

Suggested Mechanism for GP Commissioner Acute Contracts

Dr Rod Jones (ACMA, CGMA)
Healthcare Analysis & Forecasting
hcaf_rod@yahoo.co.uk mobile 07890 640399

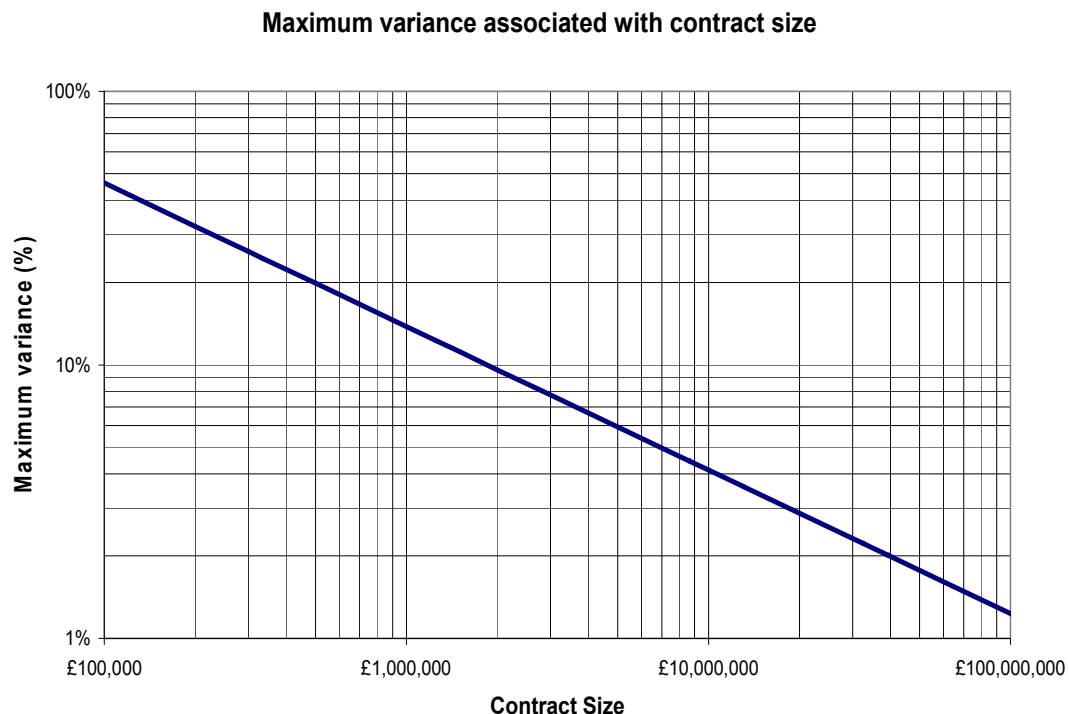
For further articles in the financial risk series please go to <http://www.hcaf.biz>

Introduction

The move toward the use of the national HRG tariff within GP commissioning has created renewed interest in the level of risk associated with healthcare contracts. Most people will be intuitively aware that smaller contracts will be more volatile. What most people are unaware of is the fact that contract financial risk in health care is in general very high.

The material in this document is intended to supplement any Department of Health guidance relating to contracts with NHS Foundation Trusts. It is hoped that this document will give GP commissioners and Trusts a clearer presentation of the real issues and the risks associated with various options.

Figure One: Simulated financial risk for different sized contracts



Financial Risk is High in Most Instances


Figure one gives an example of a computer simulation which attempts to evaluate the minimum possible level of risk by assuming simple Poisson randomness in the demand (activity) for acute services using specialty average inpatient and outpatient prices.

In practice the real variation can be up to three-times that predicted by simple Poisson randomness. This is due to the effect of the wider environment (weather, air quality, infectious outbreaks) which acts to increase the observed randomness. Specialist services (e.g. certain low-volume high-cost events covered by national agreements) have *not* been included in this simulation and hence it represents the risk associated with 'general' acute services only, e.g. emergency, overnight, day case and outpatient services. This simulation also makes the very important assumption that the true average demand is known with high accuracy – this assumption rapidly breaks down as size decreases (discussed later).

Not only does demand vary by more than simple Poisson randomness but the assumption of specialty average prices used in the simulation is very restrictive. Hence the earlier statement that financial risk was likely to be higher than volume-based risk. This arises due to the particular spread of HRG's within the overall volume.

This is illustrated in the Figure Two which is based on actual data for a large acute trust.

Figure Two: Relative volume accounted for by different HRG's.



As can be seen the extremely long tail of very low volume events acts to magnify potential financial instability. This potential financial instability is the basis of contract financial risk.

There are two alternative ways in which the implied risk may be managed, namely, each CCG manages its own risk across a number of contracts or CCG's share risk around single contracts for individual common providers.

Each CCG offsets risk across a range of contracts

The individual contracts for a CCG will cover a number of providers. These contracts will vary in size with individual risk for each contract as per Figure One. In theory the risk of over-performance in one contract should balance under-performance in others (assuming equal cash releasing/consuming flows) such that the overall risk to the CCG will be roughly equal to that of the total value of all contracts.

The main issues here are two fold:

1. Can a lower exposure to risk be obtained by joint risk sharing
2. The precision to which contracts can be formulated and the associated tolerances which trigger additional cash flows

Point No. 1 will be considered in the next section. Our main emphasis will therefore be around the precision to which contracts can be specified.

Can a single year be used to estimate demand?

The main problem here is that the estimate of demand for any contract is a function of the size of the contract, i.e. the smaller the contract the more difficult it is to discern the true average demand and hence to specify the tolerances.

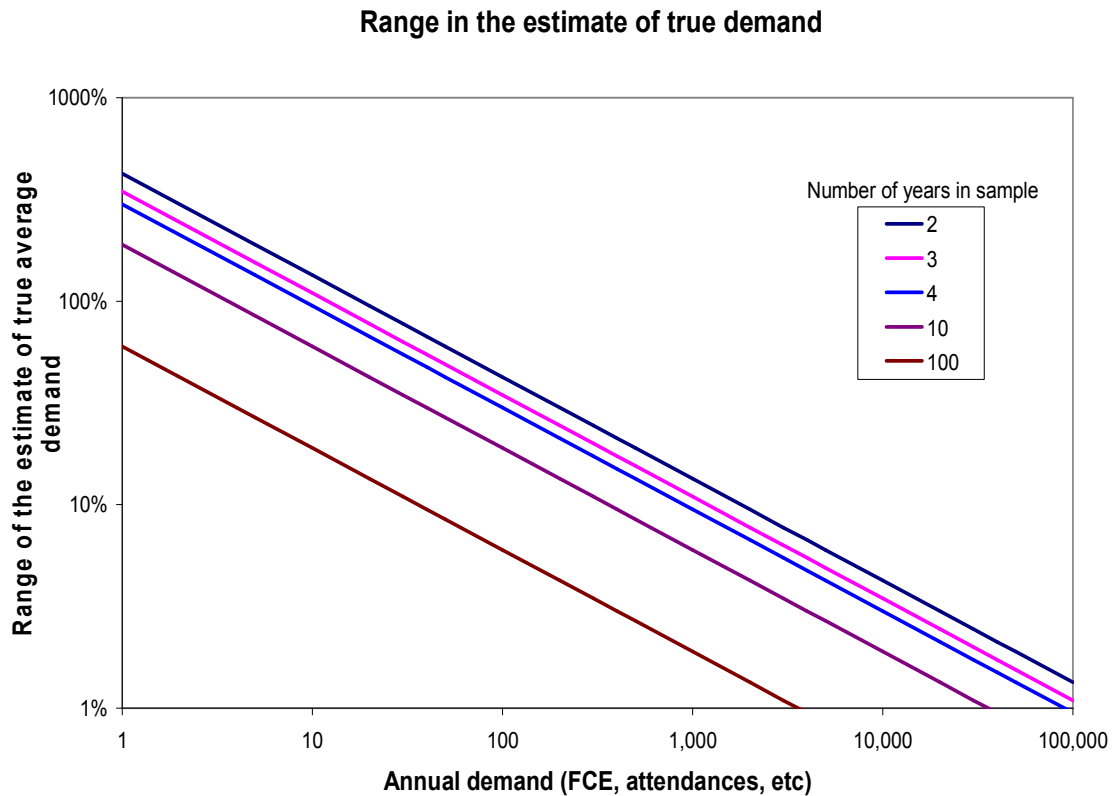
The error associated with the average can be estimated by dividing the observed standard deviation by the square root of the number of data points. From this it follows that any level of 'accuracy' implies a reasonably large number of data points.

For the X NHS Trust most estimates of demand have come from 5 years of data (although 8 years are available usually only 5 are used in the final projection). Hence our observed standard deviation around the expected average can be divided by the square root of 5 or 2.24, i.e. around half of the observed standard deviation is associated with uncertainty in the calculated average. It should also be apparent that a 1 or 2 year sample of activity (as suggested in the national guidance) will be totally inappropriate to determining the expected contract average, and, will more than likely lead to the need for contract re-negotiation as large apparent variances emerge during the progress of the year.

Figure Three illustrates the relationship between number of data points in the sample and overall size. This chart has assumed that healthcare variation is twice the level of that predicted by simple Poisson randomness. The maximum range is therefore 3-times the standard deviation of the estimate of the mean. As can be seen for all but the

largest of volumes it requires 100 years of data to achieve an acceptable estimate of the 'true' average!

Figure Three: Maximum range in the estimate of the true average



The implications of this relationship are fundamental to understanding the art of specifying a sensible contract.

Contracts covering designated HRG's

The first implication from Figure Three is that any attempt to frame a contract at (even at overall) specialty level, due to the relatively small numbers involved, will lead to the generation of numerous variances – due to the inability to specify the expected average with any degree of precision. The use of a single year (as suggested in the national guidance) is guaranteed to generate potential variances of up to 6 standard deviations!

Contracting for specified HRG's is therefore fraught with a multitude of dangers. Table One summarises the position for X NHS Trust. This table assumes that the average derived from 20 months of activity is the true average and that variation is at the level of simple Poisson randomness. Maximum variance is thus three times the square root of the expected average¹.

¹ The real world variation is likely to be up to twice this value.

Table One: Trust total volumes for the 15 HRG's²

HRG	Assumed Average (Spell)	Number on waiting list	Maximum Variance	
			(number)	(%)
B02/03	780	615	84	11%
E04	2	0	4	212%
H01	2	0	4	212%
H02	380	170	58	15%
H03	7	0	8	113%
H04	390	210	59	15%
H10	1315	643	109	8%
J02	210	5	43	21%
J03	130	5	34	26%
J04	110	5	31	29%
J05	190	5	41	22%
Q11	290	272	51	18%

Table one shows that even at Trust total level the maximum expected contract variances due to randomness in demand are at unacceptably high levels both for purchaser and provider alike.

For the Trust to be able to guarantee delivering the contracted target would imply that the number of each procedure on the waiting list at the start of the year was much higher than the maximum variance. This is because the Trust would have to supplement activity by higher than expected consumption of the waiting list should demand be lower than average – which will occur on 50% of occasions. As can be seen from Table One this is the case for the high volume procedures with longer than average waiting times.

Other smaller volume or low average waiting time procedures (for this Trust E04, H01, H03, J02 to J05) do not have sufficient buffer on the waiting list and hence volume will potentially vary (uncontrollably) between the limits set by the maximum variance. Is it fair to offer financial benefits and penalties for volumes of work where the Trust has no control over the outcome?

For the higher volume HRG's the act of splitting the Trust total down to CCG level will only make the situation far worse. Indeed attempts to manage contracts at individual HRG basis will lead to unacceptable waiting list micro-management as Trusts attempt to avoid either perceived excessive or too low activity – due entirely to processes outside of their control. This micro-management is likely to destabilise the rest of the waiting list management process at a time when elective capacity is in short supply.

In this context a contract tolerance of $\pm 5\%$ for each of the 15 HRG's, while totally consistent with simple accountancy, is clearly based on the false assumption that

² The X Trust is a very large hospital group operating over 4 sites with over 300 surgical and over 400 medical overnight beds. E03 and E15 are not performed at the Trust. Version 4 HRG has even fewer spell per HRG and hence the above represents a best case scenario.

random variation in demand does not exist. In fact CCG-based contracts for individual elements of the 15 HRG's are probably unworkable within the context of the overriding objective of overall waiting list management and achieving waiting time guarantees.

As the minimum position CCG's are therefore advised to place larger RVU-based contracts covering all 15 HRG's as a basket rather than individually. This will lead to contracts based on higher volumes and hence with lower inherent variability.

While this makes some small progress to addressing the deficiencies arising from such small number contracts the fact remains that the only rational basis for such a contract is at the level of total surgical activity (see below).

CCG's share risk around the contract with a single provider

A risk sharing arrangement is where CCGs club together to contract at an agreed proportion of the total contract cost for a single provider. Given that the level of risk is very high for all but the largest contracts it would seem prudent for CCG's to at least consider this alternative

The benefits of such an arrangement are that the individual risks are reduced to the lower level of the collective risk associated with a larger contract volume and value.

This risk can be determined by evaluating the standard deviation around the average demand for the total trust inpatient activity. In the example discussed below which is based on X NHS Trust this standard deviation (for the annual total volume) is approximately $\pm 1.4\%$. The maximum possible variance at Trust total level is therefore three times the standard deviation or $\pm 4.2\%$, however roughly 75% of all outcomes will lie in the region ± 1 standard deviation, namely, $\pm 1.4\%$.

This is illustrated in Table Two for the 7 largest CCGs. As can be seen the largest CCG has a maximum risk of $\pm 7.5\%$ (maximum risk is 3 x 1 Standard Deviation). The smallest CCG has a maximum risk of $\pm 67.5\%$. However should these CCGs agree to contract collectively then the maximum possible risk to all parties is $\pm 4.5\%$ ³

Table Two: Variation in demand as one standard deviation

CCG	Activity	1 S.D.
A	316	22.5%
B	1,181	11.6%
C	3,807	6.5%
D	5,113	5.6%
E	14,519	3.3%
F	18,784	2.9%
G	26,269	2.5%
Above 7	69,989	1.5%
Trust Total	77,885	1.4%

³ This is the most conservative estimate of financial risk. It is based on the assumption that all prices are at the overall average and that variation in demand is the only variable. True financial risk is probably twice the level suggested above. Table Two does however present the risk in the correct proportion between CCG's.

By agreeing to share costs in an agreed proportion of the total all participating CCGs have the assurance that total cost will not vary by more than + 4.5% above the expected average.

Should CCG's choose to contract separately then they must bear the potential consequence of at least a 7.5% contract over-performance (for the largest CCG) and much higher than this for the smaller CCG's.

Conclusions

Within the wider commissioning responsibilities of CCGs they may need to think 'outside the box' to achieve their objectives of financial stability.

Should you require further discussion please contact Dr Rod Jones at hacf_rod@yahoo.co.uk

Rod has been involved in commissioning since the early GP fund holder days and has two decades of experience.

Technical Appendix

How to make this work in a fair manner (applying the principles)

The obvious drawback to this arrangement is the process for determining a fair share of the total. The following process is suggested for determining the fair share.

1. Determine case mix adjusted volume using RVU's

Each CCG has a different mix of emergency/elective admissions to various specialties. It is therefore vitally important to adjust the relative volumes for case mix. It is suggested that the national average case mix weightings (RVU's) be used for this process.

This raises the issue of the timeframe over which the case mix should be averaged. In practice one year is far too short a time span and three years would normally be recommended to even out the enormous variation seen in the majority of case mix bands (due to their very small size).

2. What to do about blank or invalid codes

For the purpose of determining a relative fair share between CCGs both blank and invalid codes can be ignored. Blank and invalid codes should therefore be excluded from the analysis and the case mix adjusted volumes determined over the time period. Relative proportions between CCGs will therefore be preserved. This approach is consistent with that applied in the national tariff.

3. Adjust for changes in the waiting List

For the NHS Trust the largest CCG had seen an overall increase in their waiting list over the past 20 months. Most other CCGs have stayed relatively stationary. This probably reflects the fact that the XYZ site is a capacity bottleneck for the Trust as a whole. This will hopefully be remedied in with the opening of additional day case capacity.

However, in terms of the relative fair share this implies that historical activity for this CCG needs to be adjusted upward, i.e. its activity does not reflect the underlying demand (Demand = Activity + change in the waiting list). This adjustment is around 200 in an overall volume of 26,300, i.e. an approximate adjustment upward of 0.8%. This needs to be performed at specialty level since some specialty waiting lists have gone down and others up.

This adjustment can be performed for all CCGs using the case mix adjusted volume. This is then multiplied by (volume + change in the waiting list)/ (volume) to give the final waiting list adjusted relative share.

4. Determine the suggested contract volume and value

This step relies on the calculation of trust total volume which is then proportioned back to the consortium of CCG's choosing to share risk.

5. Agree contract tolerances

The process for calculating the contract volume has already been risk adjusted for the elective proportion of the contract.

The Trust may be willing to accept anything between ± 1 standard deviation as its share of the overall risk. This places 75% of the risk of contract variance with the Trust.

The CCG's would therefore agree to fund at HRG tariff anything greater than 1 standard deviation above the expected contract value.

Within this the elective activity would be capped by the use of control charts, i.e. the only reason for exceeding + 1 standard deviation would be due to emergency admissions.

The Trust would also agree to release money for performance less than $- 1$ standard deviation. Should this be due to unexpected low emergency admissions the Trust should be allowed to meet some of the gap via higher elective activity.