Just Deserts?[1]

Risk, Reward and Incentives in Criminal Justice

[redacted]

Since 2010, the Ministry of Justice (MoJ) has been pressing ahead with plans radically to reform offender management. The first aim of the reforms is to help the Department save money to meet its tough budget settlement,[2] under which it will be spending around one third less in 2015-16 than it was in 2010-11.[3] The second is to introduce a radically new way of commissioning services to improve the productivity of offender management services, and therefore cushion the blow of falling departmental spending.

In May, the MoJ produced the command paper *Transforming Rehabilitation: A strategy for reform*.[4] The eye-catching part of the proposals is the aim to encourage non-state providers to risk private capital on innovative services to reduce reoffending.[5] To this end, a proportion of the £3.5bn total contract value will be made on a 'payment by results' (PbR) basis: payments will depend on the subsequent reoffending rate in each contractor's area. This is the first time that such an approach to commissioning rehabilitation services has been tried at scale anywhere in the world.

The government has set much store by the capacity for PbR substantially to reduce the level of recidivism. This paper focuses on the payment mechanism proposed by MoJ in its May document *Payment Mechanism – Straw Man,* and explores whether it will in fact give providers incentives to cut reoffending.

We begin by briefly outlining the case for PbR in principle, and why the approach is seen as such an important policy innovation for boosting the stagnant productivity of public services more broadly. We go on to assess how far the current MoJ proposals fulfill these goals in practice, using financial simulation to analyse the risk-adjusted returns that investors might expect.

The results of our simulations show that, far from creating incentives to cut reoffending, the proposed regime will encourage providers to cut costs and allow reconviction rates to drift upwards. As a consequence, on current plans the scheme will do nothing to encourage innovation and improved value for taxpayers' money. This threatens to derail an important new approach to commissioning that could have wider applicability in the delivery of public services, both here and abroad. The paper concludes by offering some solutions.

Why payment by results?

For decades, the provision of public services has been hampered by central direction from Whitehall. In offender rehabilitation, as much as in employment services, professionals at the front line of service delivery are often unable to deploy their superior knowledge of their clients and their circumstances to help achieve the outcomes that everyone wants: desistance from offending behaviour in the former case, or sustainable employment in the latter. The result, in many cases, is a mediocre service that comes to be characterised by a disempowering tick-box culture for service users and professionals alike. The scope for innovation and productivity growth is very limited in such a top-down system. The reason for the hardiness of the status quo lies in the difficulty of designing a cost-effective alternative.

---

[1] This is an abridged and updated version of the paper *Paying for Results? Rethinking Probation Reform*, published by the Social Market Foundation, August 2013.
[2] The Ministry of Justice disputes the assertion that savings are part of the reform agenda.
[3] Gemma Tetlow, "Post-spending round briefing, 27 June 2013", Institute for Fiscal Studies, www.ifs.org.uk/budgets/sr2013/gemma_tetlow.pdf .
[4] Ministry of Justice, Transforming Rehabilitation: a strategy for reform (London: HMSO, 2013).
[5] Throughout this paper 'reoffending' is taken to mean 'proven reoffending', equivalent to reconviction.

Devolving much greater autonomy to front-line professionals carries two great risks. The first is that, as in any walk of life, there is variation in the skill, capability and application of people delivering the services. Centralisation is, in effect, an insurance policy against excessive variability in quality, albeit one with the unhelpful side-effect of limiting the up-side as well as the down-side of professional autonomy. The second risk is that autonomy without accountability is a recipe for spiraling costs and poor value for money – either in the services themselves or elsewhere in the public budget – since those spending the money have little responsibility for the outcomes achieved.

Conventional competitive tendering - where would-be providers bid to operate the service in question – offers no solution to these fundamental problems. Pure 'fee for service' contracts may cut costs in a static sense, but the need tightly to prescribe the services that are to be delivered offers no greater autonomy to front line professionals.[6] They therefore do not enhance dynamic efficiency.

Significant operational autonomy at the frontline is only possible if it is combined with appropriate financial accountability. For Whitehall policymakers, wherever the autonomy lies, accountability must go with it. This reframes the challenge for public services reform into one of seeking ways to devolve financial risk: if it can be devolved effectively then so can autonomy over provision.

Attaching payments to outcomes – where outcomes can be clearly defined and measured – therefore offers a powerful mechanism to achieve the devolution of operational autonomy and financial risk together. This is the theory behind paying by results. If such models can be made to work, their potential for boosting the productivity of public services could be substantial. At a time of deep public spending cuts, this goal is hugely important.

The reform plan

Under the Department's outsourcing proposals, the bulk of providers' contracts will involve a conventional 'fee for service' (FFS). But the new arrangements will also involve a significant PbR element to try to encourage innovation and better outcomes for a given amount of money.[7]

Chart 1.1 below outlines some key parts of the proposed PbR mechanism. If reoffending rates remain at their historical level (the 'baseline') in a given area of the country, or Contract Package Area (CPA), the provider operating there will receive no additional payments above the fee for service. If reoffending outcomes in its area are marginally better than the baseline the provider will not receive any success payments, since improvements could be the result of random variation (or 'statistical noise'), rather than superior underlying performance. The same is true if performance dips somewhat below past levels: no change in payment.

Only once providers cut reoffending by a statistically significant amount will they see a jump in revenue, as they receive outcome payments in respect of each additional offender who desists above the baseline level. Equally, if reconviction rates rise by a significant amount, providers face penalties. Additional payments are made for each extra non-offender above the statistical 'zone of uncertainty', before reaching a cap for providers who perform very well. So why has the Department designed the payment schedule in this way, and how will it affect incentives?

---

[6] DeAnne Julius, Public services industry review (London: Department for Business, Enterprise and Regulatory Reform, 2008).

[7] Ministry of Justice, "Rehabilitation programme - Payment Mechanism Straw Man", http://www.justice.gov.uk/downloads/rehab-prog/payment-mechanism.pdf.

Chart 1.1: Proposed PbR payment curve

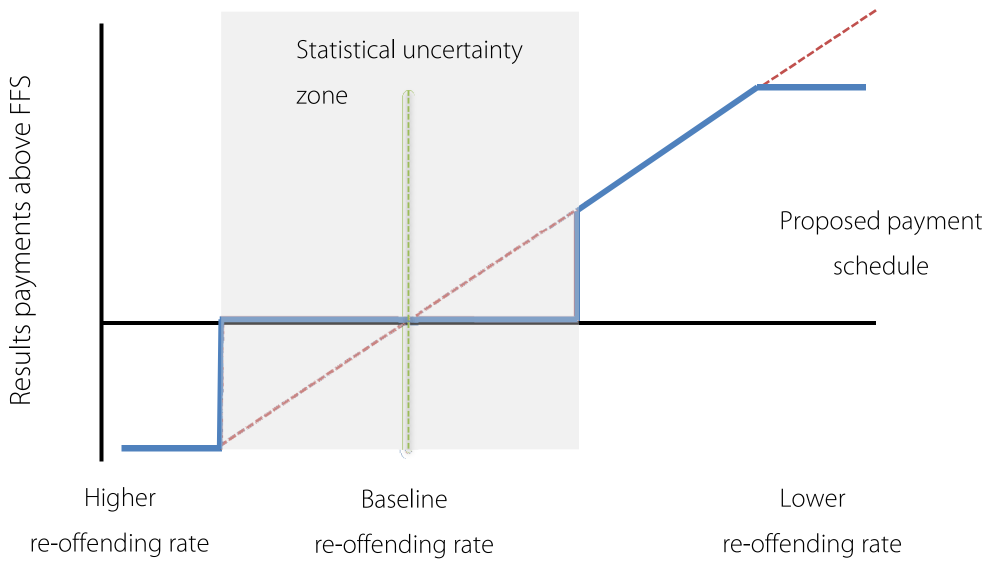

Statistical risk and incentives

Statistical uncertainty poses a big challenge for payment by results regimes in any policy area. If cohorts of offenders are small, we would expect to see a high degree of random variation in recidivism rates between different providers offering exactly the same quality of interventions. The larger the cohort, the smaller the degree of random variation we would anticipate, making it easier to identify good or bad provider performance with a high degree of statistical certainty.

If government were to pay a provider for every additional person who doesn't reoffend, above the historical reoffending baseline, MoJ fears that this would create a financial risk for the taxpayer. This is because random variation in measured reoffending rates raises the possibility that government could pay for a reduction in reoffending in a CPA that *looks* like good performance but is in fact nothing more than the result of statistical good fortune. In any one CPA this could be a risk. However, on that payment approach, it is equally likely that government will recoup payments from other providers who *look* like they have done a bad job. So across all of the CPAs this risk to the taxpayer is negligible.

Nevertheless, to avoid paying, or penalising, providers for this statistical 'noise' in any given CPA, the MoJ's proposed payment schedule involves a flat payment zone either side of the baseline reoffending rate. Over this range of reoffending outcomes, payments to providers are fixed. The width of this zone is determined by the degree of statistical uncertainty associated with measuring performance.

The Straw Man document suggests that for annual cohorts of offenders, a three percentage point reduction (or increase) in the reoffending rate would be needed to be reasonably confident that an underlying improvement had been made (or deterioration had occurred).[8] In other words, over a six percentage point zone around the baseline – three points either side - no additional payments will be made, or penalties levied, on providers for their performance. We therefore take this as the payment structure for our central case simulation.

This poses big problems for the scheme and provider incentives. Even if providers do make an

[8] "Rehabilitation programme - Payment Mechanism Straw Man", 9.

underlying improvement of, say, four percentage points, the measured figure could still come in below the three percentage point threshold required to receive outcome payments.

In seeking to guard against paying out for spurious reductions in reoffending, this approach also means providers that do succeed in improving the underlying service often won't be paid accordingly. On the other hand, providers who perform poorly, seeing underlying reoffending rise, are protected from penalties by the uniform payment in the statistical uncertainty zone. All of this has the effect of damaging incentives for providers to improve.

If reducing reoffending was a costless activity, this dynamic in the payment regime would simply mean weakened incentives for providers to cut reoffending. But since providers will actually have to spend their own money on services to achieve such improvements, the proposed payment structure could leave providers worse off than if they had spent nothing at all. Unless providers can be confident of achieving a very large reduction in reoffending, and therefore securing some reward payments, spending money of rehabilitation services looks very risky under this payment structure.

To assess the scale of the financial risks involved, we modeled the incentives for providers and investors. The simulation assesses the degree of statistical uncertainty around the measured outcome from a given level of spending on services. The model uses this to determine providers' expected (i.e. risk-adjusted) gross profit (i.e. revenue less the costs of delivering the interventions necessary to secure that revenue), given the MoJ's proposed payment structure for the PbR part of the contracts, under a set of assumptions (outlined in the Annex). The expected gross profit for a provider is determined by three factors:

1. the cost of the interventions made – more intensive interventions will cost more but should also reduce underlying reoffending by more;
2. the level of outcome payments that will result from cutting reoffending by a given number of percentage points; and
3. the probability that statistical noise will leave the measured desistance rate above or below the underlying rate.

Chart 1.2: Investor perspective on PbR offer: risk-adjusted gross profit from different levels of intervention.

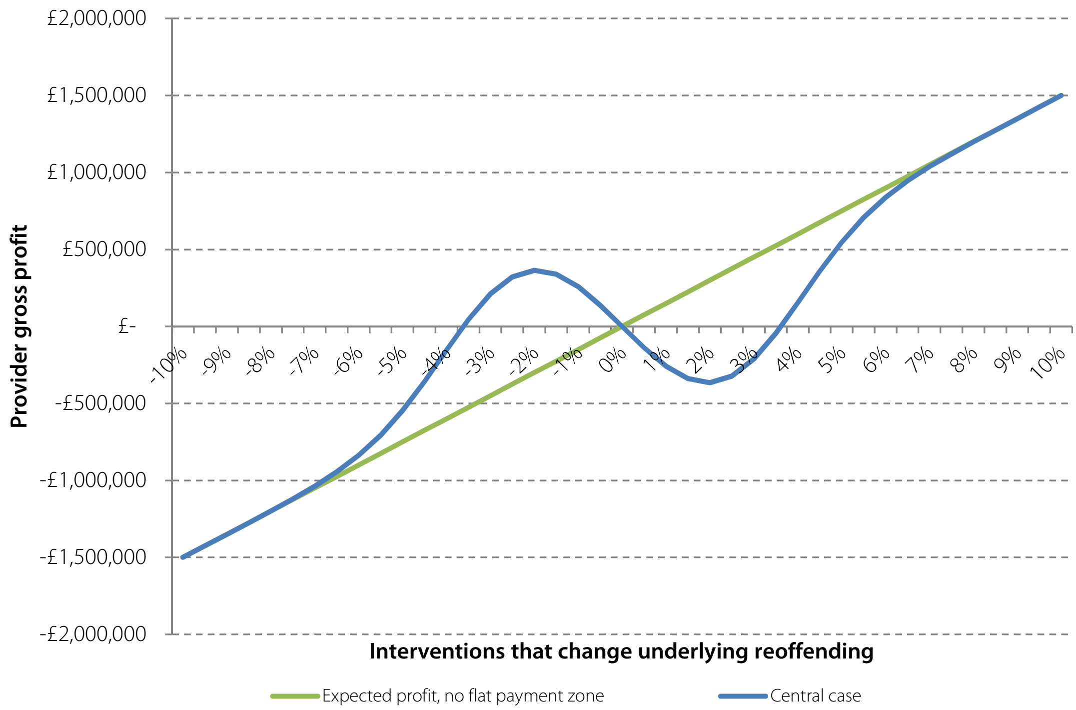

Chart 1.2 illustrates the results. The horizontal axis represents a menu of intervention levels that the provider can expect will shift the underlying desistance rate by the percentages shown. For example, if the provider chooses to hold services as they are, not investing more to cut reoffending, they can expect the reoffending rate to remain at the baseline 0%. There is a possibility that, by good or bad statistical fortune the measured reoffending rate will be higher or lower than the baseline, but those possible outcomes are equally likely and therefore cancel each other out in terms of what revenue this provider can expect (in probabilistic terms). Hence at the 0% intervention, expected gross profit for the provider is nil.

If, on the other hand, the provider decides to invest in interventions that are thought to cut reoffending by three percentage points, for example, the expected gross profit will change as shown.

From an investor's perspective, the implications of payment schedule set out by the Department are represented by the blue line. This shows the risk-adjusted gross profit from different interventions, and demonstrates that the incentives on providers are perverse. Since no payments or penalties are made for providers whose measured reoffending rate lies within the zone of statistical uncertainty, only a measured reoffending reduction of more than three percentage points will result in a payment. It shows that a provider has to aim for almost a four percentage point reoffending reduction to be confident of getting paid.[9] This is a very large hurdle for providers to overcome.

In this world, no investor would target a reduction of reoffending of, say, two percentage points, since the expenditure involved in achieving it is very unlikely to be rewarded with outcome payments, leaving the investor with a loss. In other words, efforts to cut reoffending will effectively be penalised.

Far from encouraging service enhancement, the flat payment zone in fact has the perverse effect of incentivising providers to cut spending on reducing reoffending. Again, because increases in reoffending are not penalised until they reach the point of being statistically significant, this encourages providers to strip down their service, thus saving money, and raising expected gross profit. In the example here, the profit maximising strategy is to cut services to the point where reoffending rises by two percentage points. Efforts by providers to save money, pushing up reoffending, will effectively be rewarded under this proposed payment structure. Meanwhile, funding rehabilitation interventions under this scheme is more than likely to result in losses. At the margin, the proposed scheme penalises good performance. Needless to say, this makes no sense from the perspective of either the taxpayer or the victim of crime.

By contrast, the green line in Chart 1.2 illustrates how that expected gross profit would vary under a simple alternative payment schedule where a per-person payment (or penalty) is made for each offender above or below the reoffending baseline. On this approach, without the statistical significance zone that features in the MoJ proposal, the expected gross profit unambiguously improves for every extra pound invested in rehabilitation services. It would set clear incentives that the more a provider cuts underlying reoffending the more likely they are to make a profit – a central attribute of any PbR scheme worthy of the name.

Alternative Scenario 1: Smaller cohorts – bigger risks for providers

The statistical uncertainty zone is likely to be bigger in areas with smaller cohorts, such as the CPAs of Essex and South Yorkshire, and smaller in the largest CPAs, such as London, which will

[9] Note that this is a higher degree of improvement than the 3ppt Target Level required for payment. Were the provider to target a 3ppt improvement exactly, they would face a 50% chance of being paid nothing (if the measured outcome slipped below 3ppts) despite having spent money on interventions. For this reason, on a risk-adjusted basis, it only makes sense for providers to target improvements well above the statistical significance zone.

have a cohort five times bigger. This means that the problem identified above could be less severe in London, but more severe in the smaller CPAs, the majority of which will have caseloads less than one third the size of London's.

Chart 1.3, below, illustrates how the disincentive for providers to invest in cutting reoffending changes according to the degree of statistical uncertainty one might expect in smaller and larger CPAs. If the zone were to narrow to plus/minus 1.5 percentage points in the larger CPAs, the expected gross profit curve would look like the red dotted line on the chart. If the zone were to widen to plus/minus five percentage points in the smaller CPAs, the expected gross profit curve would look like the dotted purple line on the chart. Historical data from probation trusts suggests that this degree of variation between CPA would not be unusual. The perverse incentives in smaller CPAs could therefore be huge.

Chart 1.3: Risk-adjusted gross profit and cohort size

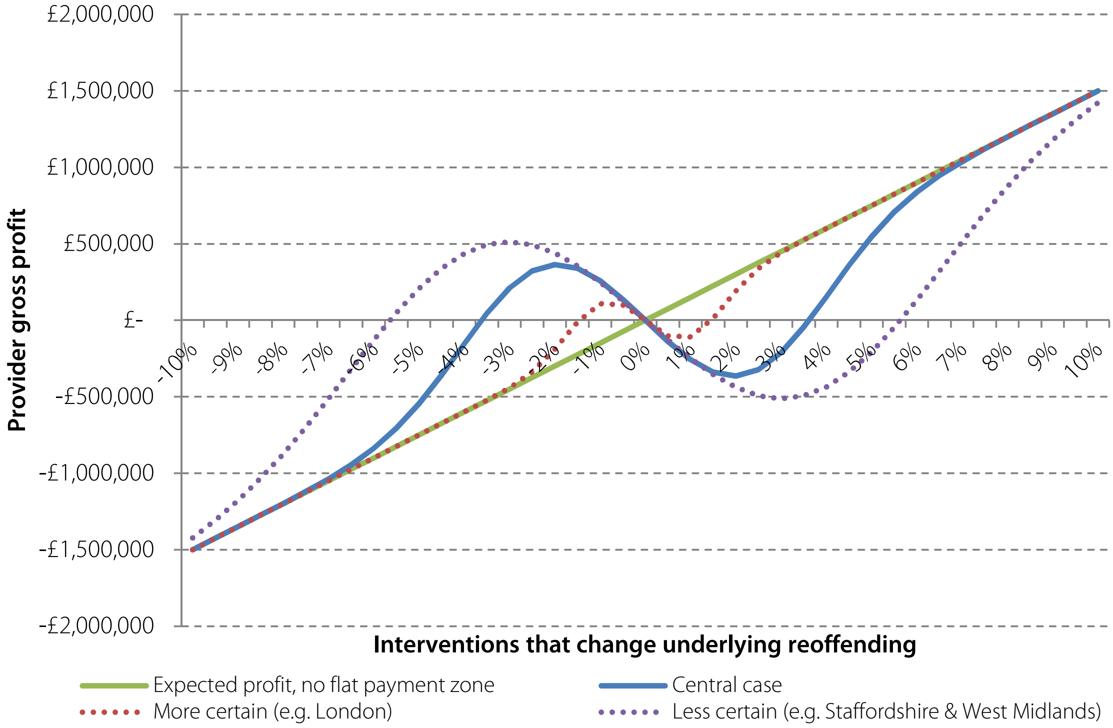

Alternative Scenario 2: Intervention costs rise for harder cases
The analysis thus far has assumed that the costs of services to reduce recidivism are uniform for each extra offender who desists. In reality the cost of successful interventions is likely to rise as services begin to target harder cases. This is an intuitively reasonable assumption: while getting some offenders accommodation and a job might help them desist, for harder cases a much wider range of services – substance abuse, training, mental health services etc. - might be required, costing substantially more.

The orange line in Chart 1.4, below, shows what the risk-adjusted gross profit curve would look like for providers if they think that costs of achieving desistance rise gradually as they move through the cohort. In this model the marginal cost of desistance is assumed to go on rising by £1,000 per offender for each half per cent of the cohort. This means it is assumed that, among the first half percentile (50 of the 10,000 offenders in this CPA), providers have to spend just £1,000 per person to achieve desistance. For the next half percentile, the costs of successful interventions rise to £2,000 per person, and so on. The total cost of reducing underlying reoffending by one percentage point is therefore £150,000 (£50,000 plus £100,000).

The consequences of this ‘rising marginal cost of desistance’ are dire for provider incentives. As the orange line illustrates, under these assumptions there is no prospect of being rewarded for cutting reoffending. This model is assumption driven, and different bidders will make different assumptions about how the cost of services varies. But in any case, rising marginal costs raise the question of whether there is any point where the risk adjusted gross profit of trying to cut reoffending turns positive under these proposals. The green line alternative, demonstrating positive incentives, again shows that this situation is not an inevitable feature of PbR.

Chart 1.4: Risk-adjusted gross profit assuming rising marginal cost of desistance

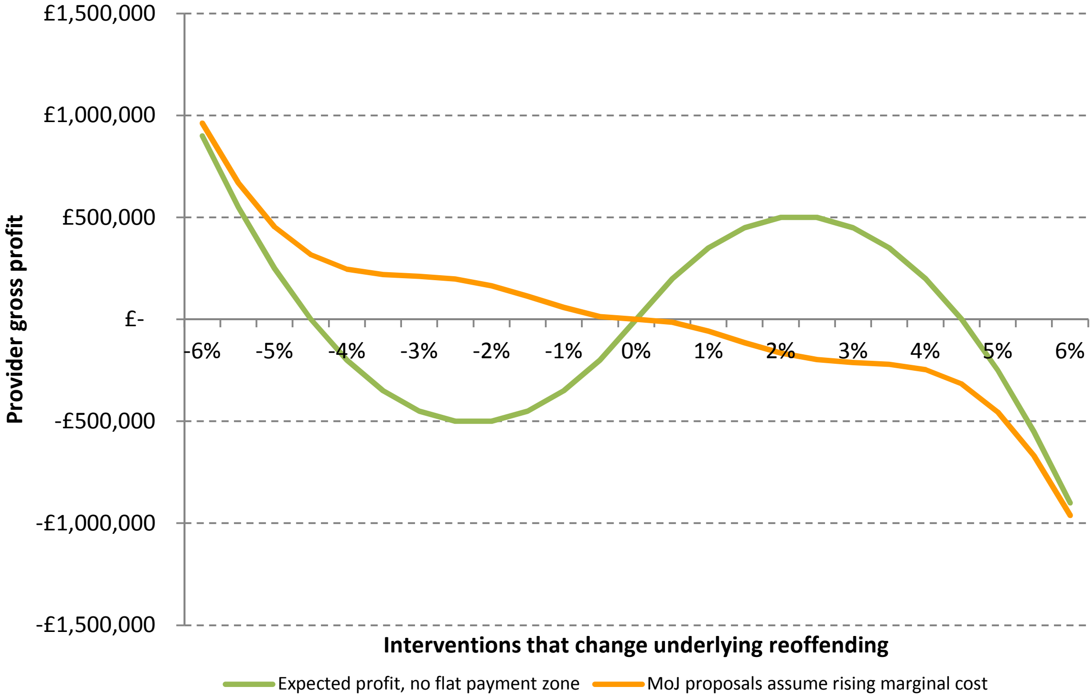

Sufficient resources?

Under our central case above, it is nevertheless possible to make a positive return on investment if really substantial in-roads into reoffending can be made. That analysis begs the question: what kind of reoffending reduction is achievable under this kind of scheme? The social impact bond pilot at Peterborough prison offers an indication.[10]

In the period from October 2008 to June 2010, prior to the start of the Peterborough pilot, one-year reconviction rates from the prison averaged 40.3%, 1.1 percentage points above the national reconviction rate of around 39.2%. Since the start of the pilot, the one-year reconviction rate, at 38.6% has been slightly below the national figure of 40.4%. A proper evaluation has yet to be completed, and these raw figures take no account of any change in the client mix in Peterborough Prison.[11] Nevertheless, this suggestion of a 2.9 percentage point reduction in reconviction rates gives a sense of the scale of improvement that might be achieved.

Whether this is achievable under the national scheme is highly questionable given the relative level of resources however. Under the Peterborough pilot, £5m is being spent on services to 3,000

[10] In the Peterborough pilot, performance is measured on the frequency of reoffending behaviour, rather than the 12 month reconviction rate proposed as the central measure for the Transforming Rehabilitation programme. Nevertheless there is obviously a strong connection between reduced frequency and desistance.
[11] The OGRS scores – which predict likelihood of reoffending - for Peterborough offenders appears to have changed by only a little since before the start of the scheme. The Peterborough caseload seems to have become only marginally less likely to re- offend (absent any different interventions), relative to the national cohort, since the start of the programme.

offenders, averaging just under £1,700 per person. By contrast, the Department's recent *Target Operating Model* suggests that the outcome payment fund in most CPAs will be around £400 per offender.[12]

With around one quarter of the resources per head, there is no reason to think that providers in the proposed national scheme will be able to do better than the Peterborough pioneers. With limited resources available, it is even more essential that provider incentives at the margin point in the right direction if the programme is to succeed.

Conclusion

The payment mechanism proposed by the MoJ is the key tool by which they hope to transform rehabilitation. Success here could also open the way to wider use of the approach to raise the productivity of public services both in the UK and elsewhere.

But as the plans stand, providers have incentives to cut costs and marginally increase reoffending rather than reduce it. This is not a scheme that will pay for results. Indeed it will likely penalise providers for trying to achieve them. These payment proposals, far from improving the productivity of probation in order to cushion the blow of falling spending, risk reducing it.

The consequences of the current design threaten to undermine the entire policy and damage the credibility of PbR as a commissioning approach. A simple solution is available, however. We propose that the payment curve be simplified to pay (or penalise) providers on a per-person basis for any observed reduction (or increase) in reoffending. The financial risk in any one CPA from dropping the proposed 'flat payment zone' would be negligible at the national level. The result would be better incentives for providers, substantially better value for money for taxpayers, and proof of concept for an approach that is perhaps the most important policy innovation in a generation.

[12] *Target Operating Model*, Ministry of Justice, September 2013, 60.

Annex – Simulation assumptions

- A cohort of 10,000 offenders – the average annual cohort size across the 21 CPAs will be around 11,000.
- A standard deviation of 1.5 percentage points' reduction in the measured reconviction rate. This assumption is based on the figure used in the Straw Man document for an annual cohort. MoJ officials have provided some historic reconviction data by probation trust, which although rudimentary, tend to support this figure. Of course, it is likely that the standard deviation will be smaller in CPAs with large cohorts, and larger in areas with small cohorts.
- Each additional outcome above the baseline is assumed to result in a £5,000 payment to providers. This equates to an assumption that just over £120m is available for outcome payments and targeted at the marginal ten per cent of offenders – around 24,000 - above the baseline. It should be stressed that this assumption is little more than an educated guess at the level of outcome payments that may be offered. Nevertheless, the assumption is not central to the point illustrated here.
- Each additional outcome above the baseline is assumed to cost £3,500. In other words it is assumed that: outcome payments are set high enough for providers to be able to make a profit; and that the cost of successful interventions does not rise as providers get further into the caseload. Both assumptions are arguably generous to the MoJ proposal.
- It is assumed that there is no cap on the outcome payments on offer. In fact MoJ has said it will set a cap, limiting the up-side for providers. This is therefore a generous simplification of the MoJ proposal, made for ease of calculation.