

Are public sector workers motivated by money?

Carol Propper

Department of Economics and CMPO, University of Bristol
CHERE Distinguished lecture March 2006

Acknowledgements

This paper was written during a visit to CHERE, UTS and I am most grateful for the hospitality of the Centre and the University of Technology Sydney

Abstract

Reforming the public sector occupies a central position in the policy agenda of many countries, Australia included. Two features of private market activity have attracted the attention of would-be reformers of the public sector – the role of choice and competition and the use of financial incentives. Modernisers argue that only by introducing these measures will we be able to improve our public services; traditionalists resist such changes arguing that choice is not appropriate when it comes to public services and that the introduction of financial rewards will drive out ‘public sector’ motivation.

The lecture focuses on research on the second of the two planks of public sector reform – the use of monetary incentives to motivate individuals who are delivering public services. While the literature suggests there are good reasons why a wholesale importation of methods from the private sector will not work, adopting the opposing view that financial incentives do not matter for, or do not motivate, public sector workers is equally unhelpful. There are reasons why we might not expect to see as much use of financial incentives in a public setting compared to a private firm. But the lesson from this is not that financial incentives don’t matter; just that they should not be as ‘high powered’ as in (some) private sector firms. I review recent evidence, drawn from a series of case studies that have been undertaken on individuals employed deep into the traditional public sector – in secondary education, government administration and public hospitals - which shows that incentives schemes can bring about increased productivity and that policy makers ignore the important of financial rewards at their peril.

Introduction

Reforming the public sector occupies a central position in the policy agenda of many countries, Australia included. Two features of private market activity have attracted the attention of would-be reformers of the public sector – the role of choice and competition and the use of financial incentives. Modernisers argue that only by introducing these measures will we be able to improve our public services; traditionalists resist such changes arguing that choice is not appropriate when it comes to public services and that the introduction of financial rewards will drive out ‘public sector’ motivation.

In this lecture I’m only briefly going to address the issue of expanding choice¹. Instead, I’m going to focus on research on the second of the two planks of public sector reform – the use of monetary incentives to motivate individuals who are delivering public services. I will argue that while the literature suggests there are good reasons why a wholesale importation of methods from the private sector will not work, adopting the opposing view that financial incentives do not matter for, or do not motivate, public sector workers is equally unhelpful. I’ll begin with a review of the arguments as to why the public sector is different and why we might not expect to see as much use of financial incentives in a public setting as in some private firms. But the lesson from this is not that financial incentives don’t matter; just that they should not be as ‘high powered’ as in (some) private sector firms. I’ll then move on to some very recent evidence, drawn from a series of case studies that have been undertaken on individuals employed deep into the traditional public sector – in secondary education, government administration and public hospitals - which shows that incentives schemes can bring about increased productivity and that policy makers ignore the importance of financial rewards at their peril.

1. Why the public sector is different

Economists have stressed that public services are services that are produced, at least in part, for the benefit of the public at large. There are three well-known reasons why an unregulated market will underprovide ‘public’ services:

¹ See Propper et al (forthcoming) for a review of the economic evidence on choice in healthcare and education.

- External benefits – there are some goods that benefit not only those who receive them but society at large. So schooling improves not only the lives of those who receive it, but also benefits society as a whole for example, by improving the productivity of workers, increasing the spread of ideas and creating a flexible labour force. Better health care not only improves the lives of those who receive it, but in the case of vaccination programmes for example, reduces the spread of disease to others.
- Egalitarianism – society cares about the consumption of goods by the poorest individuals
- Merit good – society wants to ensure that individuals consume a sufficient amount of certain goods that the individual themselves may not sufficiently value.

The existence of these concerns leads to solutions for the provision of public services that attenuate the use of the profit motive. However, there is much debate about the precise form that this provision should take. The traditional model in the UK has been centralised government provision, with attempts to achieve a fairly uniform level of provision across the country. But such arrangements have come under fire for giving consumers little choice, being inefficient and in some cases, focusing resources on those who are best able to operate the system – the educated and articulated middle classes (Le Grand 1991). The on-going reforms are an attempt to move away from this model, borrowing from the private sector of the ideas of competition, choice and financial incentives.

The use of choice and competition has been a theme running through the reform of UK public services since the late 1980s (Le Grand 1991) and is one that has been personally promoted by the leader of the current UK administration, who has made increasing choice in education one of the key planks of his last term. Other countries have also seen widespread experimentation with choice. In the US, a strongly espoused – though challenged - view is that choice in education is a ‘tide that raises all boats’ (Hoxby 2002). The general view in US healthcare is that choice reduces costs and may, in certain circumstances, improve also quality (Gaynor 2004). Nordic countries - with a strong commitment to equality - have experimented with choice as a means of increasing productivity in both their education and health care systems. But some enthusiastic adopters have had less success with such policies: for example, New Zealand’s experiments with choice in education in the late 1980s and 1990s. (Burgess et al 2005).

For the use of financial incentives for those responsible for delivering public services, the literature has identified three respects in which public services production differs

from that of the private sector. Each will affect the strength of optimal financial incentives. These are the difficulty of measuring performance, the issue of multi-tasking and the existence of multiple principles².

- *Measurability*

In many cases public services are multifaceted and, as a result, the objectives of an organisation are difficult to define and so to measure. So for example, the objectives of a school might be to provide a 'good education' but this is harder to define than, say, the production of a cars or the adequate collection of garbage. Even breaking down the overall objectives into sub-components can be difficult; there is for example, much debate about what constitutes a good measure of hospital quality. This means that it can be hard to find good measures of performance and the measures adopted may be very noisy – they impart relative little information about the effort of an employee or the organisation (Propper and Wilson 2003). As a consequence, linking rewards to the meeting of performance targets does not give effective incentives, may lead to 'gaming' and imposes unnecessary risk on employees.

- *Multi-tasking*

One of the reasons that these services are complex is because they involve several dimensions, some of which are relatively easy to measure, others of which are much harder to measure. Examples of the former are school students' performance on standardised tests and the re-admission of patients after operations. Examples of the latter are becoming culturally aware or improving the long term health of patients. This difference in measurability may mean that incentives can only linked to the easy to measure outcomes. This may lead to an excessive focus on these outcomes at the expense of other tasks. So hospitals who receive financial incentives which are linked to cost reduction may have incentives to shave quality, teachers monitored on student pass rates to reduce the effort they put into less able students.

- *Multiple principles*

Precisely because the services are public in nature, several groups in society will be interested in their provision, cost and efficiency. The tax payer is always one party: tax payers' desires to keep down costs may conflict with those who are providing the service who are interested in high quality. Within an organisation, public service providers may serve several masters: for example, doctors get professional recognition from their research activities, but the managers (and

² This section draws heavily on Besley and Ghatak (2003).

possibly the patients) of hospitals might prefer they focused their time on other, less professionally recognised tasks, such as reducing waiting times for common procedures. The economics literature stresses that in these cases - referred to as multiple principles - financial incentives should be weaker than in the case where there is only a single principle.

Recently, Tim Besley and Maitreesh Ghatak (2003) have suggested that the type of firms who produce public services have missions, which serve to attract workers who share these missions. Missions have been defined as a culture 'that is widely shared and warmly endorsed by operators and managers alike' (Wilson 1989). High-powered financial incentives might destroy these missions, or at least be a sign of poor alignment of missions. Patrick Francois (2000) has suggested that the not-for-profit form of organisation with its lack of strong financial incentives may be a good way of delivering public services precisely because the financial incentives are weak.

All these reasons are ones why we might expect financial incentives to be less widespread in public service provision than in private firms where the output is easily measured by a single number – profit - and perhaps missions are less important. But these arguments don't imply that *any* linking of employee performance to financial incentives is necessarily bad. It may be the case that the introduction of some link between pay and performance will improve productivity in public services. This is essentially an empirical matter: theory states that the incentives should be less 'high powered' - less linked to performance - but empirical studies are needed to say how strong the link can be in practice.

In fact, while the matter is essentially an empirical one, there are relatively few empirical studies of this issue (compared, for example, to the many studies of whether CEOs respond to stock options). This lack of evidence partly arises because of the inability or unwillingness of governments to carry out experiments within the public sector: changes are either introduced for all service delivers or all eligible recipients, or they are not introduced at all. This means that it is difficult to net out the effect of a scheme from other changes that may have taken place at the same time as the scheme. So much of the early evidence is not very robust. Consequently, a meta-analysis in which it is concluded that the 'bulk' of the evidence shows this or that cannot be provided as – despite the strong positions taken by the proponents or opposers of attempts to incentivise the public sector - there is no bulk (Burgess and Ratto 2003).

Instead, I present evidence for various different groups of individuals who provide public services. In selecting evidence I have focused on the more robust evidence that is taken from either actual experiments or from so-called 'natural experiments': situations in which those subject to performance related pay can be compared to an appropriate control group. Within this, I look at the use of performance related pay for two types of individual. The first type are professionals who are employed in services that we think of as traditional public service provision – health and education. These individuals often work in the public sector, but this is by no means always the case: they may be employed in not-for-profit firms (for example, teachers in private schools

in the UK, or doctors in not-for-profit hospitals in Australia) or even in for-profit firms (for example, doctors in hospitals in the US). The second type are individuals who have jobs that are closer to those of clerical employees in the private sector, but are public sector employees because the services they provide are funded and provided by the state.

Some of the schemes studied are individual reward schemes: that is, they link individual performance to individual rewards. Others are team-based schemes, in that individual rewards are linked to the performance of a team. Team based schemes are more appropriate where the output of a single employee depends heavily on the actions of other employees around them (Holmstrom 1982). Team based production is common in health care; it may also be important in other public services.

2. Evidence on incentives for public employees³

2.1 Evidence from the professions

Doctors

There is a widespread literature on the responses of doctors to changes in financial incentives, generally brought about by the changing of fee schedules. Some studies examine self-employed doctors, some examine doctors who are salaried, but a common feature is that they tend to show that if doctors are paid more, they do more (e.g. Scott 2000). Many of these studies suffer from the weakness that they are not based on experiments, so cannot distinguish between the incentive effect of schemes and the fact that certain doctors – those that value money – are attracted to schemes which link output and pay, whilst those who value leisure opt for schemes which are less high powered.

However, there is a small amount of experimental evidence, drawn from a variety of settings and countries. In Denmark, a move from payment per patient enrolled on a doctor's books to payment for each procedure undertaken (known as fee-for-service) saw family doctors increase their diagnostic and curative services and decrease their prescribing and referrals. This fits with the doctors doing more work themselves, rather than referring on to specialists or prescribing medication (Scott 2000). In Quebec, in the late 1970s, in a bid to decrease medical expenditure, the government reduced the reimbursement rate paid to doctors for a set of activities once their total expenditure on these activities had hit a target. These targets were set for three month

³ This section draws heavily on Burgess and Ratto (2003) and Propper (2005).

periods, which meant that once the doctors has hit their ceiling the financial payment they earned for each activity above this level was reduced. Doctors responded to these incentives by sharply decreasing the activity they undertook in the third month of each accounting period, some even moving to taking regular vacations every three months (Rochaix 1993). In the US, Gaynor and his co-authors exploit differences in team size across doctors working in HMOs (health maintenance organisations) to identify the impact of team based incentives (Gaynor et al 2001). They find that contracts that provide \$1 of income for every \$10 reduction in medical expenditure leads to a reduction of around 5 percent following the contract.

The UK internal market reforms provide evidence on responses of GPs to financial incentives. The 1991 reforms, described as the ‘boldest of market based reforms’ in public health care, created, out of the public sector, separate sellers and buyers of hospital-based health care (Propper 1995). One of the two sets of buyers were family doctors (known as fundholders), who were given budgets to buy elective care for the patients in their practices. Crucially, fundholders were able to keep any surplus from their budget (to be used at the level of the practice).

For reasons that had nothing to do with the scheme itself, when a GP wished to become a fundholder, they had to announce their intention and then wait for a year whilst their referrals to hospitals were counted in order to work out their budget. So the obvious question to ask was whether GPFHs increased their referrals in the year before becoming a fundholder in order to increase the size of their budget once they became a fundholder. In a study of GPs located in one area of the UK, Propper et al (2001) found that fundholders did exactly that: they increased their referrals to hospitals relative to their previous referral patterns (also controlling for changes that were happening to non-fundholders) by about 10 percent. Once they became fundholders, their referral patterns dropped by about 10% and thereafter appeared to revert to their long run normal levels. While this may have benefited both the fundholders and their patients, it is clear that those who became fundholders did respond to financial incentives. And in the process, because the total pot for buying hospital care was finite, they also took monies away from patients served by practices that weren’t fundholding⁴.

Teachers

Possibly the one areas where there are most papers on incentives schemes in the public sector is that of performance related pay for teachers, due to the recent popularity of these schemes in the US. (It is worth noting that such schemes are not

⁴ A later study examining the end of fundholding for all GPs in England came to very similar conclusions (Dusheiko et al 2004)

new: the UK had a performance related pay scheme for teachers in the late 19th century: it was abandoned for fear of giving teachers inappropriate incentives.)

Of the early evidence on teachers, Burgess and Ratto (2003) conclude that there is not a great deal which is robust and credible. However, more recent evidence appears to show positive responses of teachers to performance related pay. Lavy (2002, 2003) reports on two different experimental schemes operating in Israel in the public school system. The first examined a scheme in which schools competed on their average performance and bonuses were distributed equally to all teachers in the winning schools. This had a significant effect on pupil performance. The second was an individual (teacher) level scheme, in which teachers were ranked on their pupil's performance and the bonuses awarded to the best performing teachers. The measure of pupil performance was a sophisticated one, which allowed for the previous performance of the pupil and their socio-economic characteristics. The rewards from good performance were high: teachers in the top quarter of the distribution were awarded a bonus equivalent to around 25% of average teacher salary. A significant impact of performance was found, and interestingly there appeared to be no obvious negative spillovers to teaching of other subjects (Lavy 2003).

The UK government also introduced a performance-related pay policy for teachers, with pupil progress as one of its key criteria. This incentive scheme was explicitly teacher-based (rather than school-based). The Performance Threshold scheme was introduced in 1999 to give incentives to experienced teachers, who had been previously paid on a unified basic salary scale and could only raise their wages by taking on extra administrative duties. The scheme gave teachers who performed well in terms of pupil performance an annual bonus of £2000 per annum and put them on a new upper pay scale where they could continue to receive increments without having to take on extra duties. These increases, which could take a teacher's pay up to £30,000, were assessed against rigorous criteria and annual targets. At the time, the concept of bonuses for individual teachers was condemned by unions for being divisive and unfair.

My co-authors and I collected longitudinal data following around 180 teachers over two complete two-year teaching cycles, before and after the policy was introduced. We linked pupils to the teachers who taught them for specific subjects. For the pupils linked to the sample teachers, we collected prior attainment data, so we could control for pupil characteristics and measure the target of the scheme – pupil progress rather than the level of attainment. By looking also at the scores of those staff who were not eligible for the increased pay awards and comparing both sets of results, we concluded that higher pay did achieve better grades. Teachers eligible for the

incentive payment increased their value-added by almost half a GCSE grade per pupil relative to ineligible teachers⁵. Put another way, pupils post the scheme got scores on average half a GCSE point more than equivalent pupils taught by the same teachers before the scheme was introduced (Atkinson et al 2004).

2.2 Clerical workers in the public sector

There are individuals who are employed in the public sector, but who essentially perform clerical tasks that could also be carried out in the private sector; for example, tasks similar to those of clerical employees in a private insurance company. The economics literature gives one example of a successful public sector scheme in such a setting –the impact on Brazilian tax collection of the introduction of performance related pay (Kahn et al 2001). Reforms in 1989 involved the payment of financial incentives, based on both individual and team performance, in detecting and fining tax evaders. The amounts involved were significant, bonus were often worth more than twice the average annual salary. Examining the position before and after this reform, the scheme had a dramatic effect: fine collections per inspection after the scheme were 75 percent higher than before and while the authors do not have an experimental setting, they conclude that other factors – for example changes in the economy, are unlikely to drive this results.

This scheme is very high-powered: the size of the bonuses was very large, and arguably not of the size that many governments would introduce. The UK government introduced a more modest scheme in an attempt to boost public sector productivity for the government agency that places unemployed individuals into jobs and administers benefits to working age individuals who are out of work. The scheme ran from April 2002 to March 2003. The main features were as follows. First, the scheme was team-based scheme: the team was a district. The targets were defined at district level, all workers in the district get the bonus if the target is hit. The teams were large - there are only 90 districts covering the whole of the country, varying in size from 5 to 39 offices in the team, and from 264 and 1535 people within a team. The pilot scheme introduced the incentive structure in 17 districts, leaving 73 districts as controls. In common with many schemes, the form adopted was a step-function, based on a threshold level of performance. Workers were paid straight salary up to the threshold, a bonus was paid for hitting the threshold, and then no further increase in remuneration for further output. Thus incentives are very sharp at output levels just below the threshold but weaker further below or above.

⁵ GCSEs are a key set of exams taken by English pupils: they determine pupil's future scholastic and labour market options.

The scheme designers recognized that individuals may trade-off hitting quantity and quality targets and for this reason included targets for five different activities, one of which was a quantity target, the rest of which related to various aspects of the quality of the work undertaken. Teams were offered 1% of salary for each target hit (provided a minimum of 2 were hit) and a further bonus of 2.5% on top of that if all five were hit. So the payments were not huge – the largest that a worker could get was 7.5 percent of salary. In addition, the targets were measured with very different levels of precision. The quantity target – essentially the number of persons placed into work – is easily measured at the level of the office. But the other targets were much less precisely measured, and were, in contrast to the job placement target, measured at district, rather than office level and only based on a random sample of districts.

Our evaluation of this scheme is based on a before and after comparison of those offices located in incentivised districts with those that were not. We found that workers did respond to the relatively small incentive payments, but only for the quantity outputs. In terms of the measures of quality there was no change in behaviour: we surmise this was because it was crudely measured and the chance of any one employee's effort affecting this target very small. We also found considerable evidence of 'free-riding': individuals relying on other workers to do their tasks for them, so resulting in less productivity. We found productivity after the scheme fell in larger offices but rose in smaller offices. We surmise that this is because people working in smaller offices could monitor each other's activities more easily, so alleviating the free-riding incentives given by a team based scheme.

2.3 The importance of design

The results from these studies of public sector employees show not only that individuals in the public sector respond to monetary incentives, but also the importance of careful design: public sector workers may respond to incentives but – just as other workers - in ways that the scheme designers might not always want. One of the more general themes of the literature on measuring performance is that of 'gaming' by those whose actions are being monitored (Propper and Wilson 2003). When a target is introduced, individuals may work harder to achieve that target. But they may also divert effort from other less well measured activities to do this: a current example from the UK is that University lecturers monitored on their research papers may neglect their (less monitored) teaching. And if the target is poorly specified, individuals may put effort into appearing to meet the target, whilst not increasing productivity. When a target for numbers on hospital waiting lists was introduced in the UK, hospital administrators spent considerable effort removing dead individuals from these lists. This made the lists shorter, so that hospitals were more likely to hit their targets, but didn't mean that the living got treated any faster.

Within the pay for performance literature there has been a focus on the unintended consequences of the linking of pay to things that can be measured. For example, an examination of programmes to train unemployed individuals in the USA (the Job Partnership Training Act) that gave bonuses - in the order of 5 percent of the annual

budget - to organisations that hit annual targets of employment of its programme graduates found that training organisations manipulated the graduation of their enrollees to hit this target. Whilst such manipulation could simply be accounting, it was found that these accounting practices reduced the efficiency of the training given the program participants (Courty and Marschke 2004). In a similar vein, Asch (1990) examined a scheme in which US Navy recruiters were awarded bonuses if they achieve a quota of recruits before a cut-off date. He found that the number of recruits was higher immediately prior to the cut-off date and lowest immediately afterwards. Crucially, the quality of recruits fell as the cut-off date approached, suggesting that recruiters shaded down their selection criteria to hit their targets.

In sum, the evidence does suggest that individuals employed deep in the public sector do respond to schemes to link performance to pay and will respond quite rapidly to often small financial incentives. Some responses may increase productivity; other responses might lower it. There are caveats and I will return to these below.

But first I want to look evidence on the more general link between pay and public sector productivity, not from a scheme, but for the operation of the labour market for nurses in the UK.

3. Pay and performance in the UK NHS

As part of its efforts to increase productivity in health care, in 2001, the government introduced a rating system that measures the performance of NHS hospitals against a wide range of targets, including aspects of quality, volume and financial performance. These were called the 'star' ratings: every hospital was awarded between three (the highest rating) and zero (the lowest rating) stars. To sharpen incentives further, the government linked performance on these star ratings to the degrees of freedom to be given to hospital managers.

One striking feature of these ratings was that NHS hospitals in the North of England got more stars than those in the South. Why should there be a regional pattern to the star ratings with hospitals in the North performing better on average than those in the South? It cannot be due to differences in the health of the population since, in general, this is better in the South than the North. Research undertaken by colleagues at the University of Bristol (Burgess et al 2003) suggests that differences in pay between hospitals located in high costs areas – basically the South – and those located in low cost areas – basically the North – may be a key factor driving the differences in star ratings.

Pay in the NHS is still primarily set centrally. Hospitals in London and the South East pay more than hospitals located elsewhere. All other hospitals pay the same amount to staff at a given grade and experience. But wages in the private sector are not fixed in this way. They reflect the state of the local economy. Areas of high demand for labour will have high private sector wages; areas with low demand will have low private

sector wages. And since the pay differences between regions *within* the NHS do not fully reflect these private sector wage differentials, the differential between NHS pay and private sector pay varies considerably across the country. Hospitals in high cost areas are likely to pay less compared with the private sector than hospitals in low cost areas. This means that hospitals located in high cost areas – those where wages outside the NHS are high – face more competition for staff than hospitals located in low cost areas. Low relative pay can lead to two possible shortages: staff and quality. It can cause problems in hiring and retention, which can affect productivity and quality.

Nurses are key staff in the NHS. The NHS faces a shortage of nurses: problems of recruitment and retention are widespread and are particularly an issue in the more prosperous South. Research has shown that the number of people working in nursing is related to pay, as well as other aspects of working as a nurse, such as shift work⁶. We examined whether the gap between what nurses are paid and what they might earn outside the NHS has an impact on the star ratings – a measure of the quality – of NHS hospitals. Unpacking the star ratings into their separate components reveals that hospitals located in areas where the outside options for nurses are good have poorer outcomes in several areas.

- For the 2001 targets, the nurses' pay gap is associated with 7 of 21 targets. 4 of these are 'key' targets – the number of patients waiting for an inpatient appointment, and measures of outpatient waiting, trolley waiting and hospital cleanliness.
- For the 2002 targets, the pay gap is associated with 11 out of the 37 targets. 10 of these are performance targets, 6 based on inpatient surveys.
- Across the two years, over a third of the individual targets (18 out of 53) are negatively associated with the pay gap. Many of the targets associated with the pay gap are based on patient surveys. Patient experiences depend, to a large extent, on their contact with nursing staff. The fact that these outcomes are rated as better where staff are paid relatively more points to a link between the ratings and pay.

Recent research looking at another measure of hospital quality – the number of deaths following admissions of individuals with heart attacks - supports these findings. Hospitals located in areas where the NHS pay is poor compared to the private

⁶ Doiron and Jones (forthcoming) examine the factors behind problems of nurse retention in NSW.

alternatives have higher death rates, holding everything else constant. We also find that hospitals located in these areas have lower productivity.

The results suggest that public-private pay differentials matter and that some of the performance of NHS hospitals may be related less to managerial ability than to the labour market in which the hospital is located. This, in turn, suggests that allowing managers to raise pay in areas where it is relatively low may be one way to improve NHS productivity.

4. Conclusions

My aim in this talk has been to try to convince you that individuals employed in the provision of public services do respond to financial incentives. This does not mean that they lack what might be called ‘public sector motivation’: it is clear that teachers, nurses, doctors, job placement officers care about the provision of good quality services to their students, patient and clients. But this is not to the exclusion of also caring about financial rewards. It seems clear from the evidence that individuals located deep in the public sector do respond to targeted financial incentives, often quite small ones, and often changing their behaviour quickly. In addition, they also appear to respond (negatively) to underpayment relative to their options in the private sector.

There are caveats. Most of the evidence comes from either the operation of schemes that ran for relatively short periods or the assessments only covered short periods. We therefore do not know whether the effect of small payments would wear off over time. We also know that schemes that run over a long time period are plagued by ‘ratchet type’ effects – the linking of present performance to past performance means people lower their effort in order to make the initial performance as low as possible (there are countless examples from the former Soviet Union of such behaviour). It is also clear that schemes have to be carefully designed; the problems of measurability and multi-tasking meaning that unintended consequences may easily arise, some of which may actually reduce rather than increase productivity. Financial incentives alone are also not enough to induce productivity increases: introducing a scheme into an organisation in confusion about its objectives and mission is very unlikely to improve matters, and may well make matters worse.

But in sum, the lessons appear to be that financial incentives are one tool that government seeking to raise public sector productivity may use. More generally, a blurring of the traditional distinctions between public and private sectors is important for the success of government’s attempts to improve our health care, our education and our government administration.

References

- Adele Atkinson, Simon Burgess, Bronwyn Crosson, Paul Gregg, Carol Propper, Helen Slater and Deborah Wilson (2004) Evaluating the Impact of Performance-related Pay for Teachers in England. University of Bristol, CMPO Working Paper 04/113.
- Asch, B (1990) Do Incentives Matter? The Case of Navy recruiters, *Industrial and Labor Relations Review* 43, 89-107.
- Besley, T and Ghatak, M (2003) Incentives, Choice and Accountability in the Provision of Public Services. *Oxford Review of Economics Policy* 19,2, 235-249.
- Burgess, S and Ratto, M (2003) The Role of Incentives in the Public Sector: Issues and Evidence. *Oxford Review of Economics Policy* 19,2, 285-300.
- Burgess, S, Gossage, D and Propper, C (2003) Explaining Differences in Hospital Performance: Does the Answer Lie in the Labour Market? CMPO Working Paper No. 03/091. <http://www.bris.ac.uk/cmpo/workingpapers/wp91.pdf>
- Burgess, S, Propper, C and Wilson, D (2005) Choice: The Implications of the Economic Research. CMPO, University of Bristol March 2005.
- Courty, P and Marschke, G (2004) An Empirical Investigation of Gaming Responses to Explicit Performance Incentives *Journal of Labour Economics* 24, 1, 23-56.
- Doiron, D and Jones, G (forthcoming) Nurses Retention and Hospital Characteristics in New South Wales. *Economic Record*.
- Dusheiko, M, Gravelle, H and Jacobs, R (2004) The Effect of Practice Budgets on Patient Waiting Times: Allowing for Section Bias, *Health Economics* 13, 10, 941-958.
- Francois, P (2000) Public Service Motivation. *Journal of Public Economics* 78, 275-99.
- Gaynor M, (2004), 'Competition and quality in health care markets: what do we know, what don't we know?' Paper commissioned by the Federal Trade Commission, mimeo, Department of Public Policy, Carnegie Mellon University, Pittsburgh.
- Gaynor, M, Rebitzer J B and Taylor, L J (2001) Incentives in HMOs, Economics Working Paper Archive 340, Levy Economics Institute.
- Holmstrom, B (1982) Moral Hazard in Teams, *Bell Journal of Economics* 13, 324-40.
- Hoxby, C (2000) Does competition among public schools benefit students and taxpayers? *American Economic Review* 90(5), 1209-1238.
- Lavy, V (2002) Evaluating the Effect of Teachers Group Performance Incentives on Pupil Achievement. *Journal of Political Economy* 111, 1286-317
- Kahn, C, Silva, E and Ziliak, J (2001) Performance-based Wages in Tax Collection: The Brazilian Tax Collection Reform and its Effects. *The Economic Journal* 111 (486), 188-205.
- Lavy, V (2003) Pay for Performance: the Effect of Teachers' Financial Incentives on Students' Scholastic Outcomes' CEPR Discussion Paper No. 3862.

- Le Grand, J (1991) Quasi-Markets and Social Policy. *The Economic Journal* 101, 1256-67.
- Propper, C, Wilson, D and Burgess, S (forthcoming) Extending Choice In English Health Care: The implications of the economic evidence *Journal of Social Policy*.
- Propper, C (1995) Agency and incentives in the NHS internal market, *Social Science and Medicine*, 40 (12) 1683-1690.
- Propper, C (2005) Why Economics is Good for your Health, *Health Economics* 14 (10), 987-998.
- Propper, C and Wilson, D (2003) The Use and Usefulness of Performance Measures *Oxford Review of Economic Policy* 19 (2), 250-267.
- Propper, C, Crosson B, and A Perkins (2001) Do Doctors Respond to Financial Incentives: UK family Doctors and the GP Fundholder Scheme *Journal of Public Economics* 79, 375-398 (2001).
- Rochaix, L (1993) Financial Incentives for Physicians: the Quebec Experience. *Health Economics* 2, 163-176.
- Scott, A (2000) Economics of General Practice in Culyer, A J and Newhouse, J (eds) *The Handbook of Health Economics* Vol 1B. Amsterdam: North Holland