

What SSAP24 can tell us about accounting quality

Abstract of the Southampton Discussion

[Institute of Actuaries, 24 May 2010]

Contact

Paul Sweeting, F.I.A., E-mail: paul.j.sweeting@jpmorgan.com

This abstract relates to the following paper:

What SSAP24 can tell us about accounting quality. *British Actuarial Journal*,
doi:10.1017/S1357321711000183

Professor P. J. Sweeting, F.I.A. (introducing the paper): The purpose of my paper is to establish what SSAP24 tells us about accounting quality. I will give an overview of the research, the results and its conclusions, and the wider implications.

I looked at the extent and the value of the SSAP24 disclosures by considering what information was disclosed and the value of that disclosure.

I looked at non-financial FTSE 100 companies because, when looking at leverage and other balance sheet items, financial firms like banks and insurance companies tend to skew the sample. I covered the period from 1989 to 2005 because, although FRS17 was already in place by 2005, there were still firms making SSAP24 disclosures.

SSAP24 offered quite a lot of discretion over the values that could be assigned to the various assumptions. There was also quite a lot of variability in the extent of disclosures. Not all companies showed everything they necessarily could have, therefore giving an indication of management intervention in the content of the accounts. However, it provides information in a quantitative way about what actions management were taking.

Analysing the firms that tend to disclose more or less, or use stronger or weaker bases, in terms of “Are they more highly levered, are they larger firms or smaller firms, or are the pension schemes particularly large or small?”, provides information about the type of firm, on these measures, that is likely to take such actions. When SSAP24 is no longer available it will be possible to say, for example, that if highly levered firms were less likely to give full SSAP24 disclosures and were more likely to use weaker bases, this provides information about the accounting quality of highly levered firms – that they might tend to understate some of their liabilities.

There are quite a few caveats with this research. The hypotheses that I have used in the paper are not the only hypotheses that could be drawn for particular measures. So, for example, different post-retirement discount rates might just reflect the extent to which different firms allow for discretionary pension increases.

If you are looking at the difference between small and large firms disclosures, it might be too simplistic to attribute such differences as due to size. Small firms are often newer firms so some size effects might be age effects, or something else about the nature of the firm rather than some of the hypotheses that I have used.

Another important caveat is that I was not able, in this dataset, to analyse asset and liability bases together. If you use a low net interest rate for assets and a low net interest rate for liabilities, or a high one for assets and a high one for liabilities, those two situations are going to be broadly the same in terms of funding level. Clearly, one is going to give larger liabilities than the other. Not analysing assets and liabilities together limits the strength of the conclusions from the research but, nevertheless, they are still interesting.

First I looked at whether variables were disclosed or not. I employ a logit regression. I could equally have used probit regression, as it does the same type of the thing. It is just a way of transforming something that can take any of a range of values to something that just falls between 0 and 1.

I then looked at the situation where one or more variables were disclosed, on the assets side in particular, in relation to the net discount rate and the actuarial value of assets. For that I used an ordered logit regression, which allows any number of different choices rather than just “yes” or “no”.

Finally, I looked at the value of disclosure. In studying this I have just used standard least squares regression.

For a logit regression, start with a latent variable, which I call y^* . The question then asked is whether we can explain this latent variable in terms of variables, X_1 , X_2 , up to X_N . These can be, for example, firm size, leverage or dividend yield.

The variable y^* can take a value between minus infinity and plus infinity. However, if looking at an item and want to ask, “Was it disclosed or not?” you need to transfer y^* into something which falls between 0 and 1. So use a logistic function and create \hat{y} , which is $e^{y^*}/(1 + e^{y^*})$. This always falls between 0 and 1.

You are therefore effectively saying that if y^* , your latent variable, is less than or equal to 0, your observed variable, y , is 0, but if it is greater than 0, then y is 1. \hat{y} essentially tells you the probability that y is equal to 1.

You would use a similar method if you were looking at, say, the probability of somebody making a claim – certainly the probability of someone dying. If you are looking at generalised linear models for a post code rating, for example, you will often use logit regressions.

A probit regression is similar. This uses the inverse-normal distribution function to transfer data from the minus infinity to plus infinity range to the 0–1 interval.

Ordered logit regression starts the same way. But instead of having $y = 0$ if the latent variable is negative and $y = 1$ if it is positive, you are saying, “We will chop the latent variable into different bounds and we will say that if it is less than the smallest cut-off point, it is equivalent to $y = 0$.”

If it is between that and the next cut-off point, it is equal to $y = 1$, and so on, until you get above the top cut-off point, and that is $y = N$," for however many choices you want.

It is a little less clear exactly what \hat{y} in this case represents, but you could almost regard it as being the probability that $y = N$. You would expect your latent variable to increase in line with your cut-off points. You might say, "if $y = 0$ then the firm discloses nothing, if $y = 1$ it discloses just the net dividend risk discount rate, if $y = 2$ it discloses just the actuarial value of assets, and $y = 3$ if it discloses both of them." That does not mean that you are regarding one as being twice as important as the other. It just gives you a dummy variable for formulating the cut-off points.

If you are looking at, say, car insurance, and you want to model not just whether someone claims or not, but how many claims someone makes in a year – whether it is, say, one, two, three, or four plus – then you can have four different categories and use the same sort of generalised linear modelling approach.

For the logit analysis, I first looked at the net post-retirement discount rate. This was given either directly as a net discount rate or the nominal rate and the rate assumed for pensions increases, from which the net post-retirement discount rate is determined.

I also looked at the net dividend discount rate, and, to the extent that it was disclosed, the actuarial value of assets and other asset-related information. For these, I looked at whether they disclosed or not. I was able to do this because there were enough firms that were not disclosing everything and enough that were disclosing everything.

For the least squares regression, I looked at the net pre-retirement discount rate that was disclosed, the net post-retirement discount rate that was disclosed and the net dividend discount rate.

The results of the research indicate that firms with under-funded schemes tend to disclose less. Firms with big pension schemes disclose more. So if the market can see a pension scheme is large in relation to the company, it tends to demand more information from the firm that has that pension scheme.

Larger firms also tend to disclose more, which is what is expected from market pressure, but tend to use weaker assumptions, which probably reflects the fact that larger firms are able to exert more influence on their actuaries in this case rather than their accountants.

Profitable firms appear to disclose less. This would suggest that it is to do with tax management – that the more money you are making, the less you want to disclose because you are using a pension scheme to try to manipulate your tax position. Even though your taxes depend on what you actually pay in rather than your pension scheme disclosures, there is evidence that they are both driven by the same factors.

Highly levered firms not only disclose less, but, when they do disclose, they tend to use weaker assumptions. The conclusion I came to was that firms that would perhaps be more sensitive to bad press try to hide any bad news by disclosing less. Materiality leads to realism. So if you have something which is big and important, and the market can see it is big and important, you are more likely to have to tell the truth about it: you cannot play it down.

Market pressure means larger firms disclose more but pressure from the firm reduces realism. This is to do with the fact that larger firms use weaker assumptions. Firms tend to alter their accounts to manage tax bills. Highly levered firms have a greater tendency to manipulate the disclosure.

These are the conclusions I take from the analysis of the pensions data but they do not necessarily apply only to the pensions information. The conclusions give information that could be used when considering the results of companies in a more general sense.

There are other points, which are also relevant. The data I was looking at formed an unbalanced panel. What this means is that, while I was looking at data from 1989 to 2005, some firms would be included for a couple of years, some for ten years, and so on.

If you carry out the regression and assume all are independent variables here, then there is a serious risk of understating the standard errors. This is because these observations are not independent. A particular firm might just idiosyncratically be more likely to disclose something. In my analysis I had to allow for clustering in the standard errors by altering the standard errors to allow for the fact that these observations are not independent. This can be done in most statistical programmes.

In the tables for the logit analysis, there are quite a few significant variables amongst the unadjusted standard errors. Having established the robust standard errors, quite a few of the variables stopped being significant. This is important because it means that there are quite a few firm-specific effects. Say Firm 1 tends to disclose a variable and happens to have low leverage; Firm 2 tends to disclose less and happens to have high leverage. The firm-specific effects look more like general effects.

But why is this important for actuaries? If instead of looking at firms, you look at, say, the extent to which policyholders claim on their car insurance, and there is policyholder data going back for, say, five or six years, there are some policyholders who have been in there for more than one year. When carrying out analysis on whether they have claimed or not, by not allowing for the fact that there is a lack of independence in these observations, more significance than is actually valid may be given to some of the factors included in the analysis. It is quite important where you have a number of years' observations, particularly for policy-related data, to allow for links between these individuals and to make sure that the standard errors are adjusted accordingly. If such an adjustment was not made then, for example, the conclusion could be that there is a significant effect in driving a blue car when there is not. In this case, it is just that there is the same three people, who have had policies for the last 15 years, who happen to drive blue cars and keep crashing them. Allowing for clustering can stop these sort of mistakes.

The second wider implication concerns censored data. You can get statistically useful information from looking at the independent variables of the firms that did not disclose anything at all. You can do what is called a 'censored' regression, which takes into account the fact that firms which did not disclose data can provide information about the value of disclosures – information as useful as that from the firms that did disclose information. This can change not just the standard errors but the estimates of the parameters in the equation. If you are looking at, say, car insurance again, and you are looking at the factors that affect the size of claims, you should include the independent variables for people who have not made any claims. Not doing so makes less use of the available information.

There is some interesting material here if you are analysing firms. Some of the techniques are also useful for actuarial work, particularly on the policy-related side when you have many years of data.

The President (Mr N. B. Masters, F.I.A.): We do not have an opener this evening so instead I will make some opening remarks.

Firstly, what was the implication of the timing of the data? The table of data by year starts off with only a few readings and then it builds up to a standard 60 or 70 offices. I would conjecture that this may create a trend in the findings over time. Disclosure may improve over time. One of the things that I learnt, when I was an auditor, is that the first year of any new standard was undisciplined and companies did the minimum, because it was less hassle. Then, in the next year, the auditor said, "Well, you do realise you did not provide as much disclosure as everybody else did" and companies reassessed their approach. By the third year there was considerable unanimity. This could create a trend over time.

My second point is about the ability to manipulate or to drive the results, or in some way put pressure on the external actuaries or external auditors. There are different cultures within companies. Some companies embrace disclosure, whereas some companies resist disclosure, auditors and indeed actuaries. They feel that audit or actuarial review is something being done to them to try to catch them out.

The paper's conjecture that bigger companies were able to put more pressure on their actuaries did not gel with my personal experience. I speak from the viewpoints of both a consulting actuary and an actuarial audit partner. I found that the bigger the company, usually the bigger the audit firm or the consultancy that the company were using. I would have been very surprised, for example, if any of the sample of FTSE 100 companies were using an audit firm outside the big four. The big four firms have a lot to lose by giving in to pressure and they are pretty well organised. You tend to get the toughest audit partners on the FTSE 100 companies. So I am sceptical about the paper's conjecture in this area. Again, there is only a handful of the biggest actuarial consultancies that would be involved and the same arguments apply.

I wondered whether what was happening was slightly more subtle than direct pressure. Different actuarial firms tend to have different house views. It would be worth studying whether or not there was a migration by companies over time to different consultancies (and it would only be quite subtle with a few competitions and beauty parades). Was the consultancy with the most acceptable house view having the edge in the competition? I do not know whether you have done any analysis of which consulting actuary sat behind the various SSAP24 disclosures. It would be interesting as a generality.

A third point was the last line of the paper which said that there is more difference between individual companies than between the sectors in which they operate. That certainly gels with my experience, which is that there tends to be a management culture, and I suspect that, whether you are big, small, leveraged or not, it is much more to do with whether or not the CFO is relatively relaxed and outgoing, or is a detailed, tense and tough-minded individual who does not like anybody getting inside the numbers.

It will be interesting to see whether there were correlations in the cultural areas with different companies. Again, the techniques you have developed here would be helpful if taken into the wider arena.

Changing tack a little, the paper also suggests something more fundamental: the dangers of principles-based accounting. I take this seriously, because of the position that we are going into with BAS – that all our technical actuarial standards must be principles-based.

I have looked at principles-based accounting versus rules-based accounting for many years and I have come to the conclusion that principles-based standards are usually promulgated because they are easier to write and to get agreed. If you have to create a detailed set of rules, it takes a lot of time and effort and, inevitably, you will miss things. So it is much easier for regulators to produce principle-based standards.

What I would propose is a set of rules, together with a set of principles to catch the people who find ways round the rules. That would have saved us from not just SSAP24 – but what the author is indicating in his research, which is that there is going to be bias, and more so in principles-based accounting, because of the individual behaviour of the companies and their individual cultures.

One other point is that if you do want to improve disclosure within your own company, it is important to get the timing right when you do so.

If you do want to improve the disclosure, the most powerful approach is to sidle up to the finance director and point out that the three closest peer companies are all disclosing some item and you are not. You do not do it at the time of the disclosure; you do it in the post-audit mop up. It therefore gives everybody a chance to factor it through the next audit planning cycle.

I got a lot out of the paper. I do not pretend to be able to do all the statistical analysis but at least I now know a man who can. There was possibly poor disclosure under SSAP24 but maybe it was poor performance and good disclosure.

Finally, I was thinking about how we applied some of this, which is almost like a snapshot of behavioural finance. It occurred to me it meant there were probably different branches of behavioural finance because individuals would have behaved differently from companies. I am sure that analysts would have behaved differently again. Maybe what we have in behavioural finance is not just one population but three or even four and we might decide to analyse it in terms of population statistics instead.

Mr P. M. Greenwood, F.I.A.: During the period just after the introduction of SSAP24, I was head of the pensions research department of one of the large consulting actuarial firms and we were one of the first to run surveys of published company SSAP24 disclosures. The author's conclusions fit with my memory of those times. The actuaries who had influence on SSAP24 wanted "business as usual" and were hoping the outcome would be pension cost figures not too far away from the old basis of contributions paid on the basis of a traditional funding valuation. That ultimately was not achieved.

Initially the quality of pensions disclosures was poor and published surveys did not name companies in the summaries. Over time, with peer pressure, competition between the various surveys to grab press attention generated more mention of company names. Some adverse press comment did lead to improved disclosure and this was confirmed by later surveys.

Mr G. E. Finlay, F.I.A.: The paper discloses a number of interesting correlations between the amount of disclosure, the nature of the assumptions disclosed and the circumstances of each disclosure.

There is a further interesting debate to be had about the merits or otherwise of consistency of assumptions. The natural starting point is probably that the investing community (who are the key users of the accounting disclosures) would prefer to have similar assumptions, because this facilitates comparison. A more considered response might be that there are genuine reasons for companies in different circumstances to adopt different assumptions. Provided the analyst comparing reports and accounts across companies and time periods has sufficient information to understand why those different assumptions are used, and there is some indication of the sensitivities of results, variability in assumptions is appropriate.

The key lesson to take from the paper is that what is really needed to make accounting disclosures meaningful is a requirement not just to disclose the main assumptions, but to give some justification for them, and maybe even some indication of sensitivity of results to the key assumptions.

Professor P. J. Sweeting, F.I.A. (author of the paper): The firm of actuaries producing the SSAP24 disclosures undoubtedly had a material influence on the results. However, this data was not easily available. The SSAP24 disclosures themselves could only be collected by searching through the pensions notes to old accounts, and these notes often failed to include information on which firm of actuaries had produced the information. However, I suspect that, if included, it would have been a significant variable in all of the regressions.

Mr P. M. Greenwood, F.I.A: Internally in our survey basis we did record and analyse against actuarial firm involved. From memory there was some correlation but it was by no means 100%.

Ms R. L. Loseby, F.I.A. (closing the discussion): In closing the discussion, I would like to bring out how the paper has highlighted three areas of interest: (i) the specific conclusions and how they might be applied more generally; (ii) the techniques used; and (iii) what the research might say about the merits of principles-based regulation versus rules-based regimes.

Firstly, the author summarised his research looking at SSAP24 covering 1989–2005 and explained why it was worth looking at. This was down to the discretion that was allowed in disclosure.

The first way the research was interesting was in its conclusions about the way firms with certain characteristics behaved. In the discussion, we heard how direct experience of SSAP24 from someone involved at the time corroborated the broad conclusions in the paper, but also some notes of caution in trying to extrapolate the results too far, especially in understanding the causes of the certain behaviours, when the data did not give enough information. But the research is helpful in prompting that debate. We had in SSAP24 a unique dataset, not without its problems, which has made some interesting conclusions. The paper challenges us to think more generally about the behaviour of firms and the nature of disclosure when there is discretion and to look through the figures to management interaction and decision-making.

A second area that the author drew our attention to was the techniques that he used on the dataset: logit regression and ordered logit regression. Having such a clear description of these techniques will be helpful to actuaries in considering how they might be used in other actuarial problems, especially in general insurance. The author highlighted how the techniques he used were particularly helpful when there is an unbalanced panel of data year-on-year with groups entering and leaving the sample and how ignoring this effect could lead to erroneous conclusions being drawn.

Finally the discussion turned to what the research might say more generally about principles-based and rules-based schemes of regulation and how discretion is used. Firms used SSAP24 in different ways and certainly took advantage of the discretion available to some degree. This is a feature of principles-based regulation, but the effect of disclosure over time meant that publicity led to some level of conformity. For principles-based regulation to be effective it needs to go hand-in-hand with adequate levels of disclosure and transparency.

The President (Mr N. B. Masters, F.I.A.): It remains only for me to close the discussion. The paper and discussion approach important issues, and I am sure the subject will reappear for discussion in the future. My thanks go to Professor Paul Sweeting, to Ruth Loseby, who stepped in as closer at short notice, and to the audience.

Mr. Huw Evans, F.F.A. (written contribution):

I was intrigued by the idea that SSAP24 disclosures might be used to measure the quality and quantity of accounting disclosure for firms. However, the paper did not persuade me that these disclosures could be used in this way. It is not that I doubt the mathematics; my scepticism arises from the fact that every time the paper offers an alternative, simpler, explanation for a result, it is that alternative explanation that rings true for pensions actuaries.

The analysis may have been more persuasive had the author not focussed on net discount rates. Pension increase provisions vary from scheme to scheme and from sectors to sector. For example, a number of the pension schemes in the utilities sector were once part of the public sector and so I would expect lower than average net discount rates in this sector simply because those schemes provide a higher degree of indexation than is common in other sectors.

For most private sector pension schemes it was the changes over time to statutory minimum pension increases that caused fixed increases, for which no assumptions were needed, to be replaced by indexation, for which assumptions needed to be adopted and disclosed. One result of these changes is that a difference between net discount rates for two entities may represent, at least in part, a different distribution of the liabilities by year of accrual, rather than a different strength of assumption.

From my perspective there remains enough subjectivity in the choice of assumptions under current pensions accounting standards for a similar investigation of more recent pensions disclosures to be performed. Of course it would be necessary to take account of the duration of the liabilities and yield curves when assessing discount rates and inflation assumptions but I would be optimistic that increased level of disclosure in the current standards does make such an allowance feasible.