99.

Bruno S. Frey, Not just for the money: an economic theory of personal motivation (Cheltenham: Edward Elgar Publishing, 1997).

Harvard Business School, *Harvard Business Review on motivating people* (Cambridge MA: Harvard Business School Press, 2003).

Hugh Heclo and Aaron Wildavsky, *The private government of public money* (London: Macmillan, 1974).

Frederick Herzberg, 'Workers' needs: the same around the world', textquotedblright, *Industry Week*, 21 September 1987.

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

Bengt Holmstrom and Paul Milgrom, 'Multi-task principal-agent analysis: incentive contracts, asset ownership, and job design', *Journal of Law, Economics and Organization* 7 (1991), 24-52. Special Issue.

Christopher Hood, O. James, C. Scott, GW Jones and T. Travers, *Regulation inside government: waste watchers, quality police, and sleaze-busters* (Oxford: Oxford University Press, 1999).

Steven Kerr, 'On the folly of rewarding A while hoping for B', *Academy of the Management Executive* 9, no.1 (1995), 7-14.

Alfie Kohn, 'Why incentive plans cannot work', *Harvard Business Review* (September – October 1993), 54-63.

Edward Lazear, 'Performance pay and productivity', American Economic Review 90, no. 5 (2000), 1346-61.

Judith L. Miller, 'The board as monitor of organizational activity: the applicability of agency theory to non-profit boards', *Nonprofit Management and Leadership* 12, no. 4 (2002), 429-50.

William Niskanen, *Bureaucracy and representative government* (New York: Aldine-Atherton, 1971).

Canice Prendergast, 'The provision of incentives in firms', *Journal of Economic Literature* 37, no.1 (1999).

Kieran Poynter, 'Risk and value management', *Risk and Regulation* 7 (Summer 2004). Magazine of the ESRC Centre for Risk and Regulation

Klaus M. Schmidt, 'Managerial incentives and product market competition', *Review of Economic Studies* 64 (1997), 191-213.

R. Sheehan, 'Mission accomplishment as philanthropic effectiveness: key findings from the excellence in philanthropy project', *Nonprofit and Voluntary Sector Quarterly* 25, no. 1 (1996), 110-23.

James Q. Wilson, *Bureaucracy: what government agencies do and why they do it* (New York: Basic Books, 1989).