

How was the Impact Assessment of the proposed AGFS done?

Summary

The purpose of this note is to provide advocates with more details of how the Ministry of Justice (MoJ) calculated the impact of the proposed AGFS. I review the overall method used and give details of the data that was utilised to produce a *model* that has been used to determine both the overall costing of the proposals and their effect on different groups of barristers.

Where it seems helpful I use examples drawn from the actual data used in the process.

Key points;

- The MoJ model uses detailed AGFS data for 105,900 cases that were paid in 2014-15.
- AGFS data can be used to place the majority of these cases into a single category under the proposed scheme.
- Cases that cannot be banded using existing AGFS data include murder, financial (fraud) and drugs cases.
- To classify these cases a review of case files was used and this produced estimates of what proportion of a given type of case would fall into the relevant new category.
- AGFS data also does not contain details of the additional hearings unless a particular case has more than five of them. However under the proposals these will be paid. Court records were used to establish the number of these hearings.
- The overall cost of the proposals was calculated combining these three sources of data; AGFS, file reviews and court records.

Method

It is helpful to start by stating that the proposals have been evaluated against the current AGFS scheme for *nearly all cases* paid for in the financial year 2014-15. I have used the qualifier “nearly” because some cases were excluded from the analysis and I will give the details below. The use of a complete year of cases is an important element of the method that the MoJ adopted, because it avoids the problems that arise from using just a sample of cases and it means that the assessment exercise is extensive.

Conceptually the process of costing the new proposals is straightforward. It is to take each case paid in 2014-15 and compare the actual payment with an estimate of how much that case would be paid under the proposed new scheme. *Practically* difficulties arise because the most extensive data available for this exercise are the payment records under the current AGFS and those data do not have all of the information necessary to determine how much a case will be paid under the new proposals.

There are two important omissions from AGFS payment data in this regard. First, the data do not contain all of the information required by the new classification system. In addition to the AGFS offence code, the payment data do contain more information on the offences on the indictment which helps classifying a lot of cases under the proposed scheme. But elements that are important for certain classifications (for example the weight of drugs, the aggravating circumstances of serious violence offences and the value of dishonesty offences) and which determine the *band* into which a case falls are not recorded. Second, the present scheme includes within the main fee a number of hearings and so the details of these are not recorded, and an important aspect of the proposals is to unbundle the payment for these hearings.

To properly understand how these gaps in the data were filled it is necessary to explain in more detail the underlying AGFS data and how it can be used and augmented.

AGFS Data

As noted above the data that feeds into the assessment process is extensive. The starting point for the MoJ analysis is 112,900 payments for *cases* completed in the financial year 2014-15. Some of those payments relate to rather old cases (paid under previous incarnations of the AGFS), or are incompletely recorded (these are typically for very small sums of money and are usually ‘corrections’ to previous payments). Both of these were excluded, the first because they are not payments on current AGFS rates and the latter because there are insufficient details to form a sensible estimate of how they will be treated under the new proposals. Around 7,000 payment are excluded for these reasons, so that the analysis I am going now discuss concerns 105,900 cases – a considerable number.

Cases are classified according to what the system records as their “Bill Scenario”. Advocates will recognize these as corresponding to the type of work they are instructed to perform. The following table reproduces a part of Table 5 from the Impact Assessment, concentrating on the scenarios that make up the majority of cases.

Bill Scenario	Expenditure £m (2014-15)	Number of Cases
Trial	98	20,200
Cracked Trial	28	21,300

Guilty Plea	27	38,400
Committal for Sentence	2.2	12,600
Elected cases not proceeded	0.7	3,600
Breach of Crown Court Order	0.6	4,100
Other	4.6	5,800

The table shows, what you would expect, that in terms of expenditure the most important Bill Scenarios are Trial, Cracked Trial and Guilty Pleas, but it is important to see the whole picture and recognize that advocates receive a large number of payments for Committals, Breaches and so on.

Each of those cases has a large amount of data associated with it. If each case is thought of as a row in a spreadsheet then there are lots of columns. For example in the spreadsheet used to actually carry out the calculations there are 48 columns of data. These include the breakdown of how the case was paid under AGFS and the supporting data – such as the number of pages, the role of the advocate (QC, Leader or Led), the trial days, the court, the case identifier and so on.

Case Mapping

The greatest challenge to taking this large spreadsheet of data and imputing what every case would be paid under the proposals is the new case classification system which was determined from first principles rather than from the existing AGFS architecture. To address this challenge the MoJ sought to establish what they called a *mapping* from every case in the data to its corresponding classification in the proposed scheme.

In some instances the definition of the classification and the information in the data file is sufficient to determine one (and only one) category that the case would fall into. In other instances categorization does not actually matter – since a committal for sentence for example is paid the same whatever the category of the case. So the problem of case categorization can be broken down; find those cases where the category can be exactly determined, find all those cases where the category does not affect payment and finally focus on those cases that are left. The remaining cases are ones for which the data provides *some* indication of which category the case will fall into but not an exact indication. Such cases are mapped to multiple categories, usually but not always these are different bands within a given category.

These multiple category cases are the ones where most effort was expended. It is useful to report them in some detail so that you can understand both the extent of the issue and how it was addressed. The table below shows the most important (in terms of number and value of cases) multiple categorisations.

Possible Categories into which a case could fall							Percentage of all cases (approximate)
1.1	1.2	1.3	3.1	3.2			2%
5.1	5.2	5.3					2%
8.1	8.1	8.3	8.4	8.5	8.6	8.7	13%

The first row relates to cases of murder and attempted murder where the proposed categorization looks at various aggravating features that are not a part of the current AGFS scheme. The second row relates to dishonesty offences and the financial value of dishonesty over £100,000. The third row are drugs cases where the proposed classification relies on the weight or volume of drugs and where that is not currently recorded.

Determining how to resolve into which particular category the cases in the table would fall constituted a major piece of work for the MoJ. The approach taken was to determine how likely a case in each row of the table was to fall into each of the respective columns. Equipped with these proportions it is possible to calculate the (mathematical) expectation of the cost of a case under the new proposals. Take a murder/attempted murder as an example. If we conclude such a case is equally likely to fall into each of categories 1.1,1.2,1.3,3.1 or 3.2 its expected cost is just the average of the cost of that case if it were in each of these categories.

It should be noted that this approach does not reveal the exact cost of a given case – it gives an average for that case, but when aggregating over a large number of cases that will give a reasonable estimate of the total cost of a group of cases which is exactly what is required of the costing model.

Augmenting Data - Case File Review

A process of what was called *case file review* was used to determine the likelihood of cases falling into the different categories. As the name suggests this process involves examining individual cases files (sometimes paper records, and sometimes electronic records) for a sample of cases. The idea is to see what proportion of real murder cases would fall into category 1.1 or 1.2 or 1.3 and so on. This was a major undertaking for MoJ and involved a lot of resources. I cannot describe all the details here, but in summary a sample of cases was selected, the representativeness of the sample was assessed and then the sample was categorized. During this process it was discovered that it was sometimes crucial (and sometimes not) to distinguish between trials and guilty pleas (the latter typically being less serious cases and falling in lower categories) and the whether the case was assigned to a QC or not (the former being typically more serious).

The resulting proportions of cases obviously play a crucial role in predicting the cost of the proposed scheme and it is important to recognize that considerable effort was put in to establishing and checking these proportions. As with any such exercise there can be errors and biases but a number of checks and reviews were undertaken.

The proportions are built into the MoJ model. They are in the form of tables of numbers that are used in the calculations. An example of such a table is given below. I do not have the permission of MoJ to publish the actual numbers but I have had access to them as have the bar team who worked on the scheme. We have reviewed them and are satisfied that they are a fair reflection of the evidence available concerning categories.

In the illustrative table use symbols (the letters a – e) to indicate where numbers would appear.

Example of Results from Case File Review

Distribution of murder cases for QCs	Percentage of murder cases that will fall in this band or category
1.1	a%
1.2	b%
1.3	c%
3.1	d%
3.2	e%

There are a number of tables such as this relating to other roles of counsel and other multiple mapped cases.

The way in which these numbers affect the calculations is important. If the reality is that more cases fall into the higher categories than these figures suggest, the implication is that the model is underestimating the cost of the proposals and vice versa. It is worth noting, for example, that by far the greater majority of drugs cases were predicted to fall into the lowest payment categories. If that turns out to be incorrect, the model will have underestimated the impact of the proposed scheme and it will pay *more* than is estimated. Thus there is a risk for the MoJ, just as there are risks for the profession if the proportion of cases falling into the higher bands is lower than anticipated. To assess these risks, sensitivity analysis was conducted and is reported in the impact assessment indicating the impact on the overall cost of the scheme if these percentages were to vary.

It is worth emphasizing again that this assignment of cases to categories is based on evidence and has not been simply assumed by MoJ.

Augmenting Data - Additional Hearings

As suggested above a second difficulty in costing the new proposals was that additional hearings will be paid for but are not recorded under the AGFS system. The solution to this that the MoJ adopted was to link AGFS data on cases back to court records.

This provided a means of establishing the average number of hearings, in cases where there were less than five, broken down by offence and case outcome (guilty plea, crack, trial). Again, sensitivity analysis was conducted on these averages. Each case is therefore assigned its “fair share” of additional hearings making allowance for the differences between trials, cracks and guilty pleas in this respect.

Since courts have to account for their use of time, and because the scheme proposes to pay fixed fees for additional hearings (not dependent on the categorization of a case) this seems to be a reasonable approach. The only downside is that any one case may be paid more or less (depending on whether it had in fact more or less than its fair share of hearings) than the notional estimated cost of that case. But aggregating over all cases the overall cost should be accurately estimated. Actually since additional hearing costs are usually only a small proportion of the total cost of a case, this downside is not substantial in practice even from the point of view of examining individual cases.

Putting it all together

Once the AGFS data and the augmenting data collection exercises were complete the construction of a costing model is a simple (but rather large) number crunching exercise. For those used to using spreadsheets you can conceptualise the process as follows. For a row of data first establish the new category of case, then using information on the role of the advocate and other details *lookup* the relevant fees tables (sometimes looking up multiple categories and weighting them by the proportions from the case file review), next make an adjustment for the payment of additional hearings (looking up the relevant fees from tables) in order to produce an estimated cost of the case under the new proposed scheme. Repeat this 105,900 times.

There was in fact more to the finished exercise than this, including linking and matching the payments data with characteristics of the advocates undertaking the work to check whether there were unintended differential impacts across protected characteristics. There was also a good deal of post analysis considering different groups of cases or advocates.

What does cost neutrality mean and how can it be assessed?

Summary

The purpose of this brief note is to help advocates understand the notion of cost-neutrality as it has been applied in assessing the proposed AGFS scheme and to demonstrate that cost-neutrality cannot easily be assessed by considering small samples of cases.

Key points:

- Cost neutrality is defined relative to a particular set of cases. The MoJ has claimed that the proposed scheme is cost-neutral over all cases paid in 2014-15.
- Even if the scheme is cost neutral in that sense it may result in a substantial fee decreases or increases when applied to a different set of cases. This is true even when we consider large numbers of cases.
- Deviations from cost neutrality are inevitable if relatively small bundles (less than 1000) of cases are evaluated. This is true even if those bundles or cases are chosen truly at random.
- Smaller bundles of cases will exhibit larger apparent deviations from cost neutrality even when a scheme is truly cost neutral.
- Any real bundle of cases will be very unlikely to be a random selection from the full year's cases. This is true for any individual barrister or for a set of barristers in chambers.
- Bundles of cases that are skewed towards certain kinds of case (either by category or bill scenario) can exhibit very large apparent deviations from cost neutrality.
- By way of illustration, comparisons of chambers have revealed variation in their gains or losses under the proposed scheme of more than +/- 6%.
- Therefore, it is important to be cautious inferring that the MoJ calculations are incorrect based on the evidence of a few hundred cases.

Cost neutrality

The MoJ model only establishes that the scheme being proposed will cost the same as the present AGFS for a given set of cases – those paid in 2014-15. So cost neutrality is a narrow notion. A different set of cases to those used in the model may or may not be cost neutral under the proposals.

I have analysed the MoJ model extensively. It often comes as a surprise to those unfamiliar with these data that even relatively large numbers of genuinely randomly selected cases may exhibit large deviations from cost neutrality. I will try and give an indication of this using some examples. As already stated, the MoJ model predicts cost neutrality so that if we take all 105,900 cases the estimated cost of those cases under the new proposals is the same as their cost under the current AGFS. So *all* the calculations you see below relate to a cost neutral scheme by definition. I am assuming *for the sake of the calculations* that the model is correct, and I am examining what happens when I take samples of cases.

First I took a genuine random sample of 1000 cases. Here is one result;

Number of cases	Proposed £m	Current £m
1,000	1.6	1.7

This is a 4% reduction in cost and whilst it is quite rare for the difference over 1000 cases to be that large it can happen. A different random sample could produce a 4% increase in cost, or cost neutrality. I have done this for 1000 cases – that is a lot by most standards and many more than it is feasible to do by hand or by reference to case files. If we could only do 100 cases it would be easy to construct examples where the swings around cost neutrality would come out as +/- 10% or more. Even if we take repeated purely random samples the figures can come out skewed. The following table reports 10 random samples of 100 cases each.

Proposed £m	Current £m	Difference in %
0.24	0.25	-4.0
0.14	0.15	-6.7
0.18	0.18	0.0
0.145	0.14	3.6
0.21	0.22	-4.5
0.19	0.191	-0.5
0.16	0.18	-11.1
0.164	0.183	-10.4
0.189	0.184	2.7
0.15	0.16	-6.3

If you look at the final column, this looks to be dominated by negative figures – indicating that for that random sample the proposed scheme is a cut. There are 7 negative figures and only 2 positives (and one zero). Another 10 random samples however would produce a different pattern – perhaps with more, or even only, positive numbers. The essential point is that relative to the variation that exists in AGFS cases these are all small samples and the resulting figures are not a reliable guide to the overall impact.

This is true for random selection of cases but things are more extreme if cases are not randomly selected. Selecting truly at random is actually surprisingly difficult. Human beings naturally seek patterns and that leads them to select things that are notable, different or “interesting”. This is not the same as selecting at random.

We can easily understand what might happen by looking at the impact assessment. Suppose we skewed a sample towards guilty pleas. We know that on aggregate spend decreases for these cases and so any sample that is dominated by them will exhibit a cut – even if the scheme is overall cost neutral. The following table shows what happened when I took 5 samples of 100 cases weighted heavily towards guilty pleas.

Proposed £m	Current £m	Difference in %
0.65	0.71	-8.5
0.71	0.86	-17.4
0.67	0.8	-16.3
0.62	0.71	-12.7
0.59	0.68	-13.2

The effect of non-random selection in less extreme settings can be illustrated by some work that I have recently undertaken examining the entire case bundles of sets of chambers. Chambers often have more than 1000 cases but of course they operate in different parts of the country and have different barristers who have different specialisations, so they do not represent random selections of cases.

Here, in anonymised form, are the results of the analysis for 10 sets.

Chambers	Approximate Number of Cases	Impact of Proposed Fees in percent
A	800	+0.7
B	400	+2.1
C	700	+6.5
D	900	+3.75
E	600	+0.62
F	1,900	-6.5
G	400	+3.75
H	1,500	-3.8
I	1,500	+0.5
J	1,100	+7.3

As you can see, quite a spread of percentage effects both positive and negative, over quite large numbers of cases.

Martin Chalkley

25 February 2016