The relationship between decision volume and success rates

There is no fundamental reason why success rates, of any kind[1], should vary with decision volume. Then again, there is no fundamental reason why they should not[2]. What, if any, is the relationship between success rates and decision volume?

[1] This analysis refers to success rates specifically, but these are just one instance of a process of 'trials' and 'successes' in which only binary outcomes are possible.

[2] An exception would occur if peer review allocates funding entirely randomly with a fixed success probability. In that case we would expect to see no relationship between success rates and decision volume, or indeed between success rates and any other factor.

How are success rates and decision volumes related?

Success rates are related to decision volumes in a way that can be described by the following equation[3]:

$$success\ rate = 1 - \frac{1}{10^{an^{(b-1)}}}$$

Where *n* is the number of proposals on which a decision was made, and *a* and *b* are parameters which can be derived either from the data themselves or by fitting this equation to the data. When $b = 1$ there is no relationship between volume and success rate. All other values imply success rates that are positively ($b > 1$) or negatively ($b < 1$) associated with larger application volumes.

Figure 1 shows a funnel plot of ESRC success rates by Research Organisation for decisions made in the period 2014-15 to 2016-17, in which the expected value for each number of applications, *n*, is derived from the equation above.

[3] Derivation in the annex.

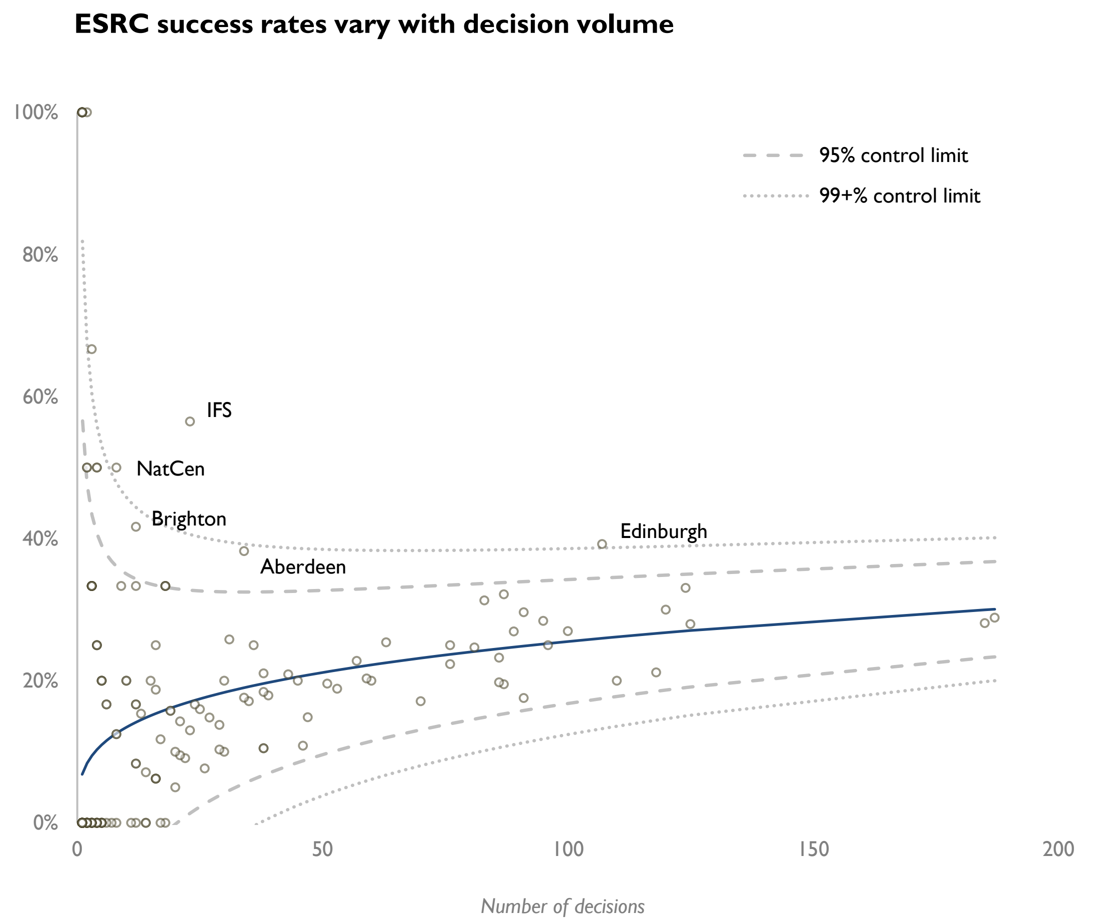

Figure 1: success rates vs. decision volume for ESRC research and Fellowship grants with decisions made in financial years 2014-15 to 2016-17. Solid blue line shows predicted success rate by number of decisions.

The values of *a* and *b* derived from the data are 0.031 and 1.31 respectively, the latter (being greater than 1) indicating a positive relationship between application volume and success rate[4].

The control limits are drawn at approximately the 95% (inner) and 99+% (outer) levels. Almost all the ROs in the data set sit comfortably within the limits and so can be thought of as having success rates which do not differ meaningfully from those we would expect to see for ROs applying that number of times.

[4] The 95% confidence interval for b is 1.24 to 1.39, suggesting that a relationship at least as strong as this is unlikely to have arisen by chance alone.

A few ROs sit outside the control limits. Most notably both Edinburgh and the IFS appear to have success rates much higher than we would expect for ROs applying as many times as they do.

Interestingly no ROs have notably low success rates. If, rather than using a success rate benchmark that varies with the number of decisions, the funnel plot had been centred on a flat overall rate of about 22% we would have identified (probably incorrectly) one additional RO as having a higher-than-expected success rate, with suggestions that the two most frequent applicant ROs might also have unusually high success rates.

Why are success rates and decision volumes related?

It appears that success is associated with volume, at least for ESRC[5]. Do higher volumes lead to increased success, does increased success lead to higher volumes, or does something else behind the scenes influence both success and volume?

It seems most plausible that it is the latter two factors that lead to the observed relationship. Simply applying more frequently is not going to increase any RO's success rate. An unconsidered increase in proposal volume is more likely to have the opposite effect to the one desired. But there is a more pragmatic reason why applying more frequently is not, on its own, a good strategy.

The curve which describes the relationship between volume and success is just that: a curve. So the rate at which the success rate increases varies with the number of decisions made[6]. It does this in a way which itself varies with the number of decisions.

In other words, while there might be an association between success rates and volumes, the strength of the association decreases at higher volumes, and it does so quite rapidly. Figure 2 shows this graphically:

[5] Examples for other organisations and other measures which follow the same pattern are in the annex.

[6] The rate being given by the derivative of the equation above with respect to *n*, namely the following, horrendously ugly, expression: $\frac{\ln(10)a(b-1)x^{(b-2)}}{10^{ax^{(b-1)}}}$. This would be only slightly simpler with base e.

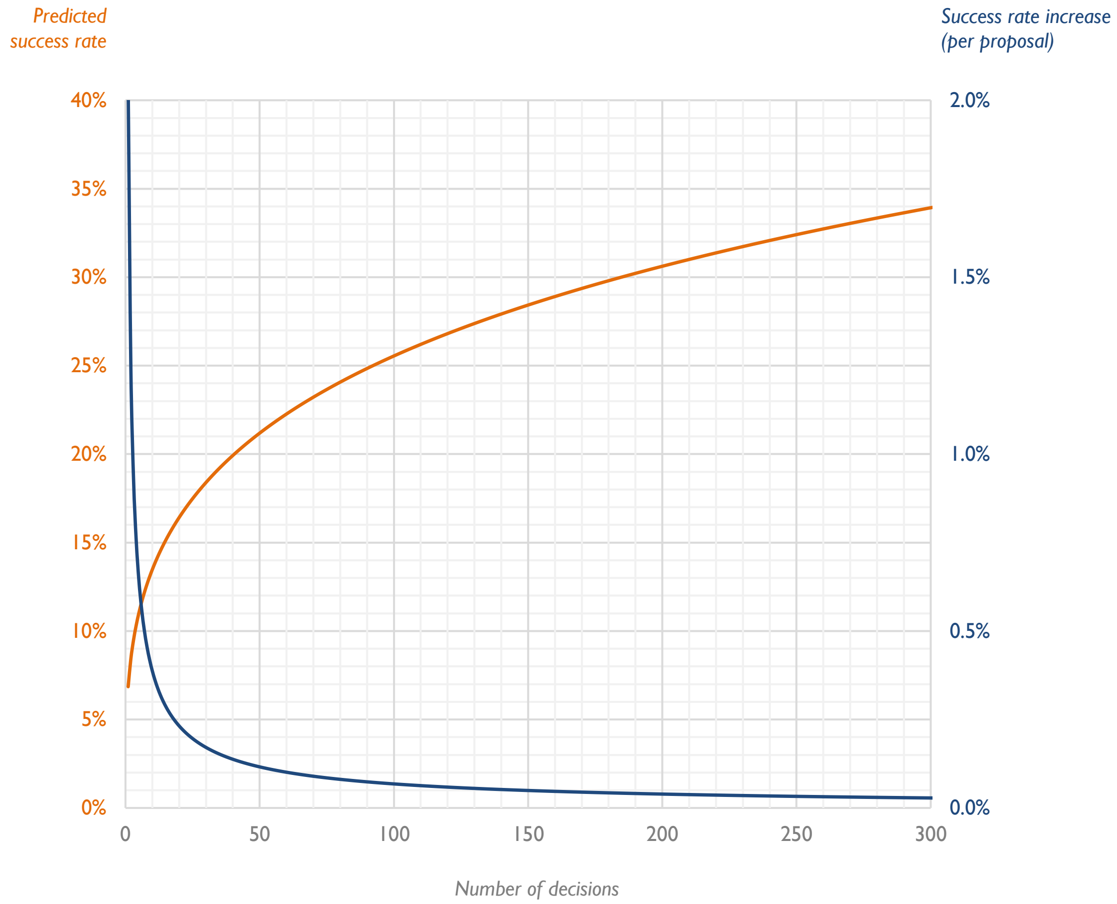

Figure 2: predicted success rate (orange, left axis) and success rate change per additional proposal (blue, right axis) for ESRC decisions shown in Figure 1.

The orange line is reproduced from Figure 1. The blue line shows the rate at which one additional proposal will increase the success rate for an RO having that number of decisions made. Several conclusions stand out:

1. Even for ROs applying very infrequently, the absolute increase in success rate is very small. It is less than 1% for all applicant ROs with three or more decisions. If the right axis was on the same scale as the left axis, the blue line would appear almost completely flat.
2. For decision volumes above 30 the association is about 0.17% per proposal and decreasing. We know that just 25% of decisions relate to ROs applying less than 10 times per year[7], implying that three quarters of applicant ROs will find themselves in a regime in which volume is not strongly associated with success.

[7] See Figure 11.6 on p. 148 of http://www.esrc.ac.uk/files/about-us/performance-information/esrc-analysis-2017/

3. ROs applying more than 100 times in the three years (so just over 30 times a year on average) have success rates that differ inherently by no more than 0.07% per proposal: really not worth considering.

While it might be tempting to see the blue line in Figure 2 as, quite literally, the learning curve for ESRC grant applications, it is not: it merely describes the data. The cause of the relationship is a separate matter entirely that cannot be understood without further investigation.

Annex

Deriving the relationship

If the success rate varies by the number of applications, *n*, made by an entity then every group of entities each applying *n* times will have a group success rate. From this we can work out, for each value of *n*, the proportion of entities applying *n* times that experienced at least one success. This will be

$$1-[(1-success\ rate_{[n]})^{n}]$$

Taking (1- success rate) to be the 'failure rate' (FR) and treating the proportion of entities undergoing *n* trials who receive at least one award as a probability *p*:

$$p = 1 - FR^{n}$$

And so

$$1 - p = FR^{n}$$

This means that the odds of experiencing at least one success if you are in the group of entities undergoing *n* trials are:

$$odds = \frac{1-FR^{n}}{FR^{n}}$$

$$odds + 1 = FR^{-n}$$

And so a sensible starting model for the relationship between the success proportion and the number of trials would be linear in *n* and something like:

$$\log(odds+1) = -\log(FR) \times n + constant$$

In fact this relationship isn't quite right and doesn't fit any observed data. based on observation of various data sets a better model is a similar power law relationship of the form:

$$log(odds+1) = an^{b}$$

With *n* being the number of trials. If using base 10, then:

$$odds = 10^{\mathrm{a}n^{b}} - 1$$

And the probability of <u>not</u> getting at least one award (call this p_{fail}) is

$$= 1 - \frac{odds}{1 + odds}$$

$$= \frac{1}{10^{an^b}}$$

And so for each value of *n*:

$$implied\ success\ rate = 1 - \left(p_{fail}{}^{1/n}\right)$$

$$= 1 - 10^{-an^{(b-1)}}$$

Or equivalently:

$$implied\ success\ rate = 1 - \frac{1}{10^{an^{(b-1)}}}$$

The underlying success proportion around which the funnel can be drawn depends on *n* and two parameters, *a* and *b*.

When $b = 1$, the success proportion is flat (and equal to $1 - 10^{-a}$) generating a traditional funnel plot. For all other values of *b* the success proportion will increase ($b > 1$) or decrease ($b < 1$) monotonically with *n* in a way that appears to match real-world success rate data well.

The original derivation of the relationship relies on actual data, and it is possible in some cases to extract quite good estimates of *a* and *b* from the data themselves. But, as the equation describes all likely relationships between success proportions and *n*, it is probably reasonable and certainly more useful to instead produce estimates of *a* and *b* by fitting to the data a curve based on this equation.

Examples from outside ESRC

This relationship fits a range of data well, for example when used to model MRC success rates in the period 2012-13 to 2014-15:

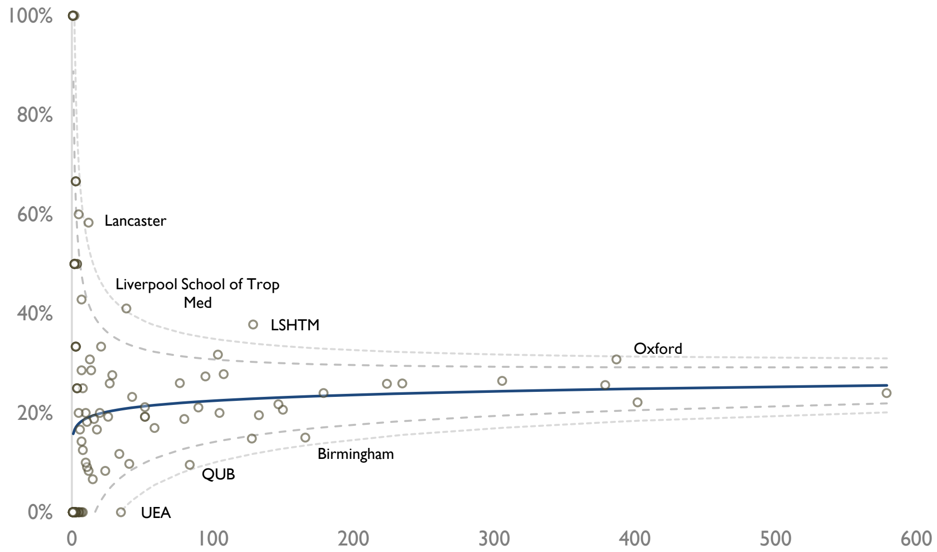

For which *a* = 0.072 and *b* = 1.092 (95% CI 1.041 to 1.147, tending to confirm the presence of an association between volume and success).

Looking beyond the Research Councils, the data for overnight total bed occupancy rates in Q1 of FY 2017-18 for England[8] by NHS Trust can be modelled in the same way:

[8] https://www.england.nhs.uk/statistics/statistical-work-areas/bed-availability-and-occupancy/bed-data-overnight/

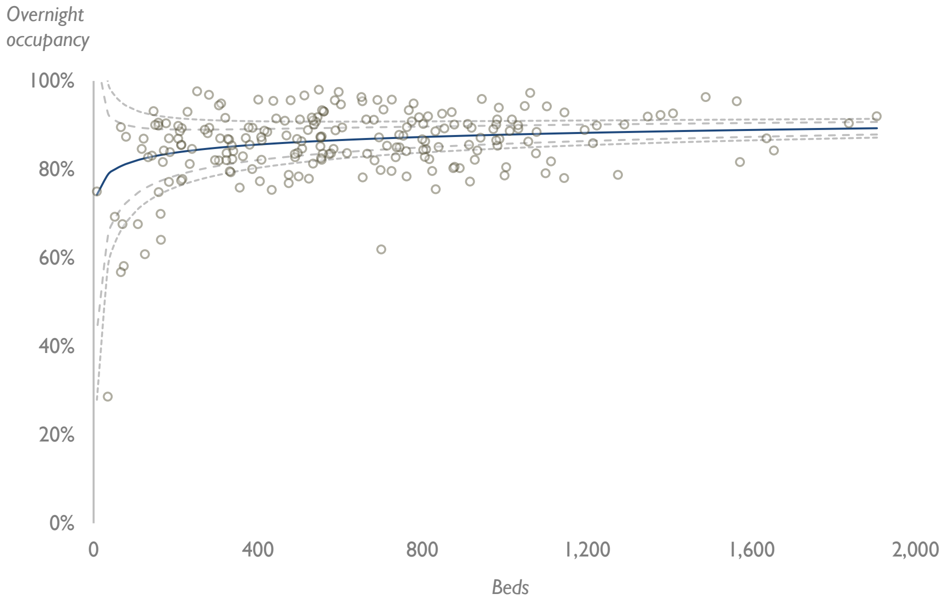

In this case *a* = 0.488 and *b* = 1.091 (95% CI 1.036 to 1.145) and there seems to be a relationship between bed occupancy and bed volume in the English NHS, with Trusts with greater numbers of beds expected to have higher occupancy rates. Leeds Teaching Hospitals NHS Trust, with 1,838 overnight beds in the period, would be identified as having an abnormally high occupancy rate with a traditional funnel plot, but on adjustment seems to be within expectations for a Trust with that number of beds.

The same relationship is found for emergency readmissions of female patients within 28 days by Local Authority of residence in 2011-12[9]:

[9] https://indicators.hscic.gov.uk/webview/ - these are the most recent data, chosen purely as an example.

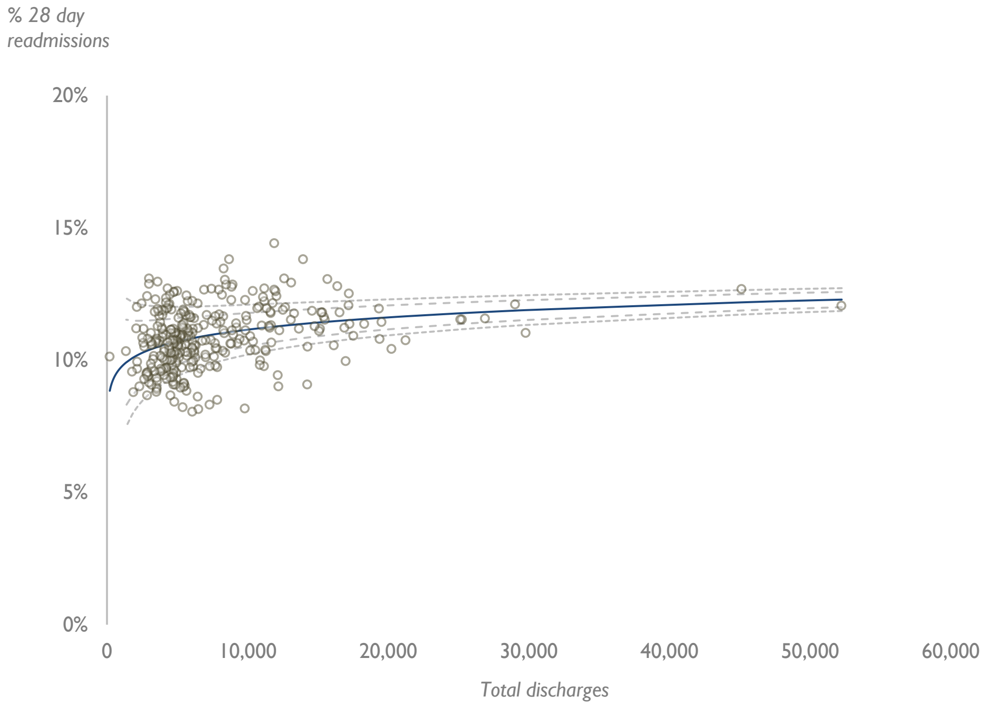

B = 1.062 (95% CI 1.046 to 1.079). In this case the two Local Authorities with the highest total discharges (Birmingham and Leeds) would have been categorised as having abnormally high readmission rates in the absence of the volume correction. Instead it appears that they are as expected for a local authority with that number of discharges.