The effect of ESRC

An analysis of the inputs and outputs of the Economic and Social Research Council's funding processes. ESRC is part of UK Research and Innovation.

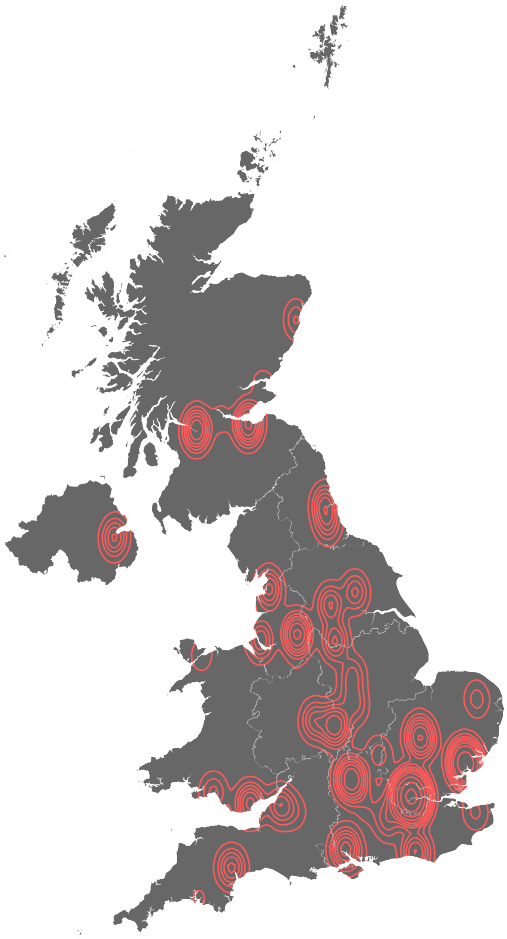

Version 2.1.4, January 2020

Economic and Social Research Council

Contents

Introduction

Much has changed since the first version of this analysis of the effect of ESRC was produced in 2017[1]. Most obviously ESRC is now part of UK Research and Innovation, an organisation that since April 2018 has combined the functions of the seven UK Research Councils, Innovate UK and Research England. Wider events, both those directly related to the work of ESRC (for example the adoption of a target of 2.4% of UK GDP to be spent on R&D) and those of potentially comparable significance but further removed (for example Brexit and the Augar review) have coincided with this transformation.

Have changes in the structure and circumstances surrounding ESRC been associated with a change in the organisation's function or functioning? If they have, this review of data derived from our grants system may show it. Even if there is nothing to see here, the original purpose of this document remains. While some may prefer to treat the operations of a funder as a black box, with all the wiring hidden, there may be others who would like to know more. And we believe that, whatever their preferences, in fact everyone associated with ESRC would benefit from a deeper understanding of the effects of the organisation. By 'exposing the wiring' (hopefully less dangerous than it sounds…) we aim to make our funding processes and cultures more explicable and, crucially, predictable for all involved.

Unpredictability is the scourge of the research funding world. Success rates, the chief and highest-profile indicator of the workings in the box, are predominantly a measure of the predictability of a decision-making process. If they are low, it is because applicants cannot reliably predict whether their idea will be funded – imbalances between demand and supply are a symptom, not a cause, of low success rates. By increasing predictability we can increase success rates, as applicants either alter their proposals in light of an improved understanding of the system, or as they decide not to submit them at all. Better information improves predictability, and so is part of the process of demand management.

Applicants reading this document may gain a slight advantage over those who do not, by avoiding failure modes that are unrelated to the quality of their ideas. That's a good reason on its own for creating and sharing it. It should also be useful to those who are in one way or another responsible for broader strategy in research organisations but who do not themselves apply for funding. By addressing issues relating to the distribution and concentration of funding and how these relate to matters of place and subject, it might help that group to situate themselves more reliably.

It should also be useful to those who are responsible for or interested in the research system as a whole. The ratio of assertion to fact in relation to a range of issues affecting the UK research base is rather high. This analysis may in some small way redress that balance – although it may not tell people what they want to hear.

The data on which it is based is quite straightforward, consisting simply of a list of research grant and Fellowship applications each with associated funding values, a submitting organisation, a Principal Investigator identifier, a primary discipline and a funding outcome (with an associated reason.) It does not include subtleties such as Co-Investigators, secondary disciplines, reviewer scores, or any of the other additional data that applicants

[1] The earlier version is available at https://esrc.ukri.org/files/about-us/performance-information/esrc-analysis-2017/. Some analysis and discussion included in it is not included in this version but may continue to be of interest.

might provide. Nor does it include any doctoral training activity. It also, unlike the first version, does not include an analysis based on protected characteristics.

Despite the simplicity of the underlying data the number of questions it addresses is surprisingly large. But much of it can be boiled down to a simple meta-question: who got what, when and for what, and what might have determined the outcome? It does this in seven sections.

Section 1 presents a brief look at the geographic distribution of resources. 'Place' is a UKRI-wide issue, and Section 1 can do little more than scratch the surface. Section 1 is also where the contentious question of concentration of funding first arises. It shows where (geographically speaking) ESRC funding went. It's fair to say that further thinking remains to be done about how the bald facts of the outcome relate to a coherent and actionable interpretation of what the 'place' agenda implies.

Section 2 is longer, and describes the organisations that are submitting proposals, how many proposals they send in and how applications are themselves distributed across organisations. These questions are of course key determinants of the outcomes seen in the preceding section: 'places' can only receive funding if they apply for it, and concentrations of funding will reflect distributions of applications.

Or they may not. Section 3 looks at the link between demand and distribution: the success rate. If there are systematic differences in success rates by organisation, the pattern of distribution of funding will differ from the pattern expected based on the pattern of applications. And as this section shows, there are systematic differences in success rates of various kinds.

Further complications to thinking about resource distribution arise because all proposals are not the same size. Section 4 describes the effect of size on the portfolio and how size differences play out across it. At this point the more structural description, free of any sense of the nature of the activities we support, is complete.

Section 5 adds more depth by introducing information on disciplines. Each proposal submitted to ESRC has an associated primary discipline. The use of a single discipline to describe a project has potential to be misleading as it underplays the complexity of the portfolio and does not allow us to say how much we actually spend on each discipline: a small fraction of a large grant might actually be more than the total of all grants with that same discipline as their primary classification. Still, the picture that emerges – of a highly varied set of projects and disciplines – is quite reasonable.

Section 6 outlines how disciplinary focus and interest varies across organisations, returning again to the idea that funding can only be allocated to a discipline if people apply for funding in that discipline. It also identifies some organisations that behave differently and provides a measure of breadth and depth of disciplinary activity – the organisational *d* index – which boils that complicated issue down to a single figure. In the face of such complexity, some reductionism is to be welcomed.

The last Section, Section 7, addresses the elephant in the room and attempts to tackle a question which provides challenge across UKRI: do funding processes result in an unjustifiable concentration of funding in either the 'golden triangle' or the Russell Group more broadly? While it can't say whether the result is unjustifiable and it cannot say why this is the case, it does show that, yes, these organisations behave differently and so have outcomes that differ noticeably. To do this though it has to try to estimate how many potential applicants for ESRC funding there truly are, layering yet more uncertainty onto everything.

Finally there are a couple of annexes, one summarising some key characteristics of the 50 organisations with the largest awarded portfolios over the financial years 2011-12 to 2018-19, and one explaining funnel plots (which are the best way by far of presenting and interpreting success rates and related data) as used extensively in the analysis. Some of the data which was used to produce this analysis is available from the ESRC website at https://esrc.ukri.org/about-us/performance-information/application-and-award-data/.

Note that the official record of ESRC success rates in 2018-19 can be found here: https://www.ukri.org/funding/funding-data/decisions-on-competitive-funding/. This analysis extends further into the past than the official data and so draws on a different, though comparable, data set. The UKRI 2018-19 data excludes from the ESRC data some grants for which ESRC was in effect acting on behalf of ESRC, and lists them as 'UKRI'. This means that this analysis and the UKRI 2018-19 data describe different things and should not be expected to come to exactly the same conclusions.

While the document should of course be read in full, we start first with a summary of the key findings.

Key findings

- Broadly speaking, ESRC funding is distributed geographically in line with the location of the UK's most research active organisations. About half is allocated in 12 geographical areas, and geographical concentration of funding has not changed over time.
- Proposal volumes vary substantially from year to year, reflecting decisions about what funding opportunities to offer more than changes in demand for funding.
- Most organisations applying for or receiving ESRC funding are UK-based HEIs, and UK-based HEIs secure the vast majority of funding. This is mostly because their proposals comprise the vast majority of those we receive, and because they tend to do better in competition.
- About half of ESRC decisions relate to responsive mode schemes. More active applicant organisations tend to make greater use of responsive schemes.
- Slightly less than half of the proposals received by ESRC are fundable. Of the fundable fraction, more than half are funded.
- Evidence relating to the concentration of proposal-writing activity is mixed. Some indicators suggest no change over time in the degree of concentration in the most active organisations. Others suggest that, over time, fewer organisations are accounting for an increasing share of proposals and grants. The identity of the most active applicant organisations varies from year to year.
- Organisational success rates and decision volumes are positively associated. This reflects the combined effects of positive associations between decision volume and both the fundability of an organisation's proposals and the chance that a fundable proposal from that organisation will be funded.
- On average, proposals and awards have tended to get larger over time, with smaller awards becoming relatively rare. There has been no trend in the relative degree of concentration of funding in larger awards over time.
- Nearly half of all funding awarded in the period 2011-12 to 2018-19 was channelled through less than 5% of the grants issued in that period, while 30% of the grants allocated in the same period distributed just 3% of the funding.
- The disciplinary composition of the ESRC portfolio has changed slightly over the last eight years, with funding awarded becoming slightly more concentrated in fewer disciplines.
- Most organisations having any contact with ESRC submit a small number of proposals in a limited range of disciplines. Only a handful of organisations have research interests which extend across the majority of ESRC's remit and which result in the submission of multiple proposals in each.
- Organisational interest in disciplines varies considerably. No discipline is dominated by just a handful of organisations' interests and behaviours.
- Disciplinary success rates vary with the level of grant activity, with more active disciplines having lower rates, but this reflects differing levels of use of responsive schemes more than anything else.
- There is no definitive evidence that membership of the Russell Group in general or the golden triangle in particular confers a benefit in funding competitions. But it is certain that those organisations apply more frequently (because they host more researchers who are individually more likely to apply) for more funding, and that they are more likely to receive an award as a result of making an application.

Section 1 – the distribution of funding

The primary effect, and purpose, of ESRC is the allocation of funding to organisations. Figure 1.1 shows the locations of all UK-based recipients of ESRC grants in the period 2011-12 to 2018-19[2] and indicates the amount of funding allocated to each in that period:

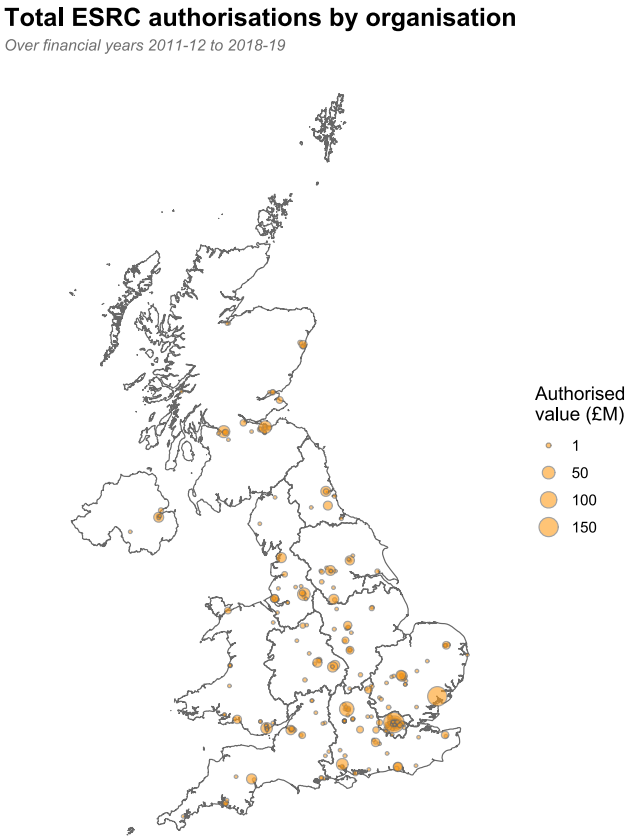

Figure 1.1: UK-based recipients of ESRC grants, 2011-12 to 2018-19. Circles scaled by total grant value allocated. Regions indicated on map are NUTS1 divisions. Contains OS data © Crown copyright and database right 2018.

As most organisations receiving ESRC grants are Higher Education Institutions (HEIs,) as most HEIs will request ESRC funding at some point and as the data covers the last eight financial years, Figure 1.1 is in effect a map of UK HEIs.

[2] 98% of the funding allocated by ESRC in the period 2011-12 to 2018-19 was awarded, initially at least, to organisations based in the UK. Some of the UK total may then have been allocated to overseas investigators or ended up overseas in some other way, but the sum will likely be small. In general, the UK-only ESRC picture and the broader ESRC picture are indistinguishable and the rest of this analysis makes no distinction between the two.

Eagle-eyed readers with a feel for how research funding tends to be allocated in the UK will note one unusual feature: the University of Essex was awarded a level of ESRC funding which is unlikely to reflect the pattern across UKRI.

The total allocation by region[3] is shown in Figure 1.2:

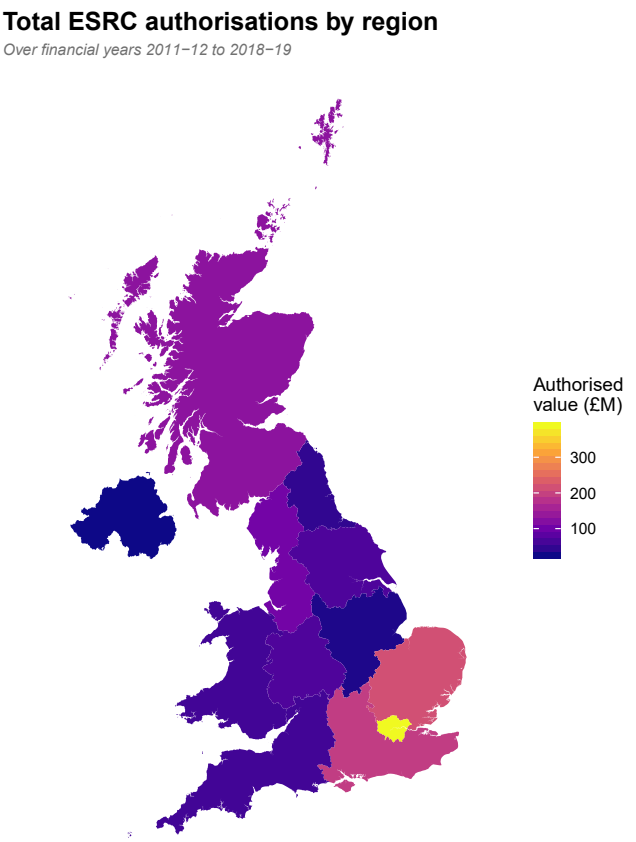

Figure 1.2: total values of ESRC grants by NUTS1 region, 2011-12 to 2018-19. Contains OS data © Crown copyright and database right 2018.

The concentration of funding in London in particular (28% of the total) and the south east of England more broadly (57% across all three regions) is quite striking. This is no surprise as many of the UK's largest and most research-active organisations are found in those regions. The question of whether the allocation is 'fair' is addressed in section 7.

Figure 1.3 shows, on the left, the density of funding awarded across the UK, based on contours outlining the smallest areas which contain within them 5%, 10%...95% of all funding awarded. The broad conclusion is the same as that suggested by Figure 1.2: the majority of the funding awarded was awarded to organisations in London and the south east of England, with smaller proportions being allocated outside those areas.

[3] The regions shown are NUTS1 regions: https://www.ons.gov.uk/methodology/geography/ukgeographies/eurostat

Based on the placement of the 50% contour (right panel) we can say that half of the total awarded by ESRC in the years 2011-12 to 2018-19 could have been reached by visiting organisations in 12 ‘places’: Birmingham, Cambridge, Cardiff, Edinburgh, Essex, Glasgow, London, Oxford, Manchester, Newcastle, Southampton and Sussex. Not all organisations within those ‘places’ will have contributed significantly to the total in the contour[4].

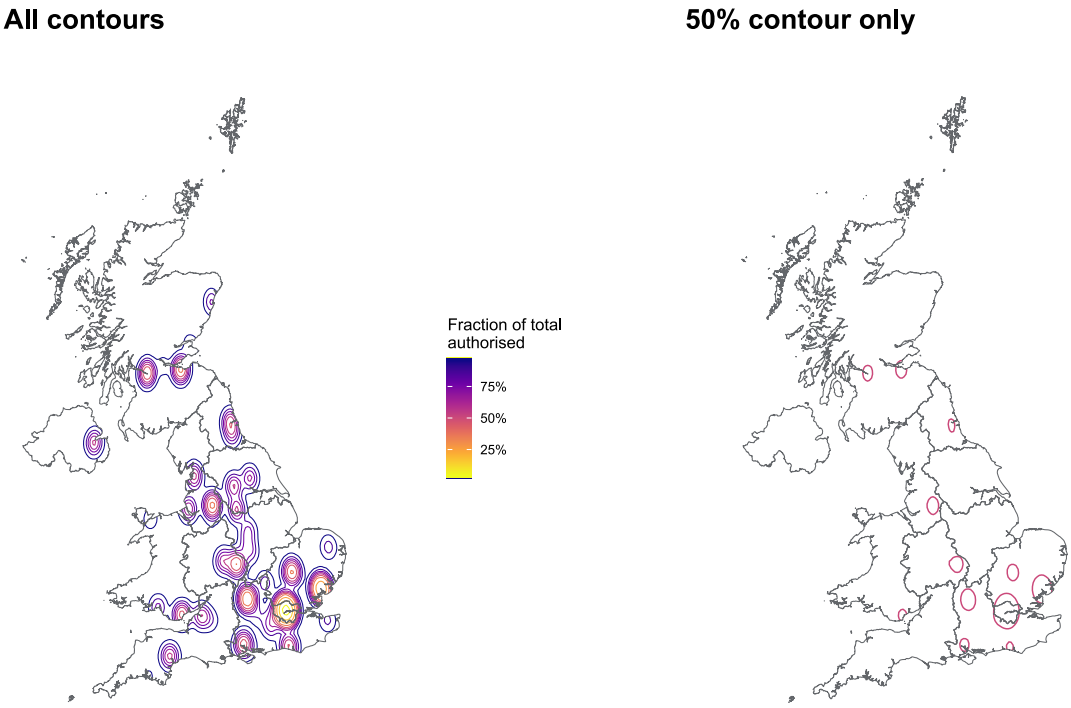

Figure 1.3: density of ESRC funding awarded, 2011-12 to 2018-19 for full range (left) and 50% contour alone (right). Contours are spaced at 5% intervals between 5% and 95% of all funding awarded. Contains OS data © Crown copyright and database right 2018.

The degree of concentration of funding in the UK’s regions can be summarised more quantitatively with a Gini coefficient by region. A Gini coefficient captures in a single figure the extent to which a resource is distributed (un)evenly among entities which might share that resource[5]. They are commonly used in relation to wealth or income distributions but can describe the allocation of any resource among any type of entity. A Gini coefficient of 0 indicates that all entities have an exactly equal share of resources, while a coefficient of 1 indicates that across all potential recipients of a resource only one actually receives anything at all, and so receives everything.

[4] Just one issue which complicates discussions around the apparently straightforward concept of place.
[5] For another explanation see https://www.bbc.co.uk/news/blogs-magazine-monitor-31847943

Figure 1.4 shows by year the Gini coefficient for the concentration of grants authorised across the UK's NUTS1 regions (in which Northern Ireland, Scotland and Wales are classed as regions.) The coefficient is quite stable, at around 0.5 for all years up until 2018-19, suggesting that overall ESRC funding has become neither more concentrated nor more diffuse regionally in the last eight years.

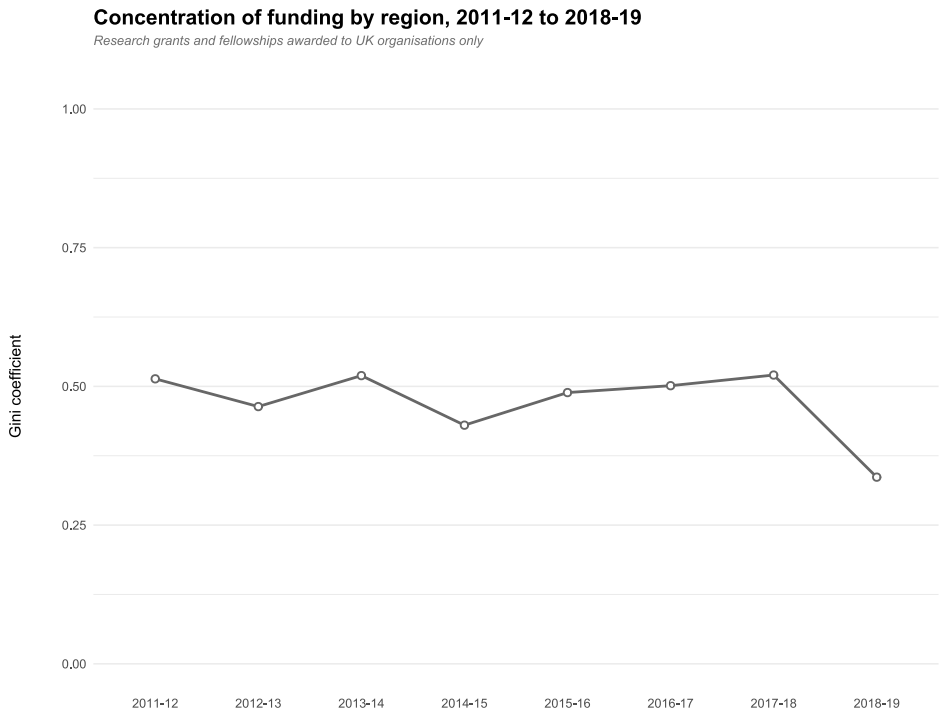

Figure 1.4: Gini coefficients for concentration in UK NUTS1 regions, for grants authorised in financial years 2011-12 to 2018-19. Non-UK grants are excluded.

The sudden drop in the coefficient in 2018-19 is caused primarily by the award of several large grants through the Global Challenges Research Fund (GCRF) to organisations which have not in the past received such large awards from ESRC. As cross-cutting UKRI funds such as GCRF, the Strategic Priorities Fund and ISCF are expected in future to become the norm there is clearly potential for this one-off change to set a new trend[6] and for the distribution of ESRC funding in the coming years to differ substantially from historical precedents.

[6] Sir Mark Walport, UKRI CEO, in evidence to the House of Commons Science and Technology Select Committee: http://data.parliament.uk/writtenevidence/committeeevidence.svc/evidencedocument/science-and-technology-committee/balance-and-effectiveness-of-research-and-innovation-spending/oral/98809.pdf, Q437.

Summary

As ESRC is primarily focused on research of the kind that tends to be done by universities rather than other organisations, it's no surprise to find that the map of ESRC funding tracks the map of the UK's universities very closely. Half the money allocated can be found in just 12 zones in the UK, those zones being centred on universities that the average person on the street might be able to name. It would take a concerted effort and significant change in strategy, process and/or ways of working to achieve any other outcome.

It is not the case that ESRC funding is becoming more concentrated in fewer regions. The pattern we see has been quite static for some time. A somewhat rash bet might be placed on the opposite being the case in the next few years, with 2018-19 being a watershed rather than an aberration. The sudden recent change in concentration is also a striking demonstration of the fact that all metrics associated with ESRC's grant awarding processes are strongly subject to the specifics and details of each year's activities. ESRC's budget is not large enough for the data it generates to feature a core behavioural pattern over which one-off incidents are laid.

Section 2 – decision volumes and organisation types

Figure 2.1 is a further example of how hard it can be to identify meaningful trends in ESRC data. It shows total in-year funding decision (i.e. proposal) volumes for ESRC since the 2011-12 financial year[7]:

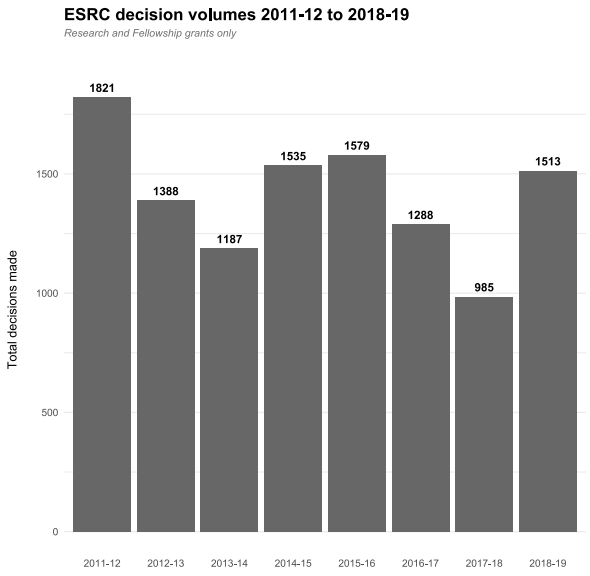

Figure 2.1: total decision volumes by year in financial years 2011-12 to 2018-19.

With the benefit of hindsight, it appears that the dip in volumes experienced in 2016-17 and 2017-18 was unrepresentative of a typical year for ESRC. Year-on-year, decision volumes routinely change by around 30%. The increase of more than 50% between 2017-18 and 2018-19 is relatively unusual and has specific, though individually uninteresting, causes.

All requests for ESRC funding are prompted by calls of one kind or another, meaning that decision volumes will to a great extent be a function of call activity. Any supposedly natural variation in volumes could only be detected within calls that operate reasonably unchanged over extended periods – in other words in responsive mode. Figure 2.2 shows total decision volumes broken down by funding mode (managed or responsive[8]):

[7] Decision volumes and success rates presented here are not the official figures, which can be found on the UKRI website: https://www.ukri.org/. Differences between the figures, if any are present, have no practical consequences.

[8] Any of the following calls: Future Research Leaders 2013, Future Research Leaders 2014, Future Research Leaders 2015, Future Research Leaders 2016, New Investigator, Research Grants (open call), SDAI open call, SDAI Phase 1, SDAI Phase 2 - 2013, SDAI Full 2015. There is no fixed definition of 'responsive' funding in ESRC or indeed across UKRI. The interpretation here is based on the extent to which the call specifies which field of research is eligible for support. None of these has any specific requirements on that front. Other calls not included also may meet that description to some extent – this is a very grey area.

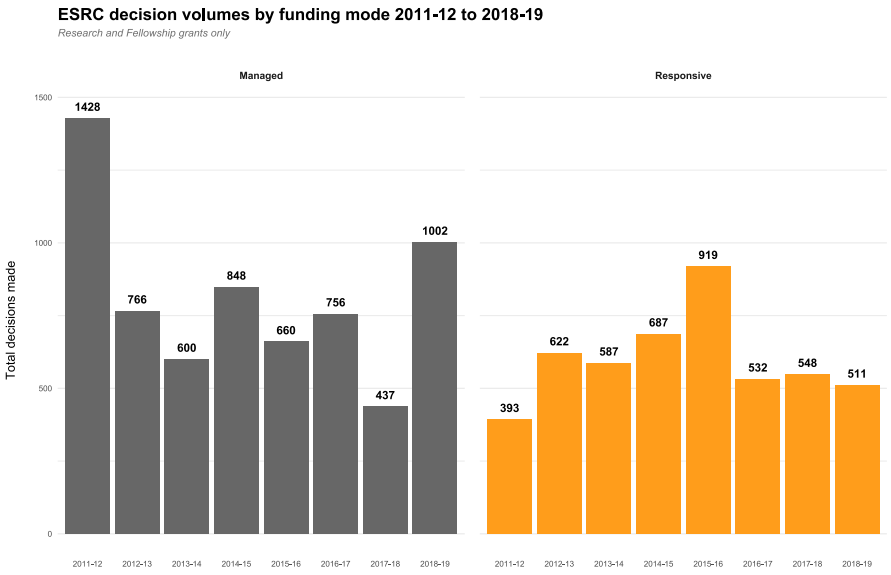

Figure 2.2: decision volumes by funding mode and year in financial years 2011-12 to 2018-19.

The human mind is very good at seeing trends where there are none, and in reality there is little to see here. Managed activities are the product of a combination of funding strategy and tactics, and so are unpredictable, although their resultant decision volumes are entirely explicable. Responsive volumes are strangely static, with the only feature of note being the spike of applications seen in 2015-16 which was caused by a change to funding limits in the 'Research Grants (open call)' scheme.

The proportion of all decisions which relates to responsive mode is irregular (Figure 2.3) but usually around 50% of the total. The single largest responsive call is the 'open call'.

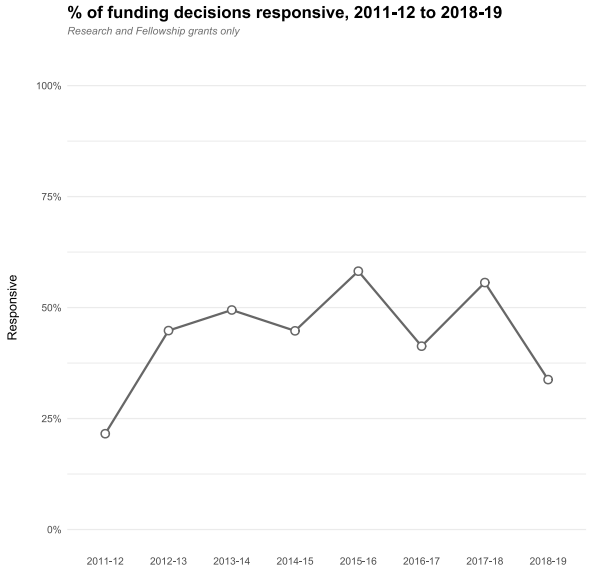

Figure 2.3: % of decisions relating to responsive mode by year in financial years 2011-12 to 2018-19.

Section 1 suggested that the vast majority of organisations seeking ESRC funding are UK HEIs. The detailed figures confirm this view of the composition of both all applicant organisations and the smaller set composed of those actually receiving funding (Figure 2.4):

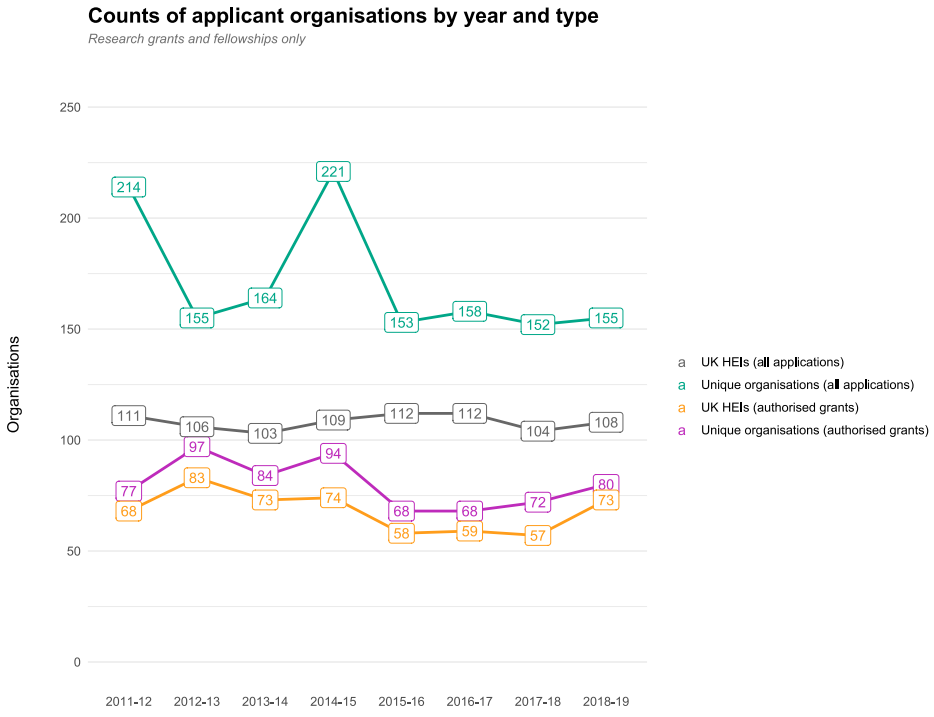

Figure 2.4: number of unique organisations applying for and receiving funding, by year and organisation type 2011-12 to 2018-19.

Generally speaking about 150 unique organisations apply for ESRC funding each year, although the maximum in the period was more than 200. Around two thirds of the UK's 150 or so HEIs apply to ESRC each year. The great majority of applicant organisations that are not UK-based HEIs applying for a grant do not get one.

UK HEIs make up more than half of all applicant organisations and the majority of funded organisations in any given year (Figure 2.5.) 80-90% of ESRC grant recipients in any given year are UK HEIs, this figure being quite stable over time.

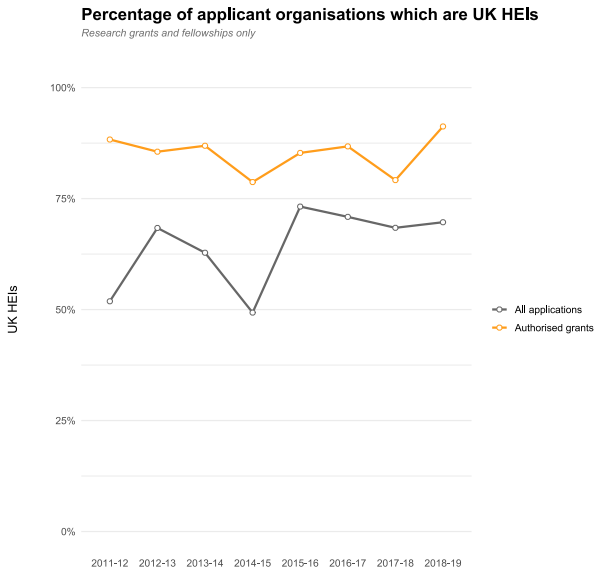

Figure 2.5: % of unique organisations applying for and receiving funding which are UK-based HEIs, by year 2011-12 to 2018-19.

The reason for the disparity in outcomes seen for UK and non-UK based organisations is apparent from Figure 2.6: organisations that are not UK-based HEIs are unlikely to apply very often.

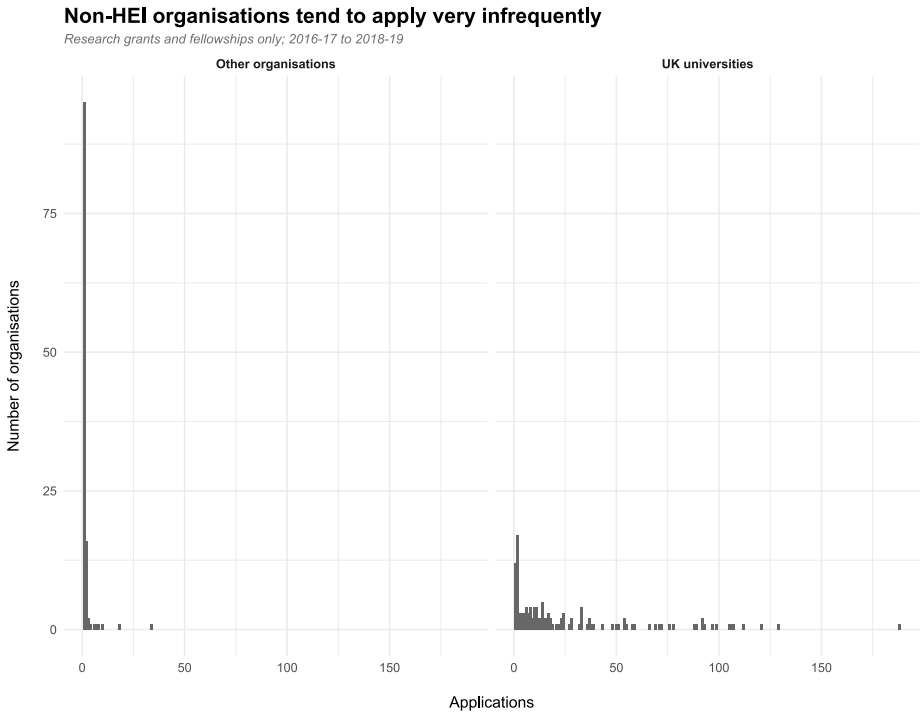

Figure 2.6: counts of unique organisations applying for funding *n* times, *n* = 1, 2, 3… by organisation type 2016-17 to 2018-19.

While several UK HEIs applied more than 100 times in total over the last three years, almost no other organisations applied on more than a handful of occasions. Fewer than 50 UK HEIs applied just once or twice, making that mode of behaviour relatively unusual in that group.

Across all years, the number of organisations applying a specific number of times in a year decreases rapidly with the number of applications (Figure 2.7):

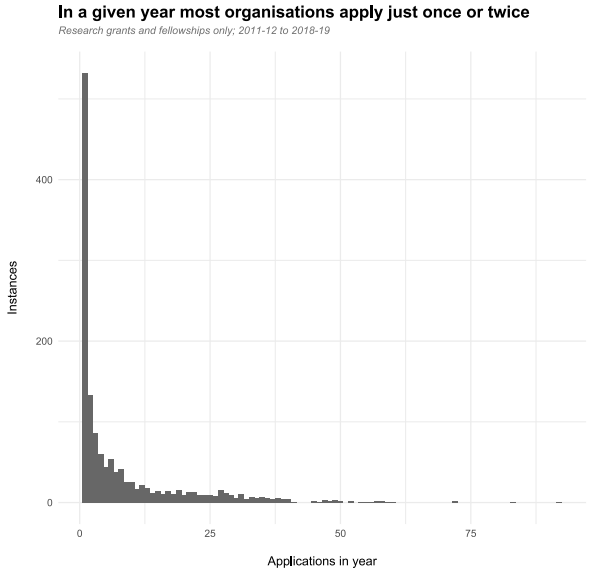

Figure 2.7: organisation count by total number of applications in the period 2011-12 to 2018-19.

One application in a year is the dominant organisational behaviour. Application years (that is the application behaviour of a specific organisation in a specific year) seeing more than ten proposals are relatively rare. But because they are, by definition, associated with multiple proposals a balance is created in which the majority of proposals received comes from organisations applying relatively frequently (Figure 2.8).

The pattern is remarkably consistent within a year. There is a relatively tidy linear relationship between the cumulative fraction of all proposals submitted and the number of proposals being submitted by an organisation, which holds until organisations submit 30 to 40 applications. At this point the cumulative total increases much more rapidly, suggesting a different type of behaviour in organisations submitting more than 30 to 40 proposals in a year.

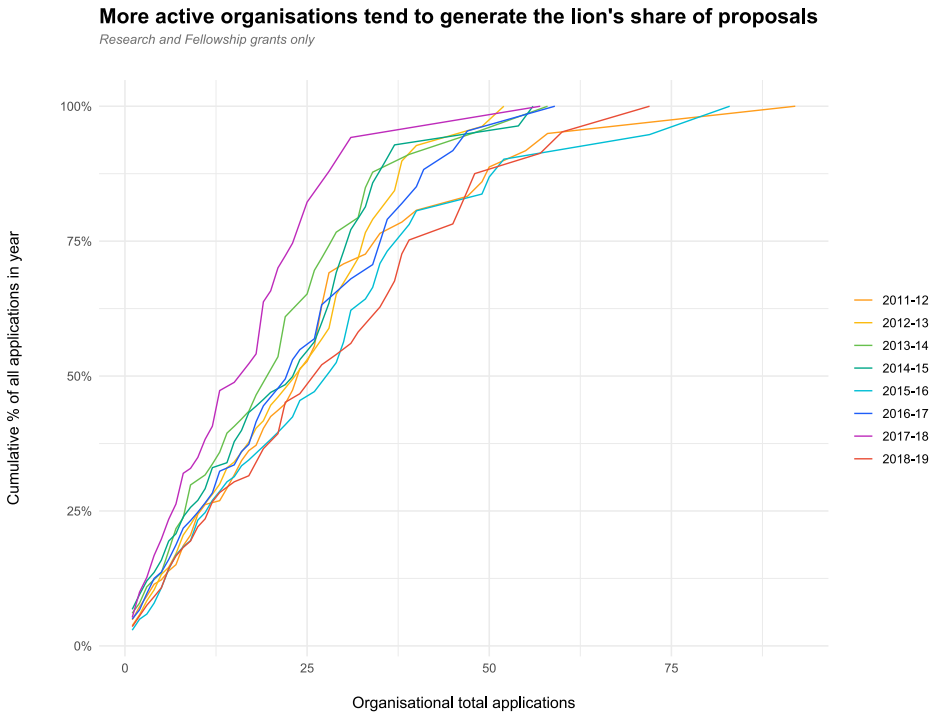

Figure 2.8: cumulative % of proposals being submitted by organisations applying up to *n* times (x axis), by year, 2011-12 to 2018-19.

In Figure 2.8 a shallower line indicates a greater concentration of proposals from more frequent applicant organisations. In general, 25% of the proposals received in a given year come from organisations submitting fewer than 10 proposals, the next 25% come from organisations submitting 10 to 25 proposals, the next 25% from organisations submitting 25 to 35 proposals, and the remaining 25% from organisations submitting more than 35 to 40 proposals. The relationship is variable over time, with the year-to-year variability in the pattern increasing for larger application numbers.

Annual variability is also apparent in the Gini coefficients (and associated Lorenz curves) describing the distribution of applications by value across organisations by year (Figure 2.9):

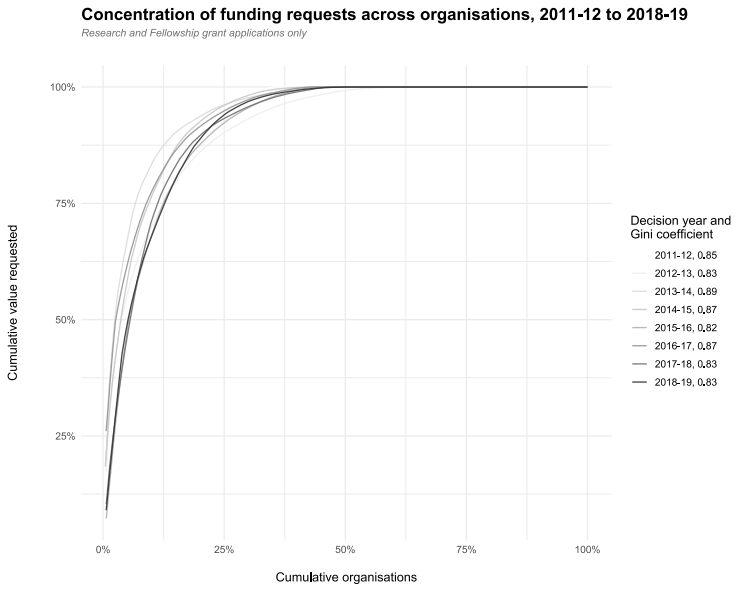

Figure 2.9: Lorenz curves and associated Gini coefficients for distribution of proposals across organisations, 2011-12 to 2018-19. Darker lines indicate more recent years.

Funding requests are highly concentrated in a relatively small proportion of all applicant organisations (with Gini coefficients summarising this concentration being more than 0.80 in any given year.) There is no obvious trend in this figure over time.

On occasion more than 25% of the total sum requested has been requested by just one organisation, but this is the result of a few large grant submissions rather than ramping up of what might be considered ‘normal business’ by any particular organisation. The specific form of each year’s curve varies, and in fact this is one weakness of the Gini coefficient as a summary measure: distributions which have qualitatively different curves may have the same coefficient.

The net result of this interplay of organisation types is that about half of all applicant organisations receive at least one grant when they apply to ESRC in any given year (Figure 2.10). This figure is reasonably stable over time. It also means that there is a 50:50 chance that any given organisation will only have received rejections in a given year. The aggregated outcome of the ESRC decision-making process is reflected in organisational and individual experience in potentially unexpected ways.

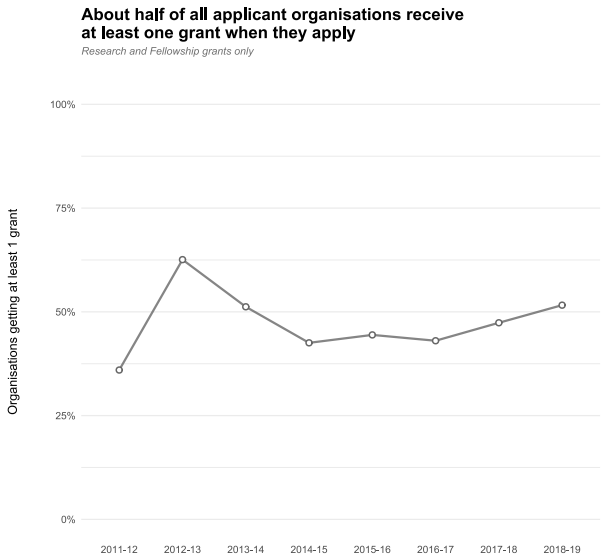

Figure 2.10: % of unique applicant organisations receiving at least one grant, by year, 2011-12 to 2018-19.

Another indicator which might signal changes in the funding distribution over time is the length of the shortest list of organisations that between them submit at least 50% of proposals (or the equivalent list of organisations that between them receive at least 50% of grants.) Both these lists have slowly been getting shorter, by about one organisation every two years, over the last eight years (Figure 2.11):

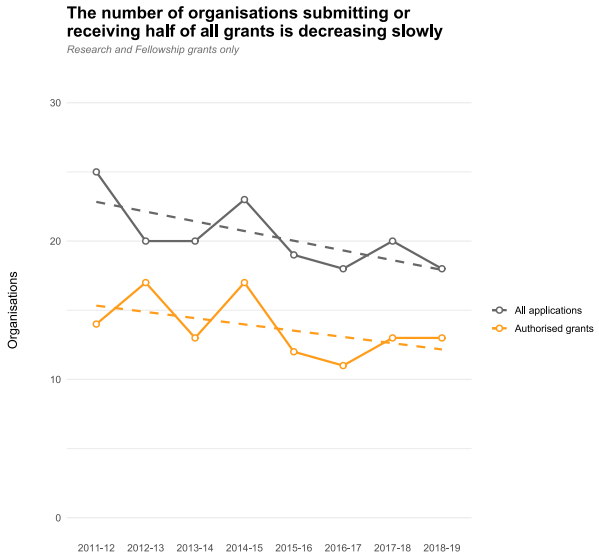

Figure 2.11: minimum number of organisations required to sum to at least 50% of the count of decisions made (grey) or grants awarded (orange) by year, 2011-12 to 2018-19.

The trend is seemingly clear, although it is just possible that what we are seeing is chance variation and that the apparent trend will in future simply drift back to earlier years' values.

All of this implies the existence of what might be described as the 'top' organisations for research requests and/or funding. While the concept is a simple one, and is often used rather freely, the actual membership of that 'top' set is quite variable (Figure 2.12):

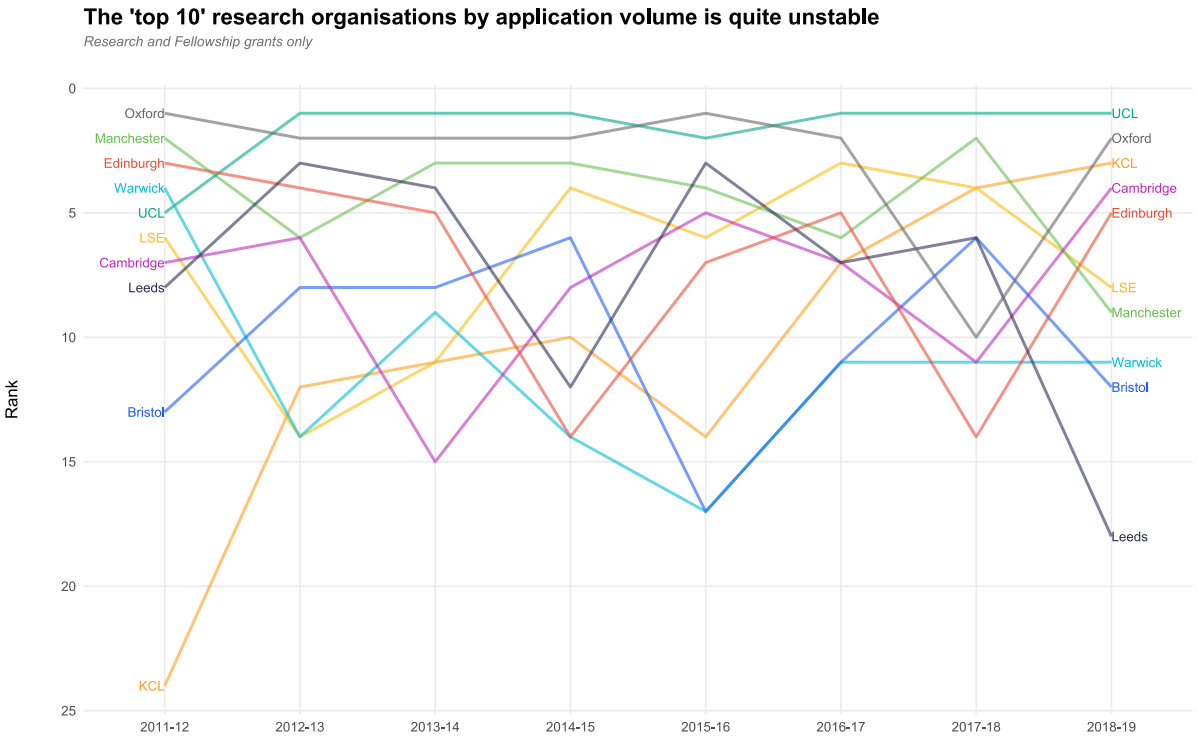

Figure 2.12: annual ranks by total decision volume for the ten organisations with the highest average ranks by that measure over the period 2011-12 to 2018-19.

Ranks based on decision volumes are rather chaotic, only UCL being found in the top 5[9] in each of the last eight years. Of the ten organisations shown only KCL may be on an upward trajectory in terms of proposal volumes, but in general the message is again one of specific circumstances dominating any underlying changes. The same is true for ranks based on funding awarded and so that chart is not shown here.

The general impression of little or no consistent change in the concentration or source of awards is confirmed by Figure 2.13, which shows Gini coefficients for concentration of grants (both applied for and awarded) across organisations over time. Perhaps counterintuitively, awarded grants are less concentrated than are grant applications. This reflects the intermittent nature of grant applications from a long tail of less research active organisations, set against a backdrop of continuous application from a few larger organisations.

[9] The smallest range of ranks for an organisation in a period could in fact be used as an indicator of the extent of variation across the system as a whole.

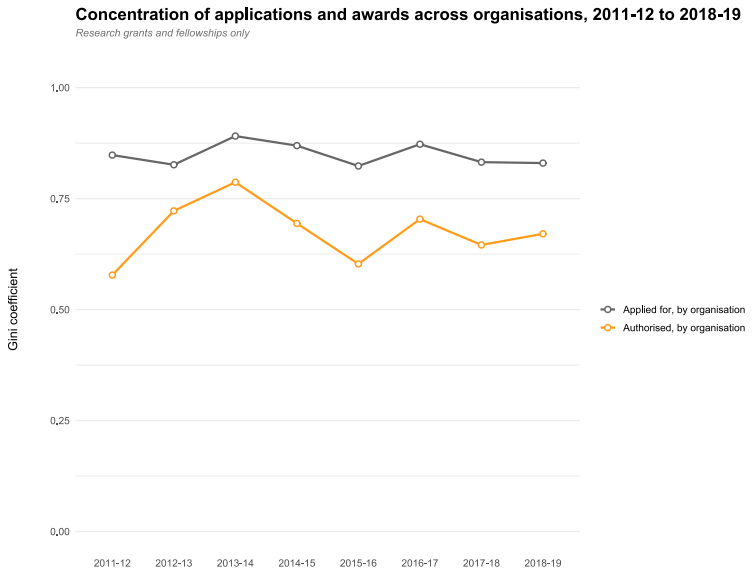

Figure 2.13: Gini coefficients for concentration of applications (grey) and awards (orange) across organisations over the period 2011-12 to 2018-19.

Figure 2.14 shows the association in each year between an organisation's rank by application volume and its rank by award volume:

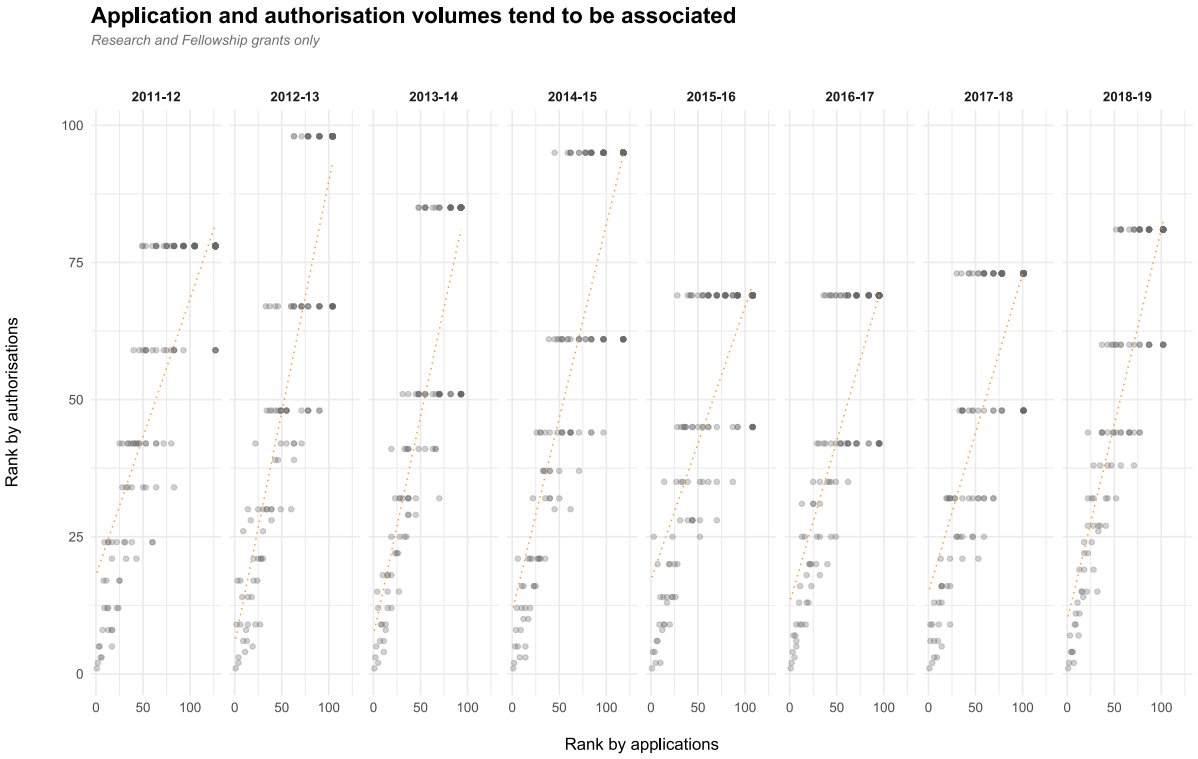

Figure 2.14: organisational rank by grants awarded vs rank by proposals submitted, 2011-12 to 2018-19. Tied ranks shown as lowest value, with indicative regression lines in orange.

While it's not surprising to see that organisations which submit more proposals tend to receive more grants, we consistently find that higher decision volumes are associated with disproportionately high award volumes. The factor that links the two is of course the success rate, which is addressed in the next section.

Summary

Demand for ESRC funding is prompted mostly by calls issued by ESRC. Any description of demand reflects primarily the actions of ESRC rather than any underlying natural process, at least for the 50% or so of decisions that do not relate to open, responsive mode, activity. It is not helpful to think of demand as being like the weather: something that just happens and about which nothing can be done. Extending the analogy, a research council complaining about demand is like a person complaining about the heat having chosen to wear a thick woolly jumper on the beach in August to avoid sunburn: the consequences of the choice are foreseeable, and a different response to the problem might be warranted.

ESRC has the second smallest budget of the Research Councils in UKRI, and yet still attracts proposals from two thirds of the UK's universities each year and awards new grants to half of those applying. The number of organisations that are not UK-based HEIs that we transact with will surprise many although the nature of those contacts – intermittent, specific, usually small-scale – is much less of a surprise.

Both demand for and supply of funding are quite concentrated across organisations. It is hard to pick out any signs of consistent change in the picture over the last eight years. Most indicators seem to vary around a stable average at random, with one year's figure being an unreliable guide to the next. The exception to this is the '50% list', which does appear to be shortening. Even if that is the case, the membership of that list changes rapidly. On balance it is hard to see how a claim that funding is increasingly being concentrated in the hands of the same few organisations could be supported.

Section 3 – rates of all kinds

Of all the indicators of the working of ESRC, or indeed any research funder, the one with the highest prominence is the overall success rate[10]. This is simply, for a specified data set, the ratio of the number (or value) of grants awarded to the number (or value) of grants applied for[11].

This simple statement hides a behind-the-scenes mass of ambiguity and process-related decisions which mean that the actual meaning of such an apparently simple figure is hard to interpret. It is not a pure number reflecting a defined process with a clear start and end point, and it is not a measurement of a natural phenomenon. Changes in success rates reflect changes in process and strategy far more than they reflect changes in the funding landscape, and they are very susceptible to intervention and, in the nicest possible sense, manipulation. Still, they matter and are deserving of some attention.

Figure 3.1 shows both the headline success rate and also the 'median experienced rate' (MER) which is the median of the success rates of the calls to which each proposal was submitted. The MER is probably a better approximation of the rate experienced by the average applicant than is the headline rate[12].

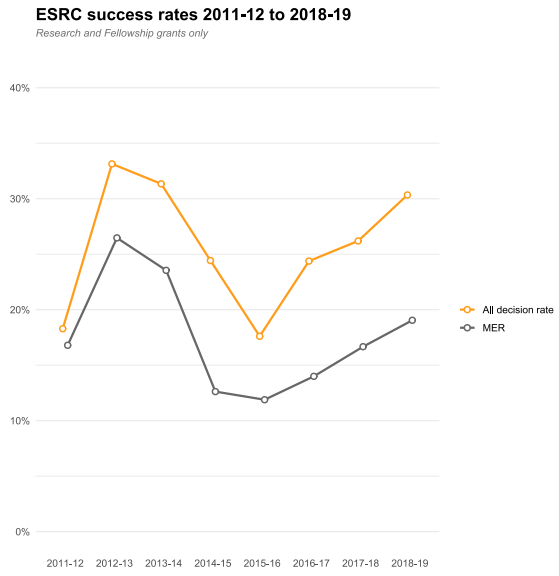

Figure 3.1: overall ESRC success rates 2011-12 to 2018-19.

[10] The rates presented here are not the official figures, which can be found on the UKRI website: https://www.ukri.org/. Differences between the figures, if any exist, are small enough to be inconsequential.

[11] This analysis uses success rates by number only, as the purpose is to try to understand what it feels like to engage with ESRC and what the effects of ESRC's processes are. The experience of an applicant applying for £100,000 is as important as the experience of an applicant applying for £30,000,000 when determining this, and so rates by number are a more useful indicator.

[12] The comparable *mean* experienced rate is exactly equivalent to the headline rate. These two rates both reflect the same concept and can be seen as specific instances of the same idea, with neither being more true than the other. There is nothing special about the headline success rate: it is just one of many related measures that can be calculated and a good case can be made for not even using it.

The rates calculated are a bit of a rollercoaster and probably go counter to what many might believe is happening, in terms both of their direction (increasing in recent years) and size (although the MER is a full 10% lower than the headline rate.)

Both the headline rate and the MER are binary: a proposal is either ‘funded’ or ‘not funded’ and a rate is calculated accordingly. Within the grants that were ‘not funded’ there is a further subdivision, of proposals which were not funded but were fundable and proposals which were not funded because they were not fundable. Figure 3.2 shows the proportion of all proposals falling into each category over the last eight financial years.

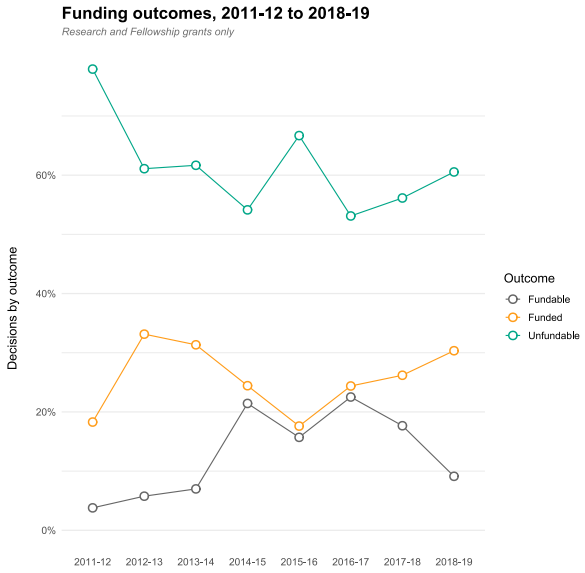

Figure 3.2: ESRC funded, fundable but not funded and unfundable rates 2011-12 to 2018-19.

The ‘fundable’ and ‘unfundable’ rates are slightly less stable than the funded rate and tend to mirror each other. The 2015-16 peak/dip is associated with a burst of responsive mode activity, itself associated with a change in scheme limits.

These more detailed rates differ little across managed and responsive activities (Figure 3.3). For both managed and responsive activities, the most common quality judgement is ‘unfundable’. Success rates vary by funding mode: they are always higher for managed activities than for responsive activities in the same year. This is a reflection of the nature and design of the schemes, not an objective indicator of a problem[13].

[13] At least, not of the problem that many might assume they indicate: a lack of funding for responsive mode. They indicate a lack of predictability surrounding the outcomes of responsive processes, relative to managed ones.

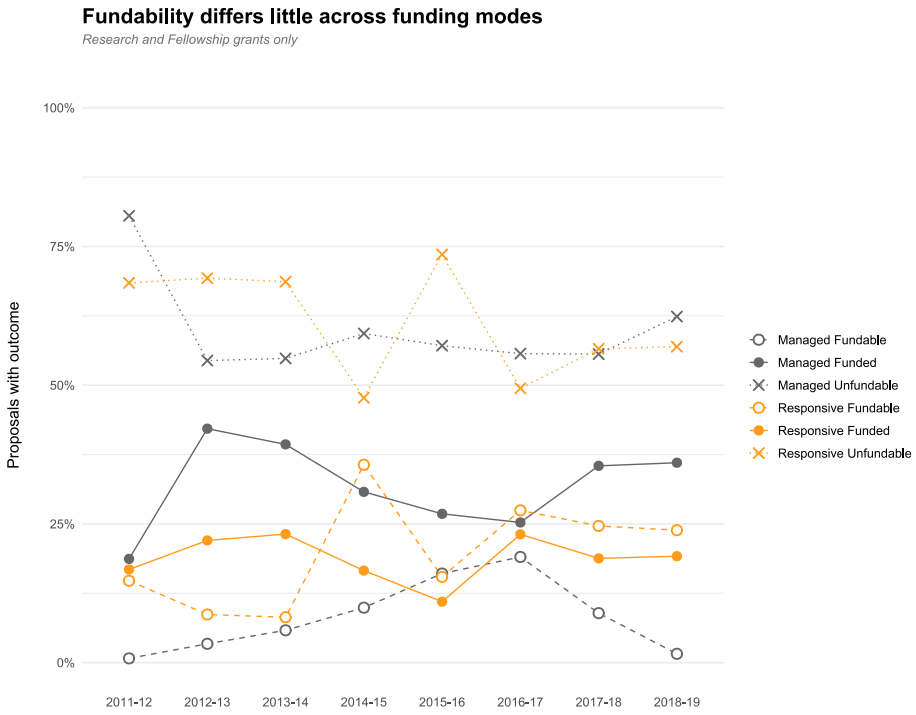

Figure 3.3: ESRC funded, fundable but not funded and unfundable rates by funding mode (managed or responsive) 2011-12 to 2018-19.

In general overall fundability rates (i.e. the total of 'funded' and 'fundable') for responsive activities are no higher or lower than those in managed mode (Figure 3.4). Both are usually lower than 50% in any given year.

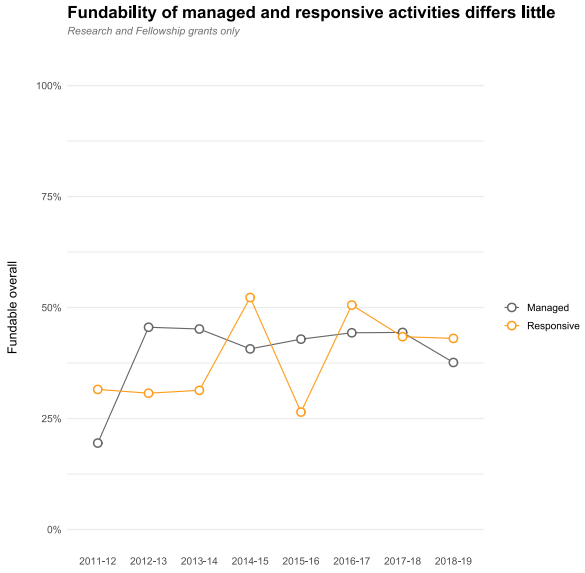

Figure 3.4: % of all proposals received that is judged to be fundable, whether funded or not, 2011-12 to 2018-19.

The lower fundability rates of responsive activities in 2015-16 is the one notable feature in this series[14].

As Figure 3.5 shows, slightly less than half of all proposals received each year are either funded or are fundable but unfunded (Figure 3.5 is the aggregate of Figure 3.4). This proportion has been quite stable since 2012-13:

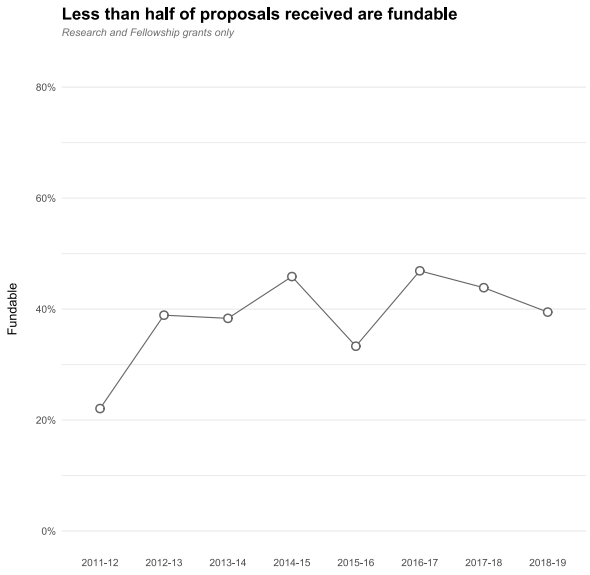

Figure 3.5: % of all proposals received that is judged to be fundable, whether funded or not, 2011-12 to 2018-19.

And of the fundable 40% or so, more than half are funded (Figure 3.6). This rate is the 'conversion rate' and it is highly variable, probably because it reflects a relative and subjective concept ('fundability') which has no clear definition, and which reflects the variety of processes in which it is employed.

[14] We might speculate that the scheme limit changes in the Open call prompted a rush of proposals that would perhaps been better left unsubmitted.

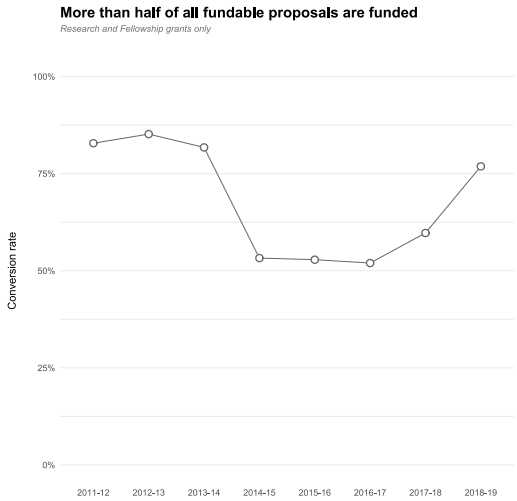

Figure 3.6: % of fundable proposals that is funded, 2011-12 to 2018-19.

Annual success rates summarise the fates of many individual proposals. They reflect the aggregated influence of a number of factors, each with varying potential to lead to a particular outcome for a particular funding application. Of all these factors the two of greatest interest are the identity of the submitting organisation and the proposal's subject area. Success rates by organisation and discipline are thus of specific interest, and in fact contain more useful and actionable information than the headline rates.

Turning first to success rates by organisation, Figure 3.7 is a funnel plot of organisational success rates across the last three years' data. Further explanation of these plots is in an annex, but their key features are:

1. a predicted rate which may or may not vary with decision volume (the orange line in Figure 3.7); and
2. two 'control limits' (grey dashed or dotted lines) which indicate the expected range of variation that can reasonably be interpreted being as compatible with chance variation around the average rate; and
3. data points (representing in Figure 3.7 individual research organisations) which, if they lie outside the limits, may reasonably be considered to be experiencing abnormal outcomes. Points within the control limits have rates which, while not necessarily exactly as expected, are within an expected range and so are not an indication of abnormal outcomes.

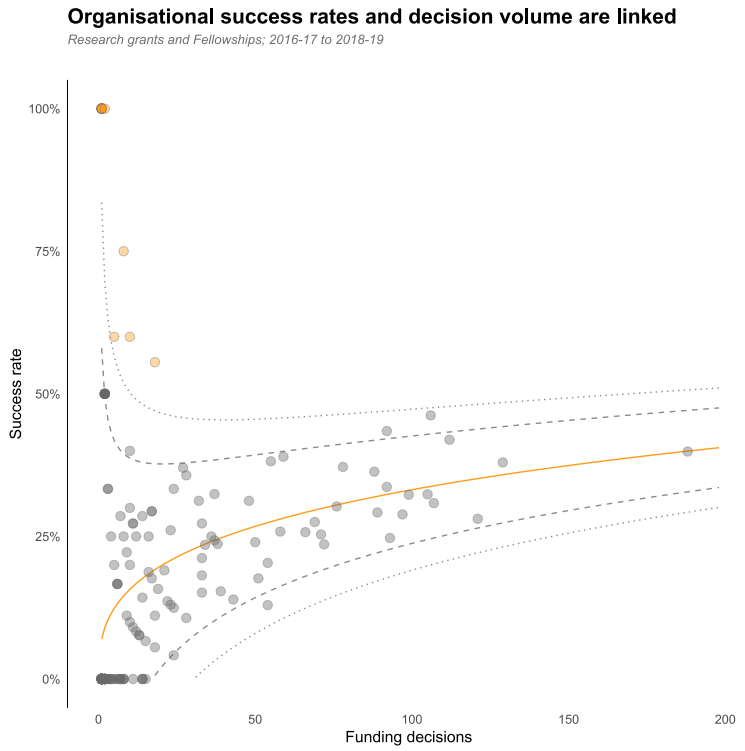

Figure 3.7: funnel plot of organisational success rates over the period 2016-17 to 2018-19. Inner, dashed, funnel shows a 95% control limit, outer (dotted) funnel shows a 99+% control limit.

Over the last three years the association between decision volume and organisational success rate has been quite strong, with higher rates being associated with larger decision volumes. A success-volume relationship is common across UKRI, and certainly not unique to ESRC[15].

Once this relationship is accounted for there are few organisations having success rates when applying for ESRC funding that are much higher than we might expect[16] and none with abnormally low rates (where abnormality is determined on the basis of the more demanding 99+% control limit.) The observed relationship between success and volume is very unlikely to be a chance result, but on their own the figures cannot tell us whether organisations submitting more proposals have higher rates, or whether organisations having higher rates submit more proposals, or whether some other factor linking the two is at work.

Similar plots can be made of both the fundability rate and the conversion rate of each organisation. The overall rate is simply the product of these two rates (Figure 3.8):

[15] It is persistent over time and is present when looking at responsive mode proposals only.
[16] The two organisations with rates above the upper 95% limit but inside the upper 99+% limit on the basis of ~100 proposals are Edinburgh and Sheffield. Those above the upper 99+% limit (and therefore with unquestionably high rates) include the IFS and NatCen.

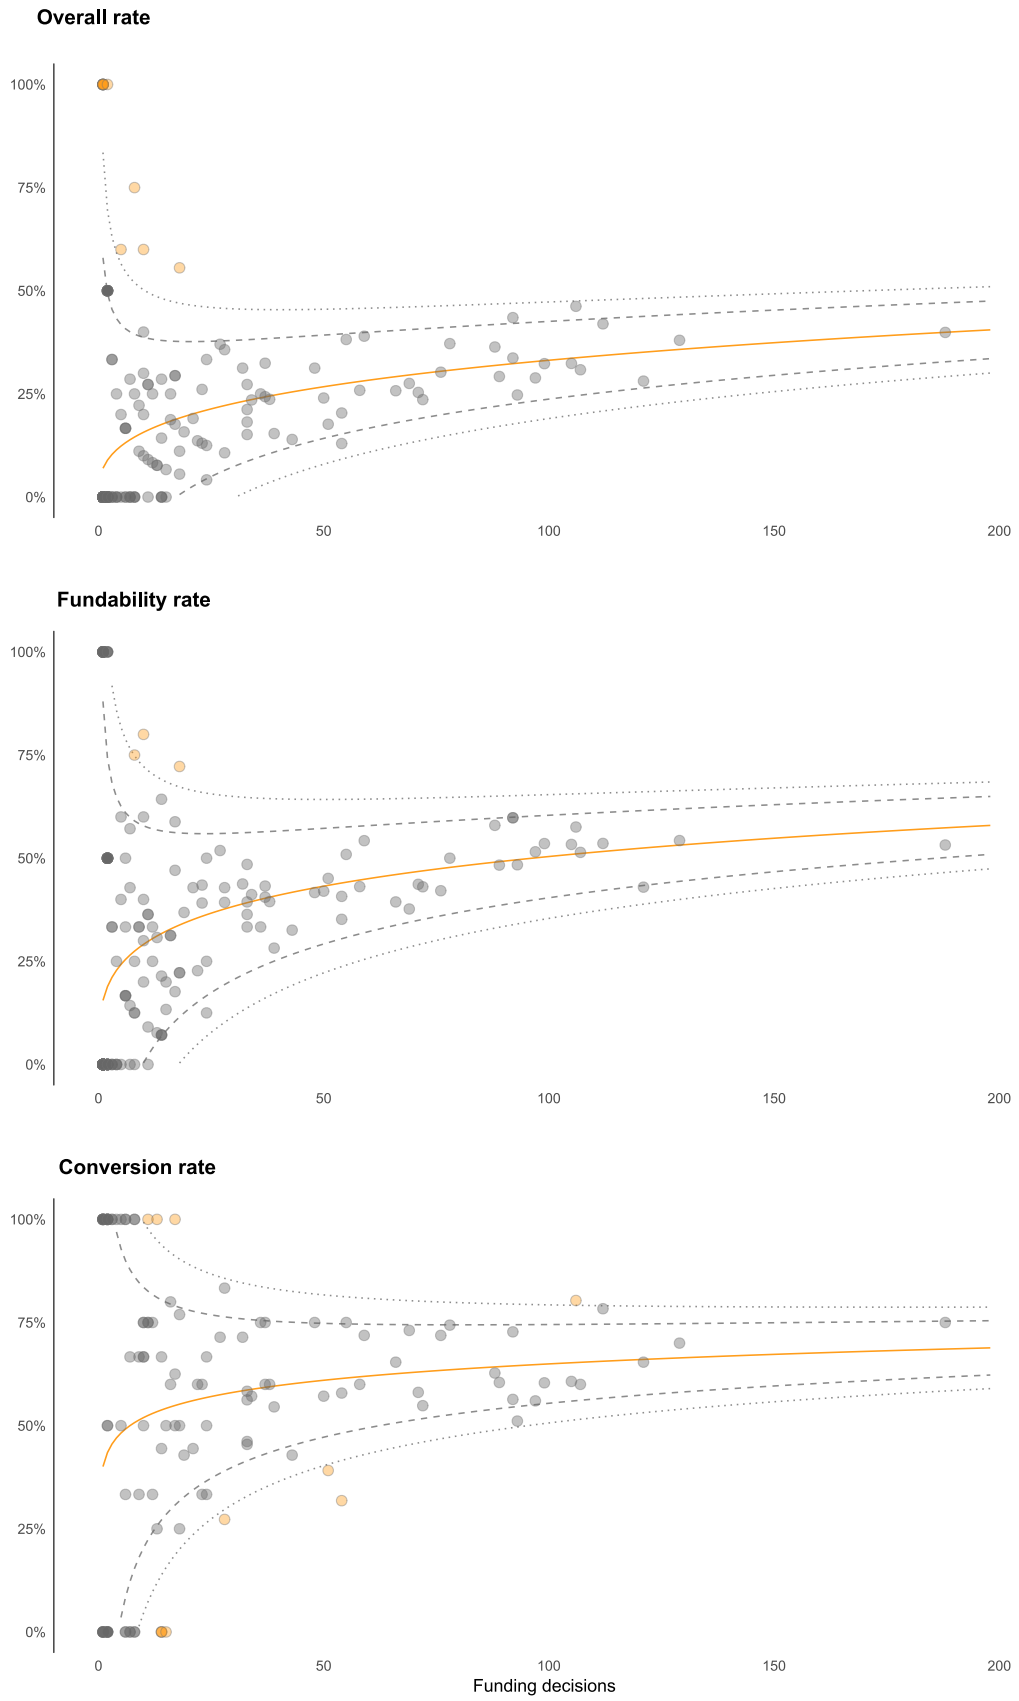

Figure 3.8: funnel plots of organisational success rates (top, duplicating Figure 3.7), fundability rates (middle) and conversion rates (bottom) over the period 2016-17 to 2018-19. Inner, dashed, funnel shows a 95% control limit, outer (dotted) funnel shows a 99+% control limit.

The association between the fundability and conversion rates and decision volume easily meets the traditional $p < .05$ significance level. In neither case are there meaningful numbers of outliers, but one organisation (QMUL) has a strikingly low conversion rate which is the driver of its low, but perhaps not unexpectedly low, overall success rate.

The immediate conclusion is that the overall association between success rates and decision volumes is driven mostly by an underlying variation in the fundability rate. An additional success-volume association in the conversion rate contributes further to the overall effect. Organisations which submit more proposals are more likely to have their proposals judged as fundable than are organisations which submit fewer proposals, and those proposals are more likely to be funded.

Organisations submitting more proposals tend to submit a higher proportion of them to responsive mode (Figure 3.9):

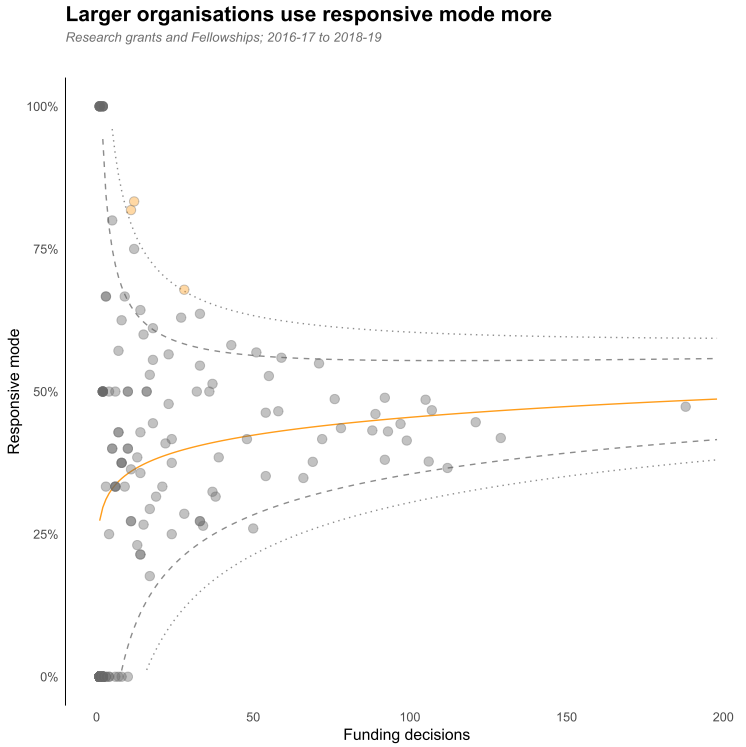

Figure 3.9: funnel plots of use of responsive mode as % of total proposals by organisation, over the period 2016-17 to 2018-19. Inner, dashed, funnel shows a 95% control limit, outer (dotted) funnel shows a 99+% control limit.

As the success rates of responsive schemes tend to be lower than those of managed calls this preference slightly moderates the observed success-volume relationship. If organisations applying less frequently made greater use of responsive mode, their success rates would probably be lower still.

It follows from Figure 3.8 that the ability to produce a fundable proposal and the chance of that proposal being funded are linked (Figure 3.10). Organisations with higher fundability rates tend to have higher conversion rates, but there is a lot of noise and much of the association is driven by outliers.

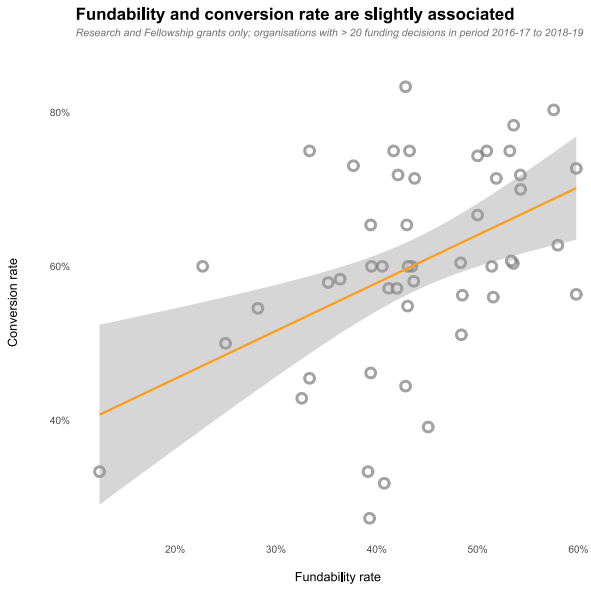

Figure 3.10: scatter plot of organisational fundability and conversion rates over the period 2016-17 to 2018-19. Organisations with < 20 decisions excluded; regression line and associated 95% confidence interval for the coefficient illustrative only.

Figure 3.11 shows the distributions of counts of organisations with particular fundability, conversion and overall success rates. These charts are in effect the density of the points that would be seen when looking back along the x axis of the corresponding funnel plots in Figure 3.8.

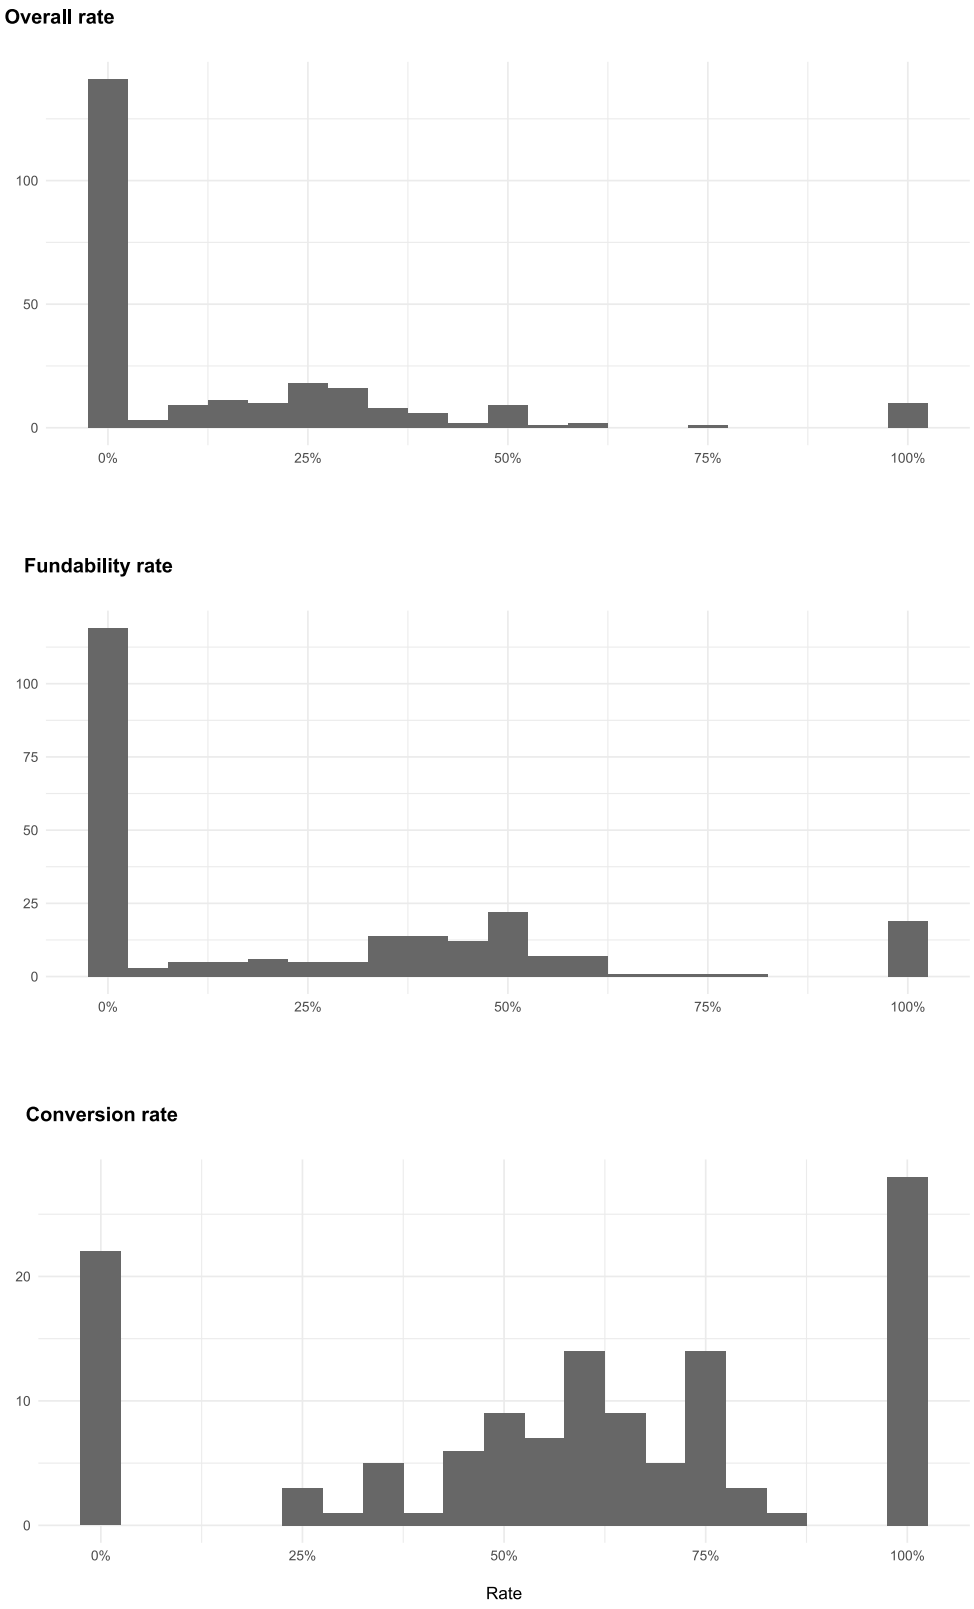

Figure 3.11: histograms of organisational overall, fundability and conversion rates over the period 2016-17 to 2018-19. Bin widths are 5%.

Over the period the most frequently occurring overall and fundability rate was 0%, but for the conversion rate the mode is 100%. Most organisations with either 0% or 100% rates will have just a single decision made in relation to them. For organisations, extreme rates are common; at the individual level they are rare.

As proposals experience one of three outcomes it is possible to describe an organisation's complete set of outcomes in a ternary plot of the kind shown in Figure 3.12.

A typical proposal will come from an organisation which sees between 40% and 60% of its proposals judged unfundable, 10% to 30% of them being fundable but not funded, and around 20% to 40% of them being funded. Many organisations experience outcomes which differ from this average, but these outlying organisations tend not to submit many proposals.

Funded, fundable and unfundable rates

Research grants and Fellowships only; 2016-17 to 2018-19

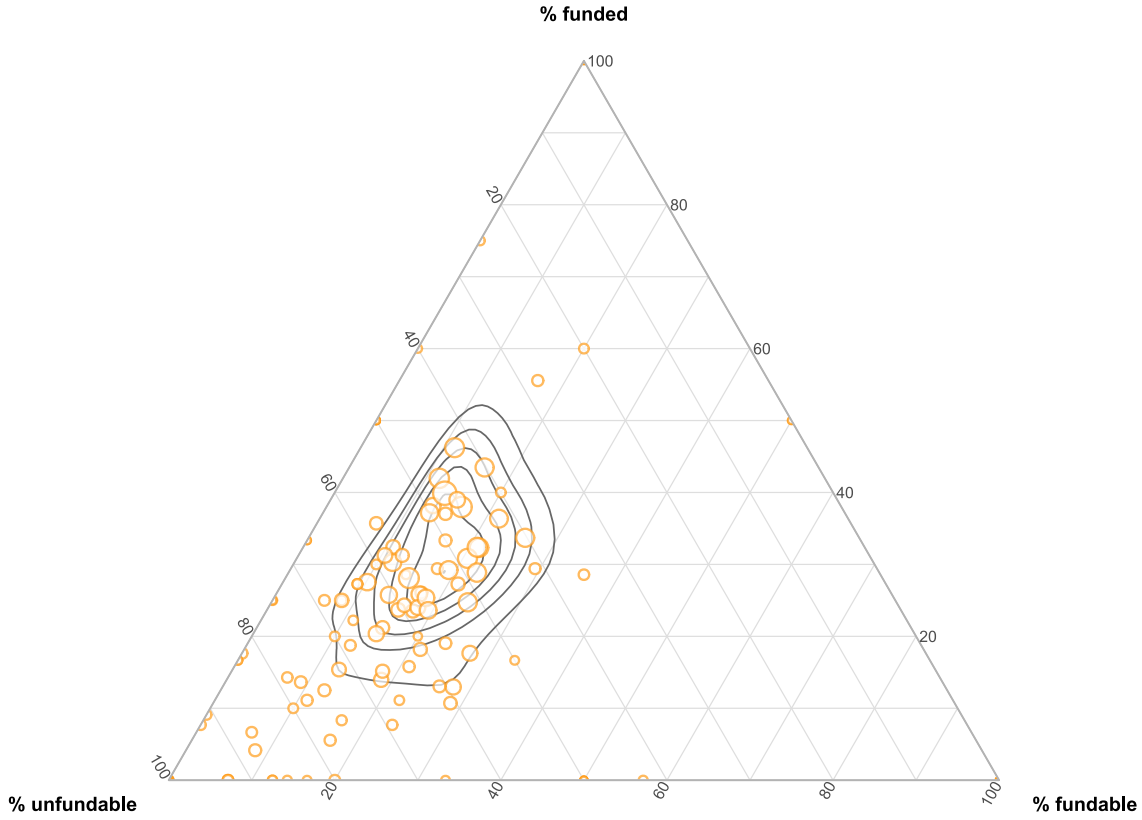

Figure 3.12: ternary diagram of organisational funded, fundable and unfundable rates over the period 2016-17 to 2018-19. Relevant axis %'s indicated as increasing from 0% to 100% towards the labelled apex. Contours show density of outcomes based on total decision volumes by organisation; circles are scaled in the same way.

Summary

Rates of any kind are potentially immensely misleading. Or at least they carry with them the potential for so many unsurfaced assumptions about what they indicate and how they are derived that discussions of them, however well-intentioned, may all too easily become incoherent and pointless.

Overall what we see has the appearance of a self-correcting system that tends to maintain rates in a narrow-ish band: a sort of homeostasis which balances a desire for higher rates with a caution on the part of potential applicants about the risk of generating too few proposals to actually attract the funding ESRC has available.

The existence of an association between success and volume is undisputable and it is seen for both the fundability rate and the conversion rate, which together create the overall rate. It seems that larger organisations have higher overall success rates both because they are more likely to produce fundable proposals and because those fundable proposals are more likely to be funded. Whether this is a sign of bias towards larger organisations, a reasonable consequence of differences in the underlying quality of the proposals, a sign of differing approaches and requirements across organisations or of some other factors we cannot yet say.

Section 4 – what is being asked for and allocated?

Research can be an expensive business and, as Figure 4.1 shows, it is getting more so over time:

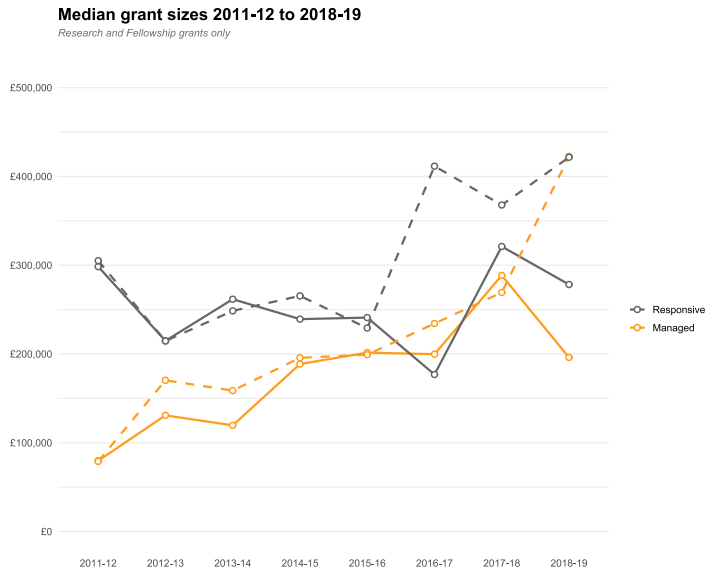

Figure 4.1: median grant size by funding mode, 2011-12 to 2018-19. Dashed lines indicate unfunded proposals, solid lines funded proposals.

The marked variability in the median value for responsive proposals is the result of the introduction and removal of differing schemes with differing size requirements, and in particular of changes in the scheme limits for the 'Open call'. Proposals submitted in relation to managed activities are much more variable in terms of the sums they request, but note that the figures in Figure 4.1 are median values which will be less influenced by the effect of a few, large, grants. Grants resulting from managed activities are in fact by far the largest of all those awarded by ESRC, and this will be the case for all the Research Councils in UKRI.

It might be tempting to assume that, with a funding volume which is fixed or nearly so, as median proposal sizes increase success rates will decrease. In fact there are so many other confounding factors at play each year (including overt demand management measures) that there is no relationship between the two variables (Figure 4.2). The journey over the last eight years has been one of gradual but not consistent drift towards larger grant sizes (the average proposal was around four times as large in 2018-19 as it was in 2011-12) accompanied by some quite random oscillation in success rates.

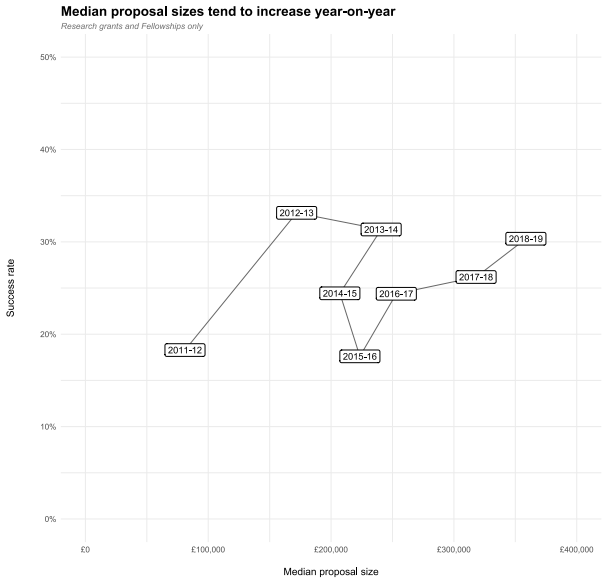

Figure 4.2: connected scatterplot of headline success rate and median proposal size across all applications (funded and unfunded) 2011-12 to 2018-19.

Figure 4.3 shows how the proportion of proposals awarded in each of six size categories has varied over time in response to changes in schemes, scheme limits, and the effects of general and sector-specific inflation:

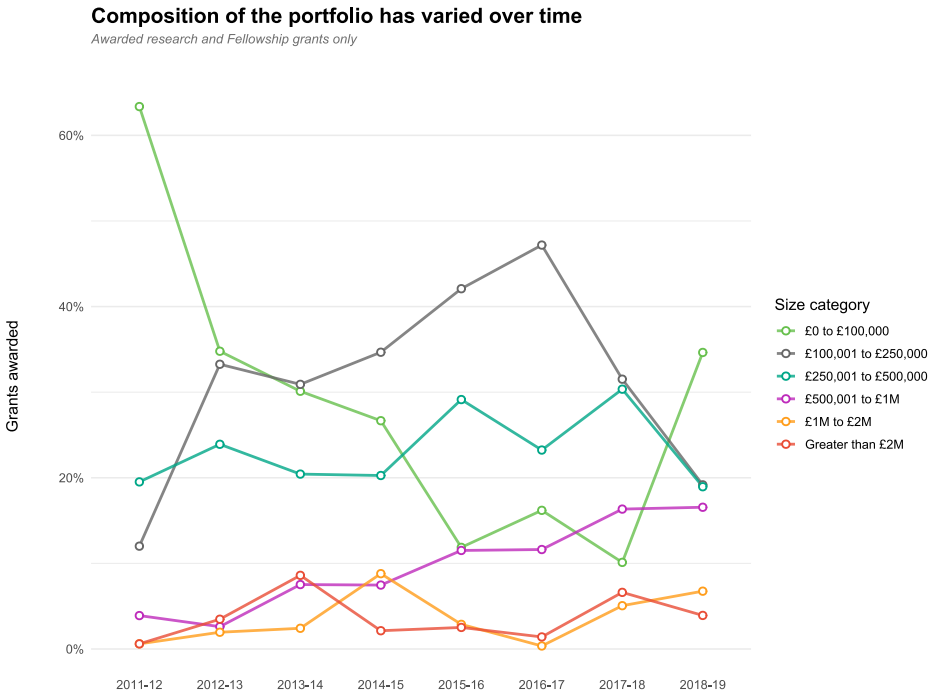

Figure 4.3: authorised grant composition by size category, 2011-12 to 2018-19.

Only one size category displays what can reasonably be described as a trend: grants in the range £500,001 to £1M have become a more common feature in the portfolio over time. Until the most recent year grants of less than £100,000 were becoming very much rarer, but 2018-19 saw the award of several postdoctoral Fellowships (an event which incidentally also had a pronounced effect on the headline success rate as the call had a measured 100% rate.) This inevitably suppressed the proportions of all other categories.

Larger grants contribute proportionally more to the composition by value of the portfolio, simply because they are larger. Figure 4.4 shows the cumulative % of the total value of grants awarded since 2011-12 contributed by each grant, when those grants are ordered by size, largest to smallest, left to right. Proportions in each value category are also indicated.

Nearly half of all funding awarded in the period 2011-12 to 2018-19 was awarded through less than 5% of grants awarded (to be precise, 46% in the 4% that fall into category F.) A further 10% was found in a further 4% of grants, meaning that more than half of all funding was allocated through less than 10% of grants[17]. At the other end of the spectrum, the smallest 30% of grants distributed just 3% of the funding.

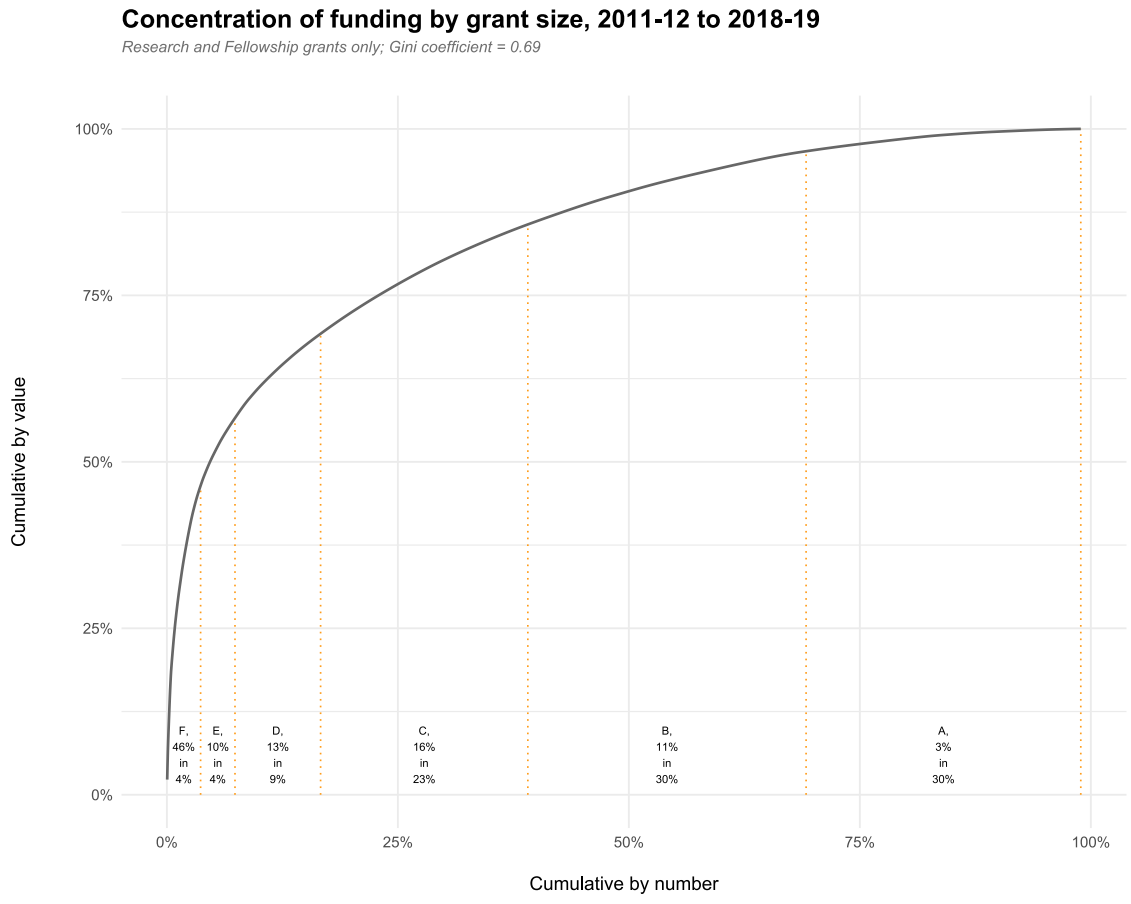

Figure 4.4: size composition of portfolio awarded 2011-12 to 2018-19. Categories: A = £1 to £100,000; B = £100,001 to £250,000; C = £250,001 to £500,000; D = £500,001 to £1M; E = £1M to £2M; F = >£2M.

[17] Note that the data excludes training grants, so across ESRC the actual concentration of resources in larger grants is greater than this.

Figure 4.4 shows the behaviour across a long time period; Figure 4.5 shows Gini coefficients for an equivalent chart of authorisations in each year:

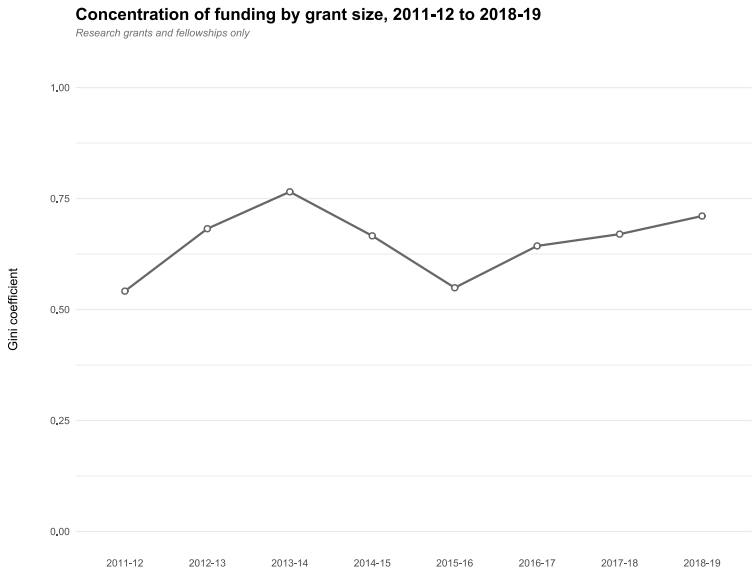

Figure 4.5: Gini coefficient by grant size for research grants and Fellowships authorised in financial years 2011-12 to 2018-19.

As might be predicted on the basis of Figure 4.3 there is no particular trend over the period, indicating that ESRC funding is no more or less concentrated in larger grants now than it has been in the past[18]. This does not mean that nothing has changed or that there is no variability in the detail. As shown in Figure 4.6, the concentration of funding by grant size is very different in managed and responsive modes, with managed funding being much more highly concentrated in fewer grants than is responsive funding. A change in the balance of funding across modes will affect the distribution of funding by grant size across ESRC, and this will undoubtedly feed into the changes seen in Figure 4.5.

The end result of all these decisions and allocations is a portfolio of grants distributed across organisations. The actual allocation of grants to organisations in the period 2016-17 to 2018-19 is shown as a treemap in Figure 4.7. This representation of the portfolio relates closely to the summary of award concentration across organisations given in Figure 2.13, and shows what a Gini coefficient of (in this case) 0.71 associated with a grant portfolio actually looks like.

[18] This is not the same as saying that grants have not got larger. A portfolio of projects each of which was ten times as large as those awarded would have exactly the same Gini coefficient. The coefficient measures a relative property of the distribution, not the absolute size of its components.

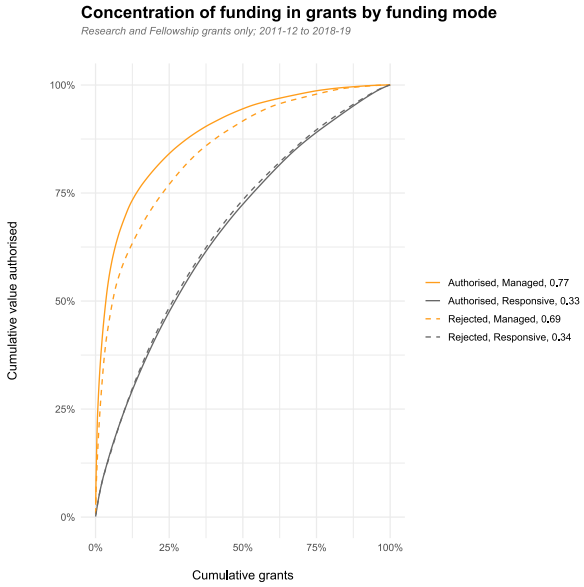

Figure 4.6: Gini coefficient by grant size for managed (orange) and responsive (grey) activities both authorised (solid lines) and rejected (dashed lines) across financial years 2011-12 to 2018-19.

Grant authorisations by research organisation and funding mode, 2016-17 to 2018-19

Research grants and fellowships only

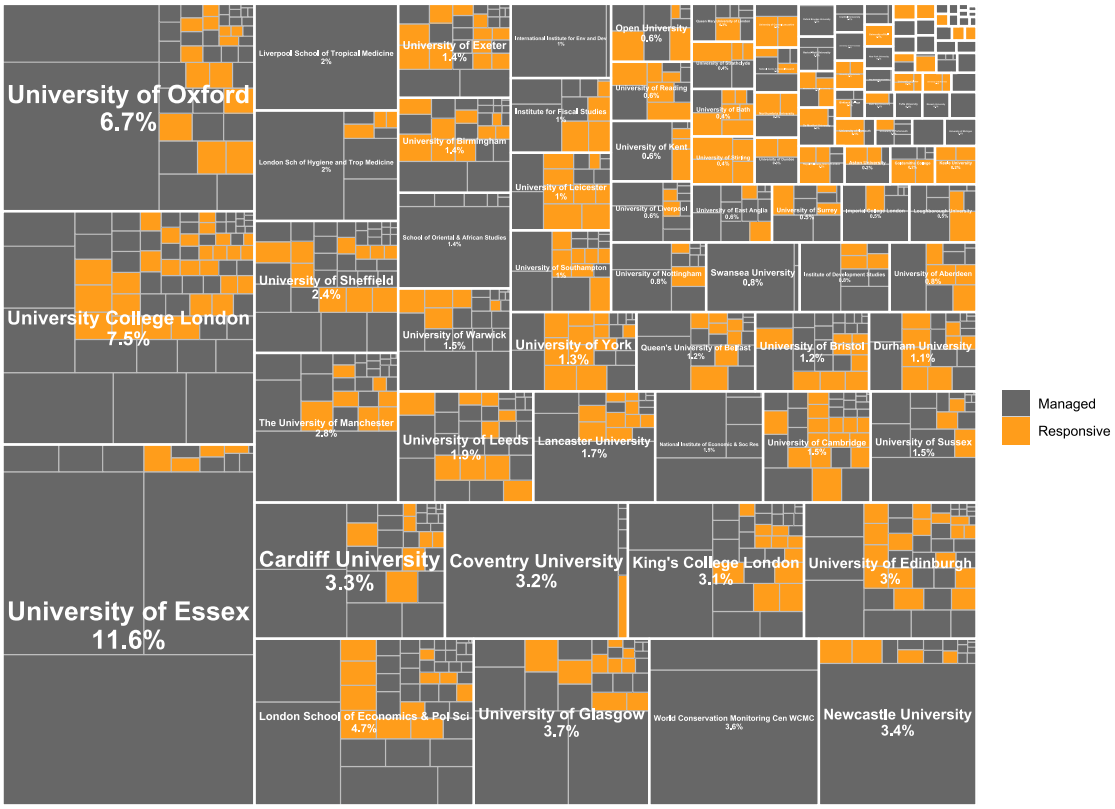

Figure 4.7: treemap of grants authorised 2016-17 to 2018-19, by organisation and funding mode. Lowest-level blocks are individual grants scaled by grant value, percentages are % of total value authorised in that period.

Clearly an organisation's grant portfolio may vary considerably in terms both of its total size, the sizes of the grants which it incorporates and its preferred mode of funding (managed or responsive). It may also vary in terms of the research that is actually supported by it, and this is the focus of the next section.

Summary

General inflation, sector-specific inflation and changes in policy combine to mean that ESRC applications and grants are getting larger on average, with smaller grants becoming generally rarer. It is in the nature of these things that a fraction of the grants allocate a majority of funding, in our case about half the funding in just under 5% of grants awarded. It's less inevitable that the degree of concentration should change over time, and for ESRC there are no signs of this happening.

Section 5 – what is being done?

So far we have seen where the funding goes, in terms of organisations and regions, how it is distributed and the outcomes of the decisions that lead to allocation. But the purpose of the funding is to carry out research, so what is the money actually being used for?

The most familiar way of describing research is with reference to its discipline. This section focuses on describing the disciplinary composition of the ESRC portfolio and how various characteristics and measures vary by discipline. 19 core disciplines are used to classify ESRC grants:

- Area Studies
- Demography
- Development studies
- Economics
- Education
- Environmental planning
- History [this is social and economic history only]
- Human Geography
- Law & legal studies
- Linguistics
- Management & business studies
- Pol. sci. & internat. studies
- Psychology
- Science and Technology Studies
- Social anthropology
- Social policy
- Social work
- Sociology
- Tools, technologies & methods

Other classifications (from a set available across UKRI) may also be used to classify ESRC grants and may feature in the analysis – when they do they are grouped as 'Other'.

Proposals may have more than one discipline allocated to them but here we focus solely on the indicated primary discipline. This approach means that, despite appearances, none of the figures given can be interpreted as 'how much ESRC allocated to discipline X', as any fractional allocations will go unreported[19].

Disciplinary portfolios may vary in size quite significantly, with some disciplines being very small indeed (Figure 5.1, which is a re-organisation on disciplinary lines of the individual grants in Figure 4.7):

[19] It's tempting to assume that any inaccuracies will cancel each other out, with an understatement of allocation through one grant being cancelled out by an overstatement on another, but this is unlikely to be the case as disciplines will tend to be associated at varying strengths. A precise explanation of spend by discipline has to deploy the clumsy wording: 'how much did ESRC allocate through grants which had discipline X as their primary discipline?'

Grant authorisations by discipline and funding mode, 2016-17 to 2018-19

Research grants and fellowships only

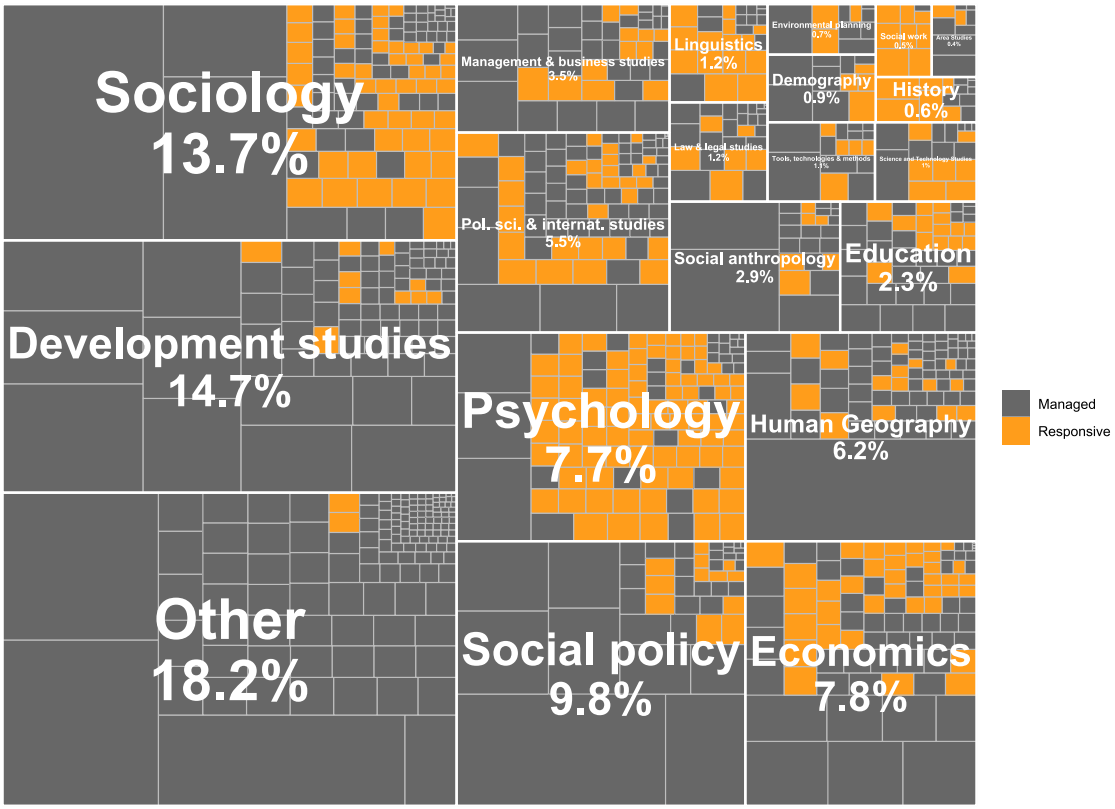

Figure 5.1: treemap of grants authorised 2016-17 to 2018-19, by primary discipline and funding mode. Lowest-level blocks are individual grants scaled by grant value, percentages are % of total value authorised in that period.

The ‘Other’ category is populated with a mixture of grants that cannot be classified against any discipline (for example Impact Acceleration Accounts) and grants that cannot readily be classified against ESRC disciplines[20].

As well as varying in size, a disciplinary portfolio may (as with an organisational portfolio) vary in its characteristics. Some disciplines rely more heavily on responsive mode than do others (this will be looked at in more detail later,) some comprise a small number of large grants while others feature a large number of small grants.

Picking out some particularly noticeable contrasts, the ‘Psychology’ portfolio is relatively homogenous in terms of grant size and is dominated by responsive mode activity, while ‘Social policy’ is the opposite, comprising a broad range of grant sizes, mostly the result of managed calls. ‘Pol. Sci & internat. studies’ received no genuinely substantial grants, while more than half of the spending on ‘Sociology’ is channelled through just three large grants[21]. The effect of GCRF on the ‘Development studies’ portfolio is obvious.

[20] It also includes grants with no classification.

[21] In fact there is one very large grant which is included here as ‘Other’ but which might better be classified under ‘Sociology’. It is worth emphasising that the imprecision inherent in such a simple description means that it is unwise to draw very specific conclusions.

Figure 5.2 shows how the proportion of awards classified against the seven largest disciplines has varied in each of the last eight years:

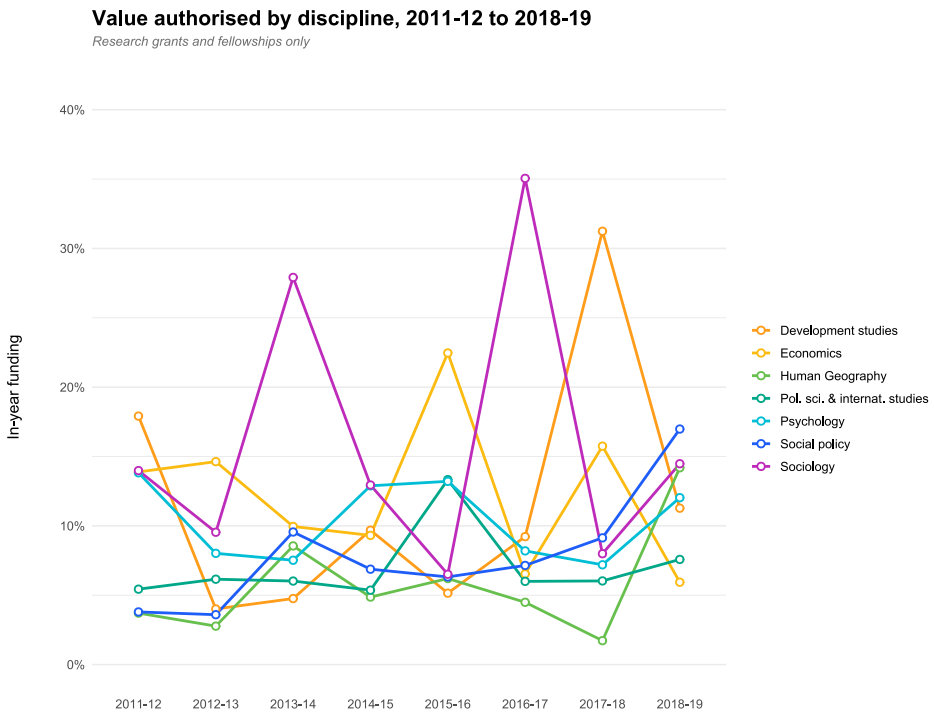

Figure 5.2: % of grants authorised (by value) by discipline 2011-12 to 2018-19 for the seven largest disciplines (excluding 'Other').

Within the limits of the data, and other than the rise of 'Development studies', there do not appear to be any meaningful trends in the composition of the portfolio over time.

A Gini coefficient summarising the distribution of grants awarded by subject suggests the same general conclusion (Figure 5.3), but it gives some indication that over time the allocation of funding is becoming more concentrated in fewer disciplines.

The change is very slight, but it is reasonably consistent. It also does not inevitably result in a portfolio which is itself more concentrated over time (although this is possible) as the disciplines in which funding is concentrated in one year will not necessarily be the same as those boosted in the subsequent year. The allocation process is becoming more 'blocky', but the portfolio will not necessarily follow the same trend.

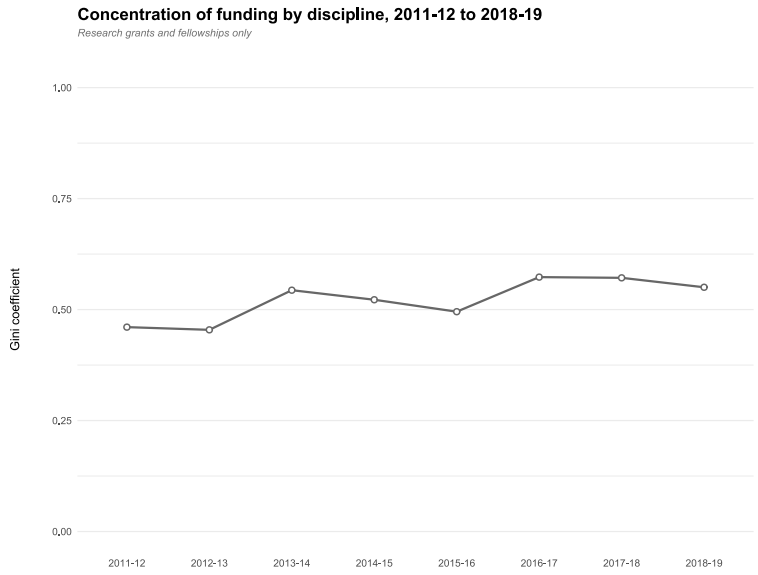

Figure 5.3: Gini coefficient for concentration of funding awarded across disciplines, 2011-12 to 2018-19.

Different disciplines apply in differing ways. Some submit larger proposals than others, and some tend to receive larger grants than others (Figure 5.4):

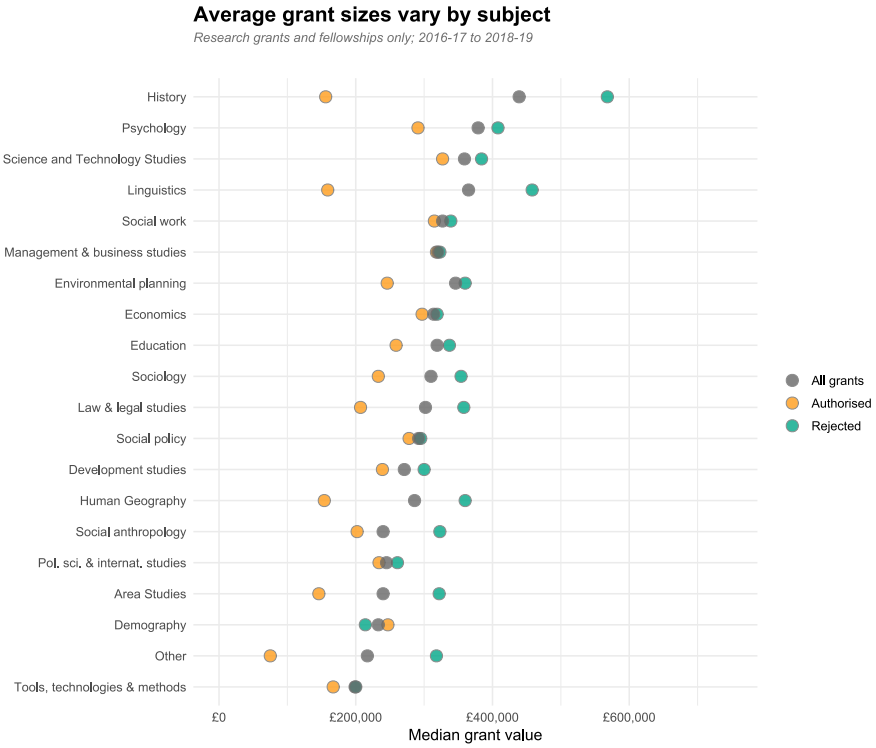

Figure 5.4: median grant size by discipline, 2016-17 to 2018-19, for all (grey), authorised (orange) and rejected (teal) grants, ordered by average of all three measures.

Research proposals tend to request on average between £200,000 and £400,000. With the exception of 'Demography', in all cases authorised grants were on average smaller than rejected grants in the last three years[22].

The top of the table is populated mostly by disciplines which rely on responsive mode funding (see also Figure 6.12.) 'History' (which in ESRC's case is limited to economic and social history rather than the broader subject) has the largest median application size but almost the smallest median authorised grant size. As a very small discipline in ESRC terms it is clearly susceptible to the specifics of individual grants.

Figure 5.5 shows the associated Gini coefficients by size for grants for each discipline, quantifying the impression of the distribution of grants by size given by Figure 5.1.

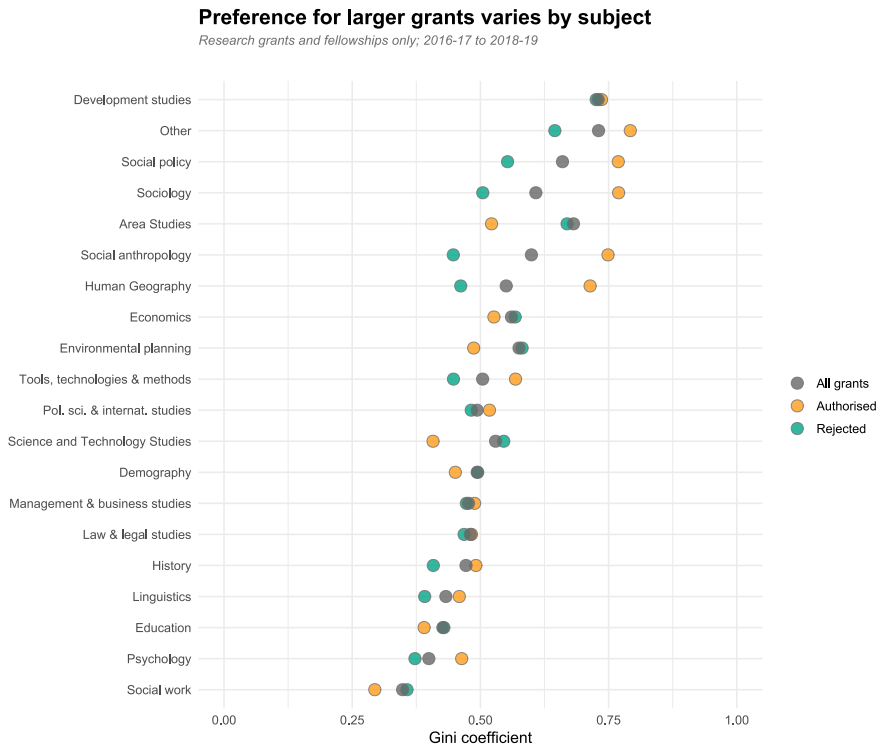

Figure 5.5: Gini coefficient by discipline, 2016-17 to 2018-19, for all (grey), authorised (orange) and rejected (teal) grants, ordered by average of all three measures.

Higher coefficients indicate a greater concentration of resource requests and/or awards within larger grants. The influence of GCRF is again apparent, with most of the highest Gini coefficients being associated with disciplines that are likely to be used to describe GCRF-funded activities. Of the larger disciplines, 'Psychology' stands out as having the most homogeneous set of grants, in terms of their size at least.

Most commonly funding awarded is more concentrated in larger grants than is the case for rejected proposals, but the balance is fine. To some extent this difference in degree of concentration, found for most disciplines, will reflect the lower success rates for responsive

[22] To a great extent this reflects the award of several postdoctoral fellowships across most disciplines in 2018-19, rather than a bias against larger grants. Another example of the specifics of a particular activity having a strong influence on what might at first be taken to be a trend or consistent truth.

activities (which are more homogenous in terms of proposal size.) Rejected proposals are enriched in grants of similar size.

Summary

Between them the disciplines of ‘Development studies’ and ‘Other’ (which is of course not a discipline) were allocated almost a third of the value of all new ESRC grants awarded in the period 2016-17 to 2018-19. The rise in prominence of ‘Development studies’ is fuelled by GCRF; the size of the ‘Other’ portfolio reflects investments in cross-cutting issues such as impact as well as cross-UKRI investments. Other disciplines’ portfolios are constructed in a range of ways, some laboriously responsive grant by responsive grant and some top-down, shaped by calls.

The data we see describes what happened. It does not explain why it happened. The extent to which the composition and genesis of each discipline’s portfolio is a matter of preference or a matter of opportunity (or lack of) is unclear, though it can safely be concluded that some disciplines are more likely to be able to identify with a particular ESRC priority than others.

Section 6 – the behaviour of organisations and disciplines

The shape of the portfolio is determined to a great extent by the choices potential applicant organisations or disciplinary groups make and are able to make. If we (ESRC) do not issue the calls, no one can apply through them. If applicants choose not to apply, they cannot be funded. If they apply a lot, they might change the shape of the portfolio (or they might not.) And if an organisation or discipline's behaviour differs from the behaviour of its peers, that behaviour may affect the shape of its portfolio.

An understanding of organisational and disciplinary behaviour and how they are linked helps when it comes to understanding ESRC's portfolio of projects, and that is the aim of this section.

Starting with some basic counting, Figure 6.1 shows the number of unique organisations applying, at least once, for funding in each of the primary disciplines[23]:

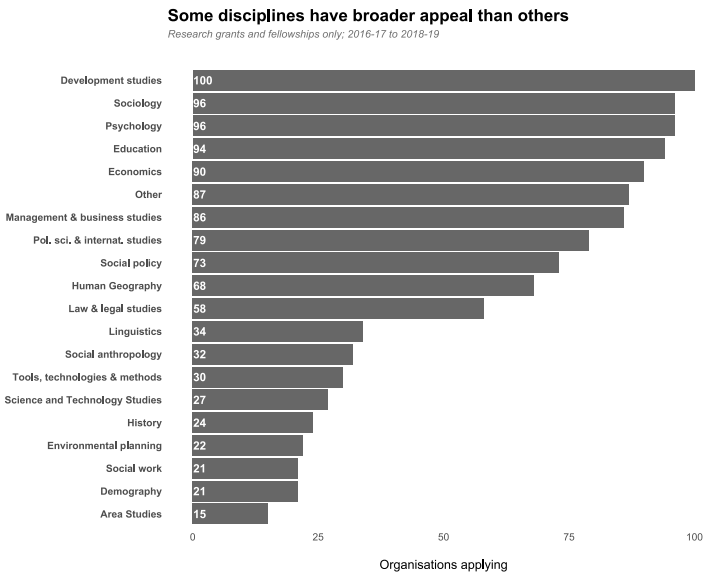

Figure 6.1: unique organisations applying by discipline, 2016-17 to 2018-19

In terms of the number of organisations applying to them, ESRC disciplines seem to be divided into two camps: those in the last three years with more than 50 interested organisations, and those with fewer than 40. 'Development studies' has the broadest appeal (perhaps because it is more relevant to non-UK organisations than are other disciplines) and 'Area studies' the least. The figures suggest an association between diversity of applicant organisations and undergraduate teaching, with some of the most popular first degree subjects featuring high on the list.

Figure 6.2 summarises the distribution of demand for funding in each ESRC discipline, in terms of the extent to which the total value of funding requests in that discipline was

[23] It is worth highlighting again that this relates solely to the organisation of the Principal Investigator on the grant that contributed to the count and ignores Co-Investigators. The number of organisations actually engaged in these grants will be greater than shown here.

distributed across all organisations applying for funding in that discipline as summarised by a Gini coefficient.

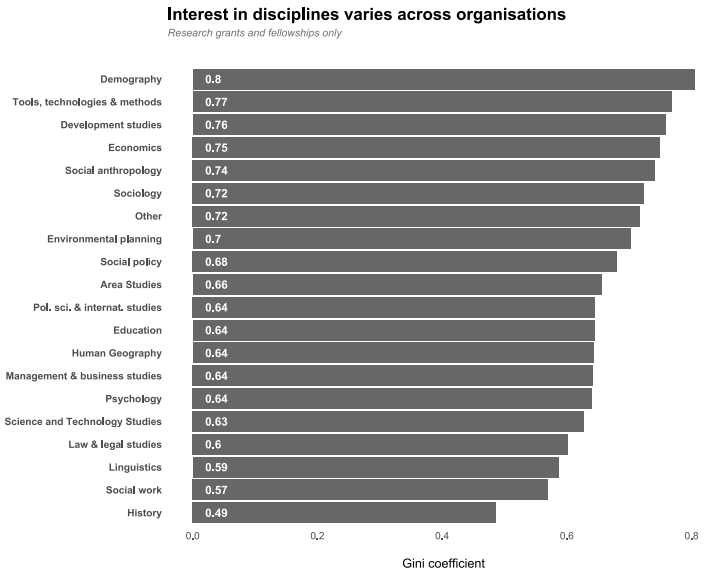

Figure 6.2: Gini coefficient for concentration of applications by value across organisations, by discipline, 2011-12 to 2018-19

In Figure 6.2 disciplines with larger coefficients see a greater proportion of their funding requests led by a smaller group of organisations. The greatest concentrations of interest are found in ‘Demography’, ‘Tools, technologies & methods’ and ‘Development studies’, each of them disciplines dominated by larger grants. The most even distributions of requests are found in ‘History’, ‘Social work’ and ‘Linguistics’ although even these have relatively high levels of concentration of demand in relatively few organisations.

As Figure 6.3 shows, most organisations submit proposals in just one or two disciplines:

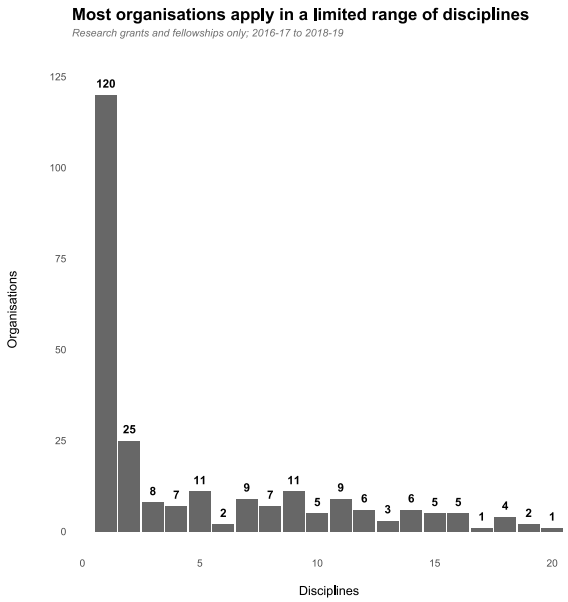

Figure 6.3: counts of organisations by number of disciplines on applications submitted 2016-17 to 2018-19.

Only three organisations applied for funding at least once in all 19 ESRC core disciplines, although nearly 50 organisations applied at least once in each of ten or more disciplines.

When organisations seek funding in a discipline, in general they apply infrequently and so have a low mean average number of proposals per discipline (Figure 6.4):

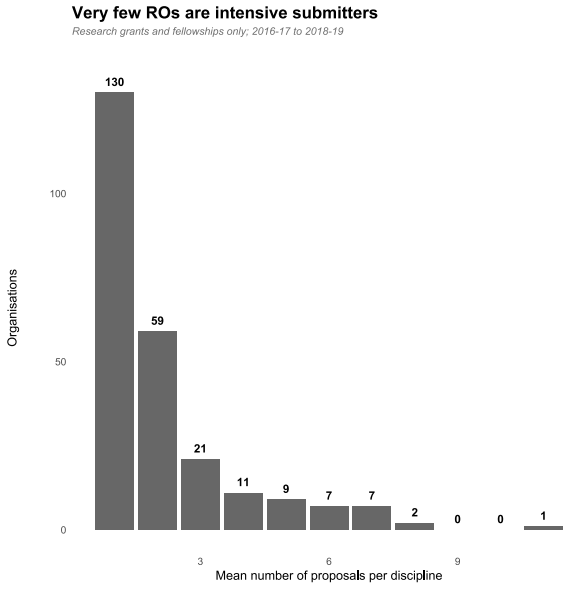

Figure 6.4: mean number of proposals submitted in each discipline by organisation, for applications submitted 2016-17 to 2018-19.

Most organisations submit just one or two proposals in each discipline in which they show an interest. At the other end of the spectrum of behaviour we see just 17 organisations which submit on average half a dozen or more proposals in each active discipline.

Some organisations submit a large (in ESRC terms) average number of proposals in a relatively few areas of interest, in effect operating in those areas as if they were much more active and prominent organisations. These are some of the 'islands of expertise' often referred to in broader discussions of research concentration[24], and they show up quite clearly in Figure 6.5:

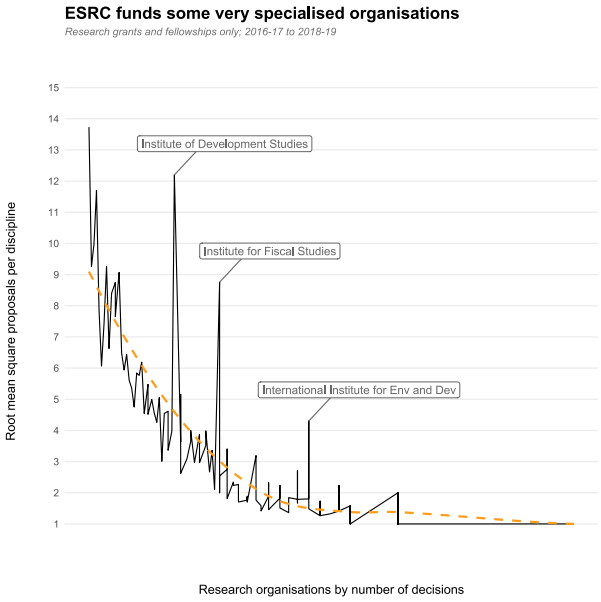

Figure 6.5: root mean square average number of proposals submitted in each discipline by organisation application total (ordered along the x axis,) for applications submitted 2016-17 to 2018-19. Dashed trend line is illustrative only.

The IDS, IFS and IIED all stand out from the background noise, having about 12, 9 and 4 proposals per discipline respectively[25]. These averages are far larger than the total number of applications each submits might suggest we would see from them.

It's not exactly surprising to find that the pattern of disciplinary interests and activity varies across applicant organisations. Figure 6.6 shows just some of the full complexity of the landscape. It is a heatmap showing the percentage of proposals within a discipline that came from each of 26 organisations which submitted at least 50 proposals in the period 2016-17 to 2018-19 (the figure of 50 is arbitrary and chosen for the purposes of the chart only.)

[24] And often in a vague, handwaving way which reflects opinion more than fact. Figure 6.5 is an attempt to see whether it is possible to detect them in a reasoned way rather than just asserting their existence.

[25] The root mean square average is used here as it is more sensitive to disciplinary concentration within an organisation than the simple average, but the figure is quite similar when using a straightforward mean for each organisation.

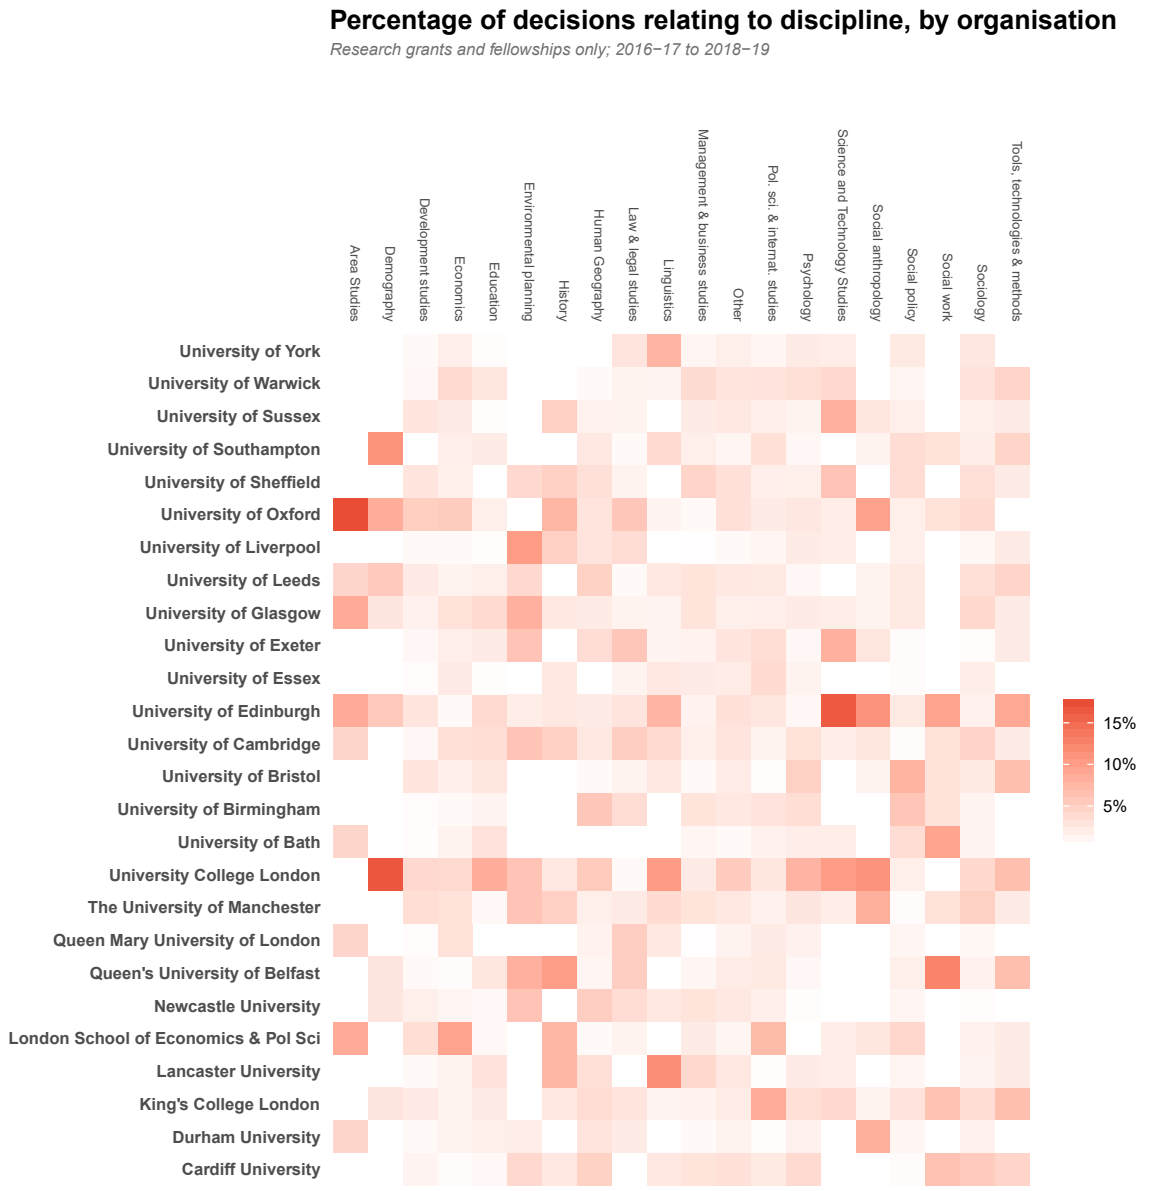

Figure 6.6: heatmap of % of all proposals in discipline submitted by organisation, for organisations submitting > 50 proposals in period 2016-17 to 2018-19. Rows and columns ordered alphabetically.

Darker rows indicate an organisation that is prominent across many disciplines. Edinburgh, Oxford and UCL stand out to some extent but really the key message is that in terms of their ESRC proposal activity even the largest organisations do not completely dominate any particular disciplinary landscape.

Columns indicate the concentration of proposals within organisations submitting proposals in that discipline. These will tend to reflect the Gini coefficients in Figure 6.2, (for example, ‘Demography’ is obviously concentrated in a relatively few organisations in both figures) but are in this case based on counts of applications rather than values. Again, the relative lack of concentration of interest in organisations is apparent (and note that the chart could be extended for a further hundred or so rows.)

Despite this complexity one simple but rather unsurprising rule of thumb holds quite well: the more organisations there are interested in a discipline, the more proposals ESRC sees relating to that discipline (Figure 6.7):

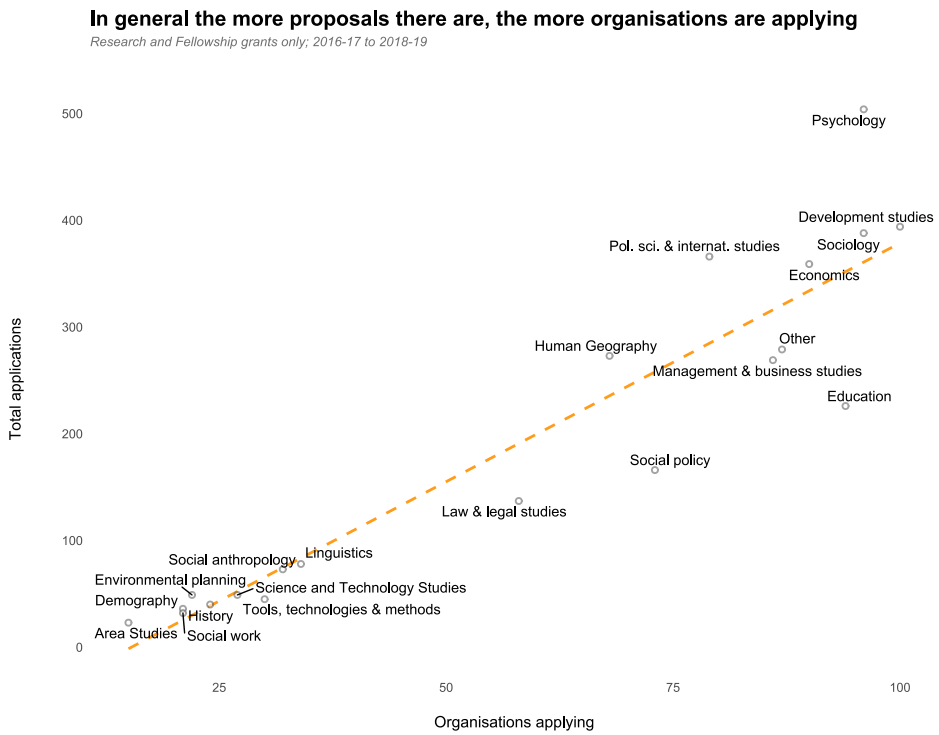

Figure 6.7: number of proposals received vs number of organisations applying, by discipline, in period 2016-17 to 2018-19.

While most of the variability in disciplinary application numbers is associated with variability in the number of organisations applying ($R^2 = 0.85$, $p < .001$), some disciplines behave differently. 'Psychology' submits many more proposals than the number of organisations with an interest in it might suggest, as does (to a lesser extent) 'Pol. sci. & internat. studies'. 'Education' and 'Social policy' both under-submit relative to the number of organisations with an interest in them.

While the building blocks of the analysis so far are simple – organisations applying for particular numbers of projects in particular disciplines – they can be combined in complicated ways. Some simplification will be helpful. If we define, for each organisation, a *d* index which is the largest number such that in a given period that organisation has applied at least *d* times in each of at least *d* disciplines[26], we can combine the concepts of breadth and scale of activity into a single figure for each organisation. Figure 6.8 shows the distribution of *d* indices across all organisations in the period 2016-17 to 2018-19.

[26]Obviously this is the well-known *h* index, but applied to proposals in disciplines rather than citations of publications. The index will change on the basis of the time period it covers and so is not an unvarying measure – it is useful for comparisons only. *d* indices for 50 of the largest organisations are given in annex A.

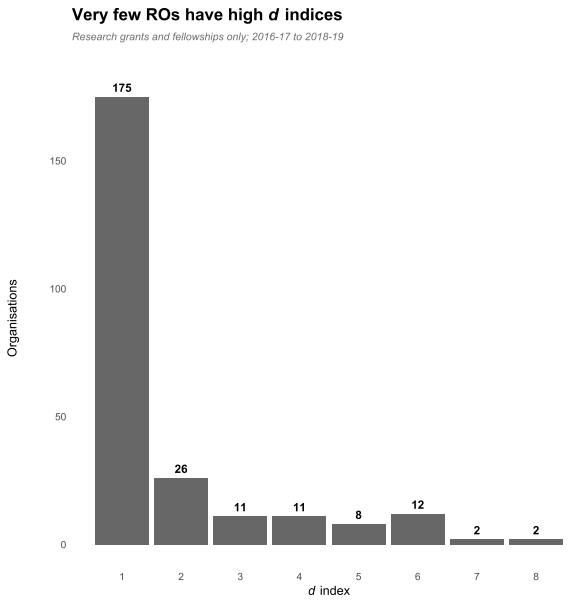

Figure 6.8: organisational *d* indices in the period 2016-17 to 2018-19.

By this measure the combination of breadth and volume of proposal activity is relatively rare, the vast majority of organisations having a *d* index of just one. Oxford and UCL had the highest *d* indices (both 8) with Cambridge and Sheffield coming just slightly lower. The same logic can be applied to funded grants only, to give a *df* index (Figure 6.9):

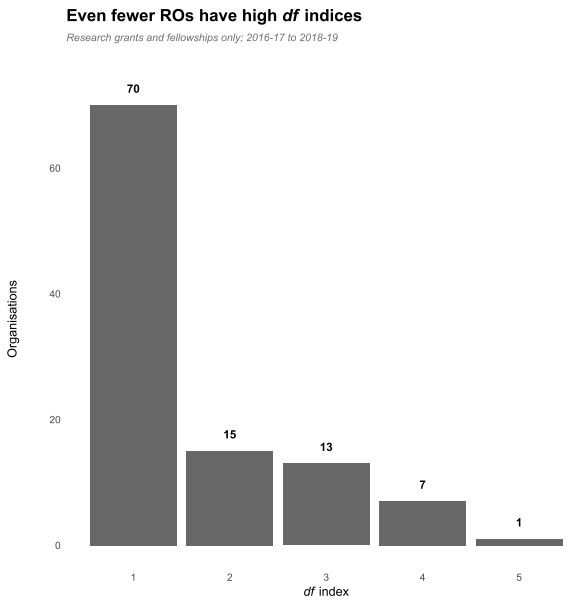

Figure 6.9: organisational *df* indices in the period 2016-17 to 2018-19.

Again, one funded proposal in one discipline is the most common organisational outcome, although one organisation (UCL) managed to secure at least five grants in each of at least five disciplines.

Finally, when looking at the behaviour of organisations it is worth restating that their interest in responsive mode funding varies widely (Figure 6.10):

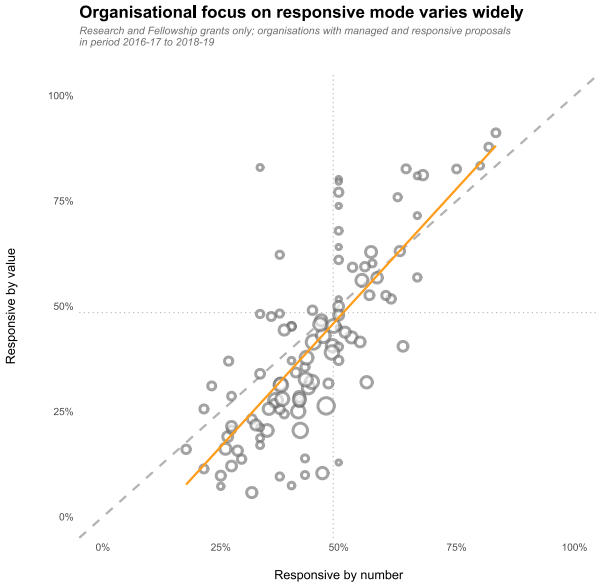

Figure 6.10: % of grants by value through responsive mode vs. % by number, by organisation in the period 2016-17 to 2018-19. Points scaled by total decision volume. Red line is fitted slope (for illustration only), dashed line is line of equality.

In general reliance on responsive activity by value scales with reliance on it by number. Some of the more active organisations request a lower proportion of their funding from responsive schemes than might be expected based on their responsive volumes. These organisations will be those applying for particularly large grants and the pattern thus suggests some divergent behaviour. None of the most active applicant organisations has a particularly high reliance on responsive mode (See Figure 3.9).

Disciplines can have an associated success rate just as logically as can organisations, and a funnel plot constructed in the same way as that in Figure 3.7 can be produced for disciplines (Figure 6.11):

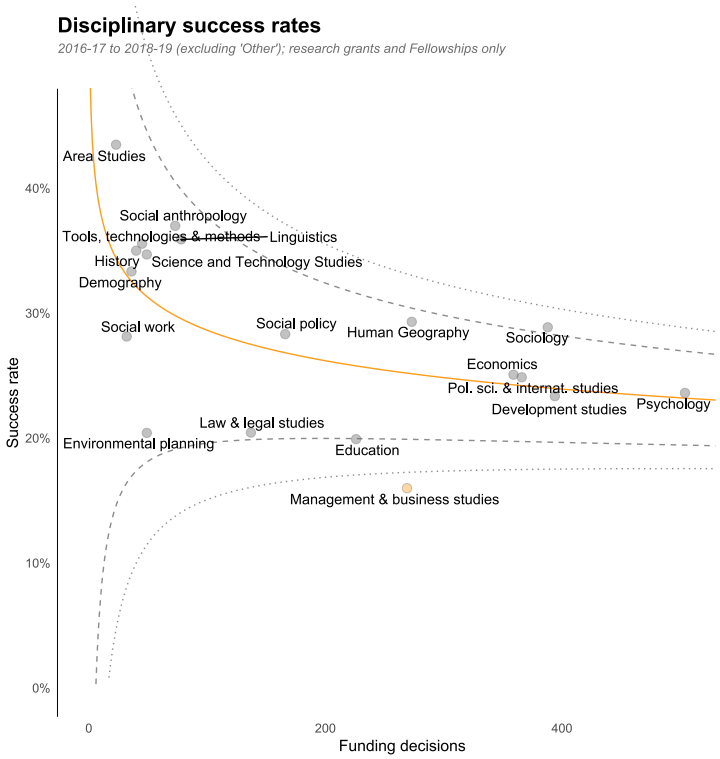

Figure 6.11: funnel plot by discipline (excluding 'Other'[27]) in the period 2016-17 to 2018-19. Inner, dashed, funnel shows a 95% control limit, outer (dotted) funnel shows a 99+% control limit.

There is a negative association between decision volume and success rate by discipline, and only one outlier: 'Management & business studies', which has an unexpectedly low success rate.

This of course does not necessarily mean that submitting lots of proposals lowers some underlying disciplinary success rate. The individual characteristic behaviours of disciplines are again in play. Specifically, the relatively low (but not unexpectedly low) success rate seen for 'Psychology' reflects to a great extent the fact that, as Figure 6.12 shows, as a discipline 'Psychology' is unusually reliant on responsive mode funding:

[27] The exclusion of 'Other' has a material effect on the position and shape of the funnel as the chart is derived entirely from the data included in it.

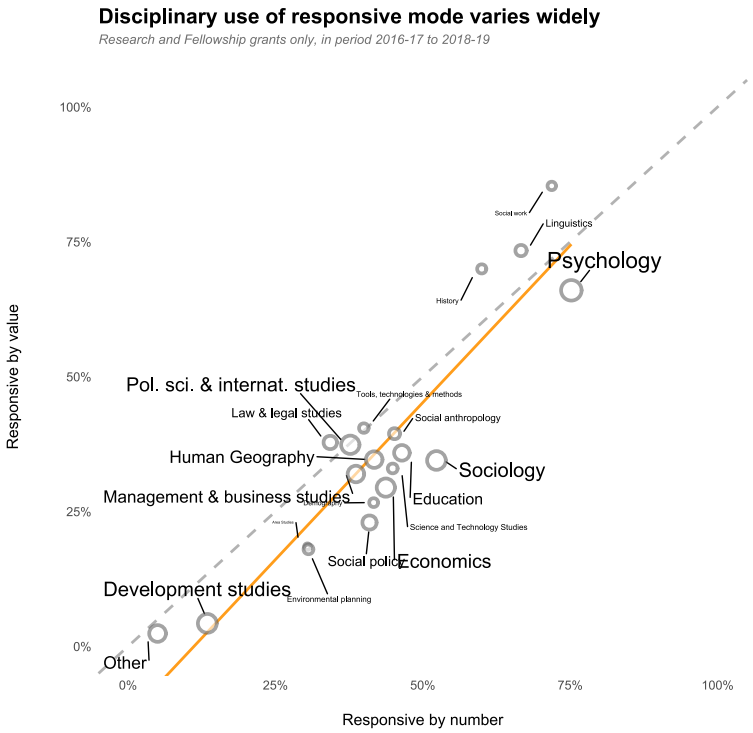

Figure 6.12: % of grants by value submitted to responsive mode vs. % by number, by discipline in the period 2016-17 to 2018-19. Points scaled by total decision volume.

While most disciplines, including most of the larger ones, submit around 40-50% of their proposals by both number and value to responsive mode calls, 'Psychology' submits closer to three quarters of its total. And as responsive schemes tend to have lower success rates than managed activities, this will suppress the disciplinary success rate significantly.

This argument is of course circular. A scheme or mode does not simply 'have' a lower success rate – its rate reflects the balance of the number of grants awarded in relation to the number of proposals submitted, and a scheme or mode which is a particular focus for a discipline will, if that discipline is large, have a suppressed success rate because the funding available for that scheme or mode is limited. It is the balance between the total funding requested by a discipline and the total value potentially available in the schemes or modes to which it applies that shapes the outcome. This is shown in Figure 6.13, as the ratio of the total amount awarded through all calls to which a discipline applied to the total amount requested by that discipline in those calls.

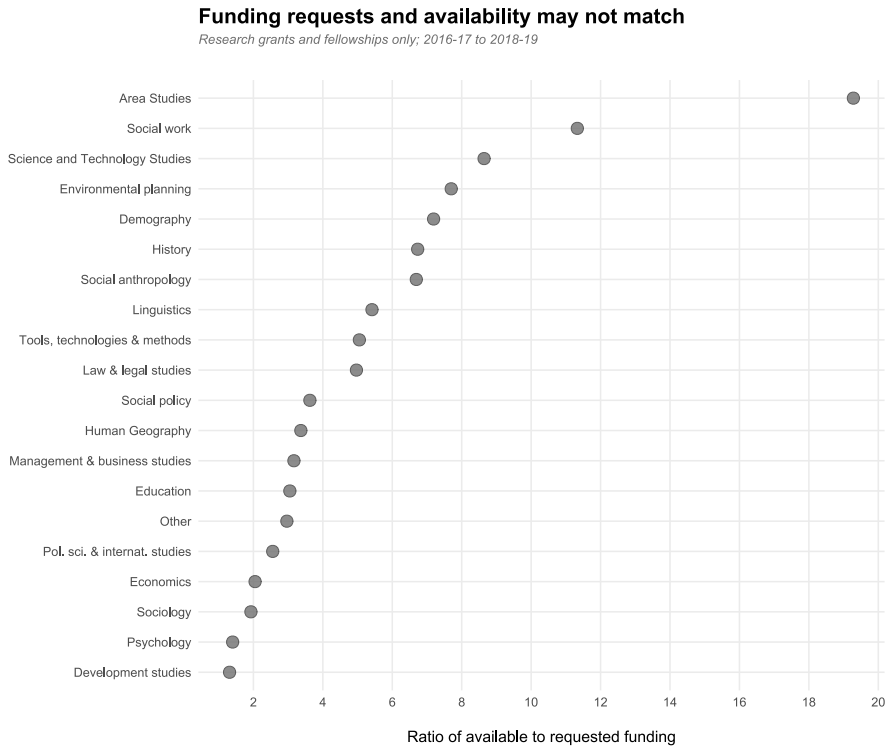

Figure 6.13: ratio of total funding awarded in all calls applied to by a discipline to total funding requested across all calls by that discipline in the period 2016-17 to 2018-19.

A large fish in a small pond may struggle to satisfy itself, and ESRC has some very large fish swimming in some relatively small ponds. For example 'Psychology' (40%,) 'Sociology' (34%) and 'Economics' (33%) each requested a substantial fraction of all the funding that was in principle available to that discipline. The size of the ponds is not a given and will reflect both disciplinary preferences or behaviours and the opportunities that ESRC decides to offer[28].

Figure 6.14 shows how the extent to which a discipline's requests for funding are concentrated in just a few calls, the ratio of the funding it requests to the funding available and the proportion of proposals that it submits to responsive mode relate. It is by definition rather a complicated chart, but then the behaviour it tries to unpick is also rather complicated.

[28] With both the funder and the discipline asking why the other doesn't simply make the pond in question bigger, but in entirely different ways.

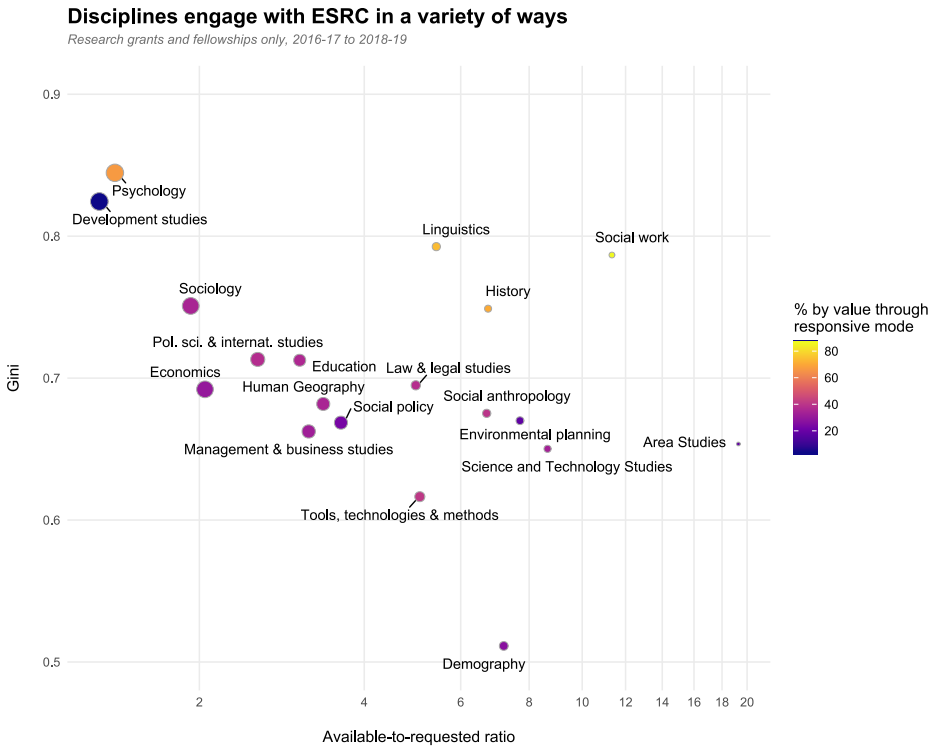

Figure 6.14: Gini coefficient summarising concentration of funding requests by call *vs.* ratio of total funding awarded in all calls applied to by a discipline to total funding requested across all calls by that discipline in the period (log scale) 2016-17 to 2018-19. Percent of funding requested that was requested via responsive mode calls shown by colour.

Disciplines may be grouped in the following way:

- Disciplines which focus on just a few calls, have a strong focus on responsive mode and which request a significant volume of funding relative to the level of funding available
    - Psychology
- Disciplines which focus on just a few calls, have a strong focus on responsive mode and which request a small volume of funding relative to the level of funding available
    - History, Linguistics, Social work
- Disciplines which focus on just a few calls, have a strong focus on managed mode and which request a significant volume of funding relative to the level of funding available
    - Development studies
- Disciplines which respond to a range of calls, balance funding across responsive and managed funding and which request a significant volume of funding relative to the level of funding available
    - Economics, Education, Sociology, Pol. Sci. & Internat. Studies
- Disciplines which respond to a range of calls, balance funding across responsive and managed funding and which request a small volume of funding relative to the level of funding available
    - All the rest, with Demography and Area Studies at the extremes.

What other consequences might a focus on responsive mode have for a discipline? As Figure 6.15 shows, the answer is few or none. Within responsive mode there is no success-volume relationship. There are also few or no disciplines that have unexpected responsive mode success rates (which, in the absence of a success-volume relationship, means few or none that have rates that differ significantly as a result of chance variation from the overall average.)

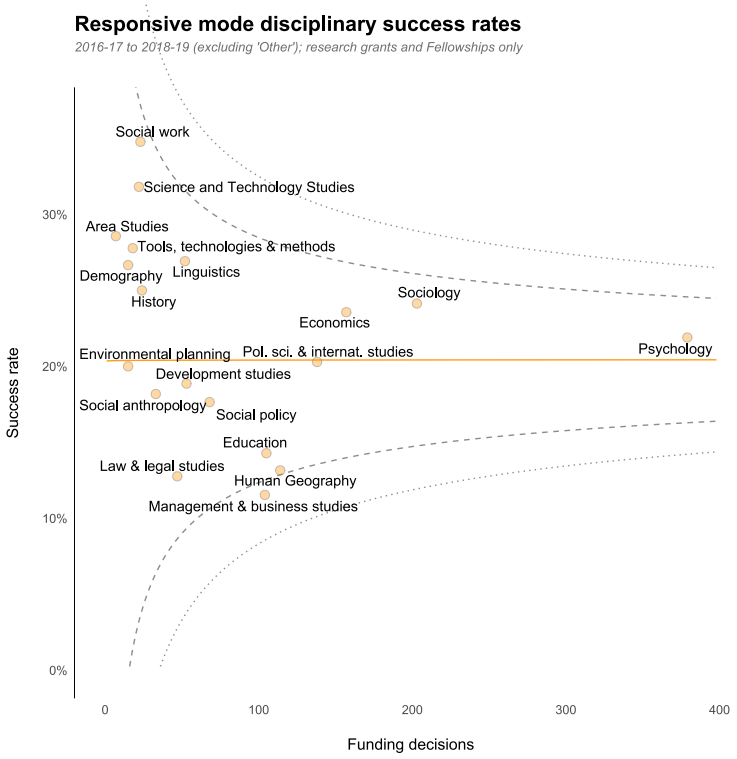

Figure 6.15: funnel plot by discipline in the period 2016-17 to 2018-19. Inner, dashed, funnel shows a 95% control limit, outer (dotted) funnel shows a 99+% control limit.

The one possible exception is again ‘Management & business studies’. The overall low success rate for MBS has a component which is derived from responsive mode activities (although MBS does not have an unusual reliance on them.)

Summary

Only the very largest institutions generate ESRC proposals in a way that might be described as a flow. For the most part, the application stream as a whole is shallow and broad, with only the smallest deep-and-rapid channel at its centre. Most organisations only see ESRC funding as a requirement for some things. The conditions are ideal for the portfolio to meander slowly rather than undergo rapid change, and this is just what we see.

Interest in disciplines is manifested across the UK and beyond. Few disciplines are the preserve of just a handful of places or of very niche interest. ESRC does not need to worry that if an organisation stops applying for funding, a particular discipline will disappear from its portfolio (though the details and quality of what is left might of course vary.)

Section 7 – organisational groupings and ‘fairness’

Much of this document is about differences in inputs, outcomes and distributions across the funding landscape. This section is different as it tackles a specific question: whether there are any differences which arise as a result of (or perhaps are merely associated with) membership of the UK research base’s most prominent organisational groupings. These are of course the Russell Group[29] and what is usually defined as a subset of the Russell Group: the ‘golden triangle’[30].

The geographic concentration of ESRC funding in London and the south east of England seen in Figures 1.2 and 1.3 is, as already noted, primarily a result of the placement of the organisations that we fund. Whether we would fund them if they were physically located elsewhere is a moot point. What we do know is that, aside from the exceptional case of the University of Essex, most of the recipients of the greatest amounts of ESRC funding are members of one or both groups (Figure 7.1):

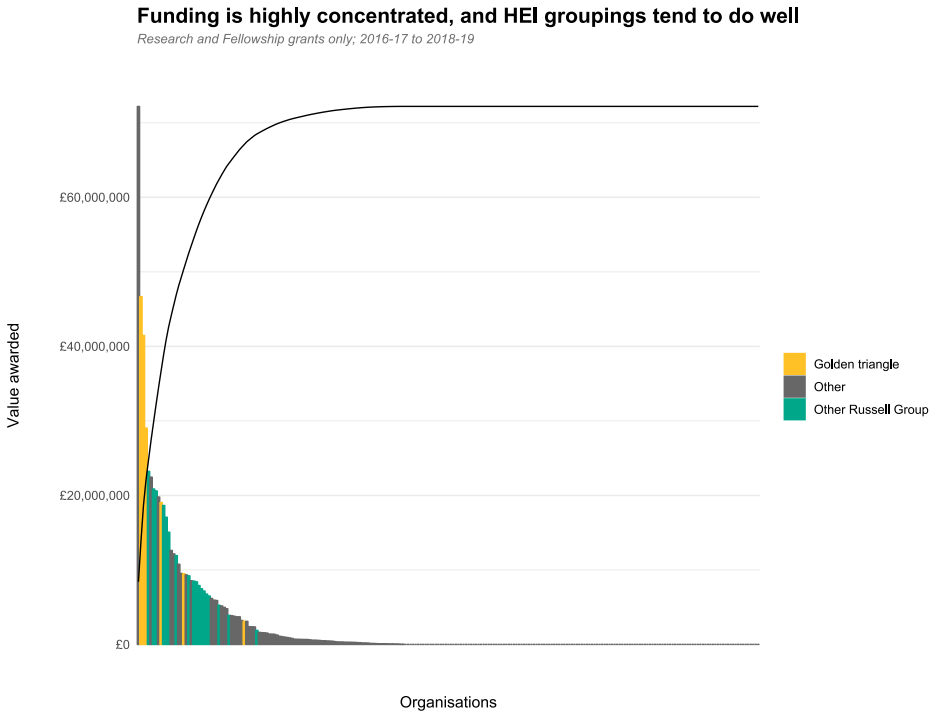

Figure 7.1: funding awarded by organisation in the period 2016-17 to 2018-19, organisations ordered from largest to smallest total value awarded. Cumulative % of total shown by black line.

Of the members of the golden triangle, Imperial College has by far the smallest ESRC portfolio. The smallest Russell Group ESRC portfolio is that of QMUL (an organisation which, as already noted, has an unusually low conversion rate.) Other than the University of

[29] Membership can be found at https://russellgroup.ac.uk/; the list of members used for this analysis was the one live in June 2019.
[30] Usually, and in this analysis, taken to comprise Cambridge, Imperial College, KCL, LSE, Oxford and UCL. Purely geographic interpretations of the name which include organisations that no one would accuse of benefitting from organisational favouritism make little sense.

Essex, the organisations with the largest ESRC portfolios that are members of neither group are ranked highly almost solely because of funding awarded through GCRF. Without that funding it is likely that the largest non-group member other than Essex would be Lancaster, ranked about 12th. 50% of the total awarded in the period 2016-17 to 2018-19 went to the top 10 organisations, four of which were golden triangle organisations.

We would expect many of the organisations in these groups to receive a significant amount of research funding from ESRC (or any other funder) simply because they are large. The question really is whether there are any factors other than scale that determine their take. As Figure 7.2 shows, in general when organisations within the golden triangle and the broader Russell Group request funding, they request more than do organisations outside those groups. The same difference is apparent for rejected proposals, but not for authorised grants[31].

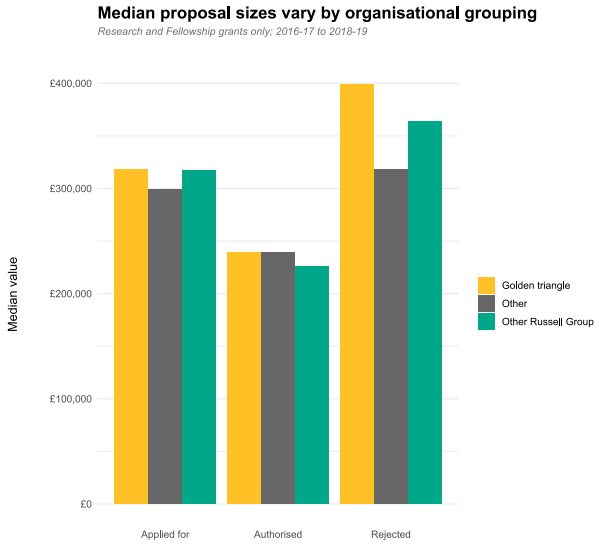

Figure 7.2: median applied for, authorised and rejected grant sizes by group membership in the period 2016-17 to 2018-19

As golden triangle and other Russell Group organisations are relatively large (in terms of their decision volumes) their success rates benefit from the underlying positive association between success rate and decision volume[32] (Figure 7.3):

[31] The fact that the authorised value for the Russell Group is actually *less* than that for non-group members is partly explained by the award of a large number of relatively small postdoctoral Fellowships in 2018-19 which were allocated within ESRC DTPs (which are concentrated in the Russell Group) and also the award of funding to IAAs. Overall it is reasonable to conclude that there is little difference in the median authorised value across organisations, but this conclusion is sensitive to the specifics of the activities to which it relates.

[32] As the success-volume relationship exists throughout the set of organisations it is safe to say that it is not caused by the existence of these groups. It may be amplified by it, but it would exist even without a group effect.

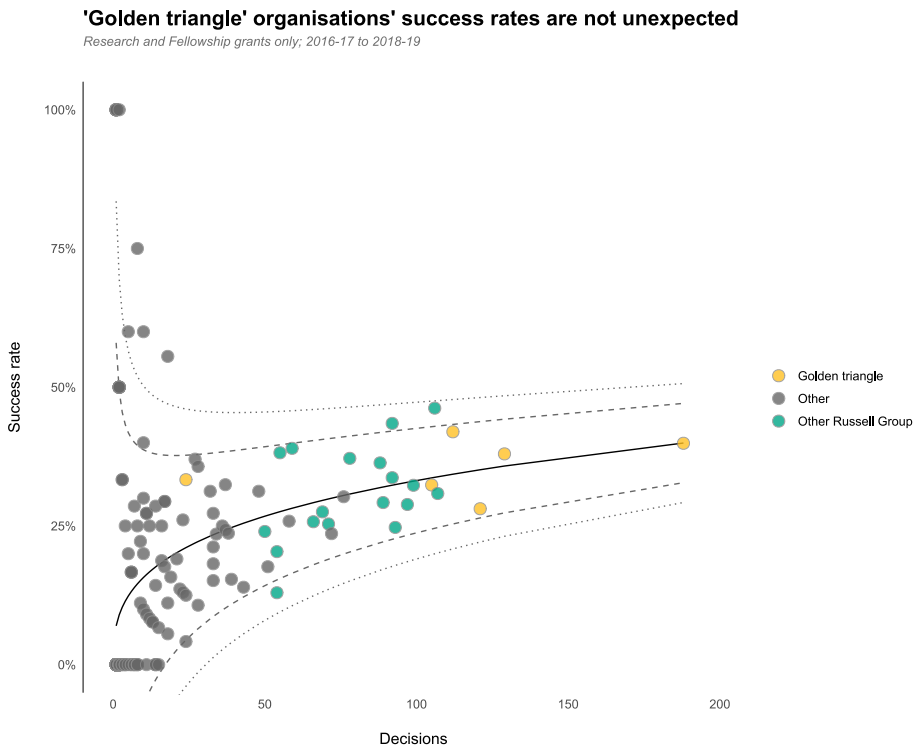

Figure 7.3: funnel plot of organisational success rates over the period 2016-17 to 2018-19 by organisational grouping (same data as Figure 3.7). Inner, dashed, funnel shows a 95% control limit, outer (dotted) funnel shows a 99+% control limit.

No golden triangle organisations have unexpectedly high success rates. Depending on the definition of ‘unexpectedly’ used this may also be true for the broader Russell Group, but it is worth noting that two organisations (Edinburgh and Sheffield) have rates that sit above the 95% control limit in Figure 7.3, with Edinburgh’s being very far outside it. The only organisation of any kind below the lower 95% control limit is actually a Russell Group organisation (QMUL.)

That’s the broader picture, but we can more directly ask whether the success rates of the three groups differ. Figure 7.4 shows that they do, both in the three possible inter-group combinations and also in a comparison of the Russell Group as a whole with all other organisations.

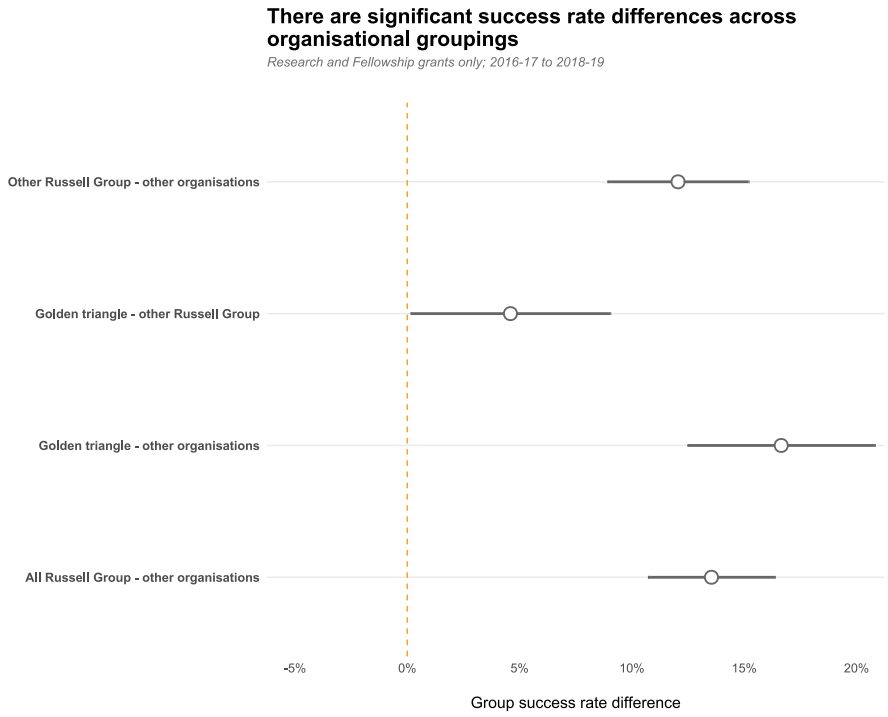

Figure 7.4: inter-group differences in overall success rate, 2016-17 to 2018-19. 95% confidence intervals for difference in rates shown for indicative purposes.

Success rates of the golden triangle tend to be higher than those of the other members of the Russell Group, and the collective rate of the other members of the Russell Group tends to be higher than that of all other organisations. It follows logically that the golden triangle also has a higher rate than organisations which are not members of the Russell Group, and that, as a whole, the Russell Group rate is higher than that of non-Russell Group members.

Taken together the data suggests that organisations in either the golden triangle or the broader Russell Group apply more frequently, and for more funding, than do other organisations. When they apply, they tend to be more likely to receive funding than are other organisations. But the success rate enhancement is a manifestation of a deeper relationship between application volume and success rate, where causality is unclear. We do not know whether more active organisations have higher success rates, or whether organisations which fundamentally have higher success rates are more active, or whether some other factor is driving what we see.

It is however possible to see whether the success-volume relationship is changing over time. The specific form of the relationship is such that a coefficient, *b*, is an indicator of the strength of that relationship, with larger values of *b* indicating a stronger relationship[33]. Figure 7.5 shows how the value of *b* associated with organisational overall and fundability rates[34] for ESRC proposals has changed over time:

[33] See annex B
[34] The same data for the conversion rate is not shown for technical reasons.

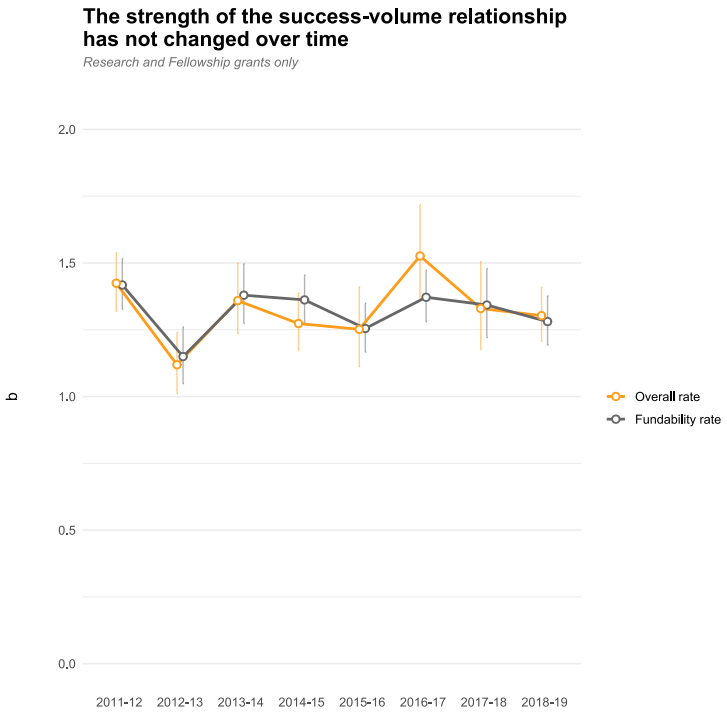

Figure 7.5: coefficient *b* from nonlinear least-squares regression of overall success and fundability rates onto organisational decision volume (as described in the annex) over the period 2011-12 to 2018-19. 95% confidence intervals shown.

While it varies slightly year to year, the relationship between these two rates and decision volume has neither strengthened nor weakened over time, and so neither has the extent to which these organisational groupings benefit from it.

An effect of organisational grouping might also manifest itself in the level of funding allocated per researcher in each organisation. While it is on the face of it simple to determine this using published counts of researchers, in reality that simplistic approach is flawed as it assumes that everyone in a research-related role in an eligible organisation is a potential applicant. This is clearly not the case.

Some researchers who are, on paper, potential applicants may in fact secure all the funding they need from other sources; others may apply so infrequently that they are in effect not countable; while others may simply not feel the need to apply or believe that the potential rewards do not justify the effort. These outcomes may be more likely in social science than many other areas of research[35] as in general the costs of research may be lower, meaning that it is more viable to have an ESRC-free career in social science than it is to avoid applying, for example, to EPSRC when working in the field of condensed matter physics. In general, published data on academic headcounts are least useful in less resource-intensive subjects[36].

[35] And they may of course be a function of group membership, which introduces an element of circularity into the argument.
[36] The situation is even worse for ESRC as ESRC funds many non-HEIs, allows international investigators and is the main conduit for GCRF grants, in principle making the entire global population of social scientists 'potential applicants' and emphasising the need to be discriminating when identifying the applicant population.

Thankfully the data we have can provide an estimate of the actual population of researchers who genuinely are potential ESRC applicants[37]. Based on applications made in the financial years 2011-12 to 2017-18, the total number of potential ESRC applicants active globally over the whole period was at least 24,000[38]. At the organisational level Oxford had the highest number of potential ESRC applicants, perhaps more than 1,300, with UCL (~900) and Cambridge (~800) some way behind. The organisation with the largest population of potential ESRC applicants but which is not a member of either the golden triangle or the Russell Group more broadly was Lancaster (~600.)

Based on these figures, a more realistic estimate of the value awarded per potential PI, by recipient organisation, in the period is shown in Figure 7.6:

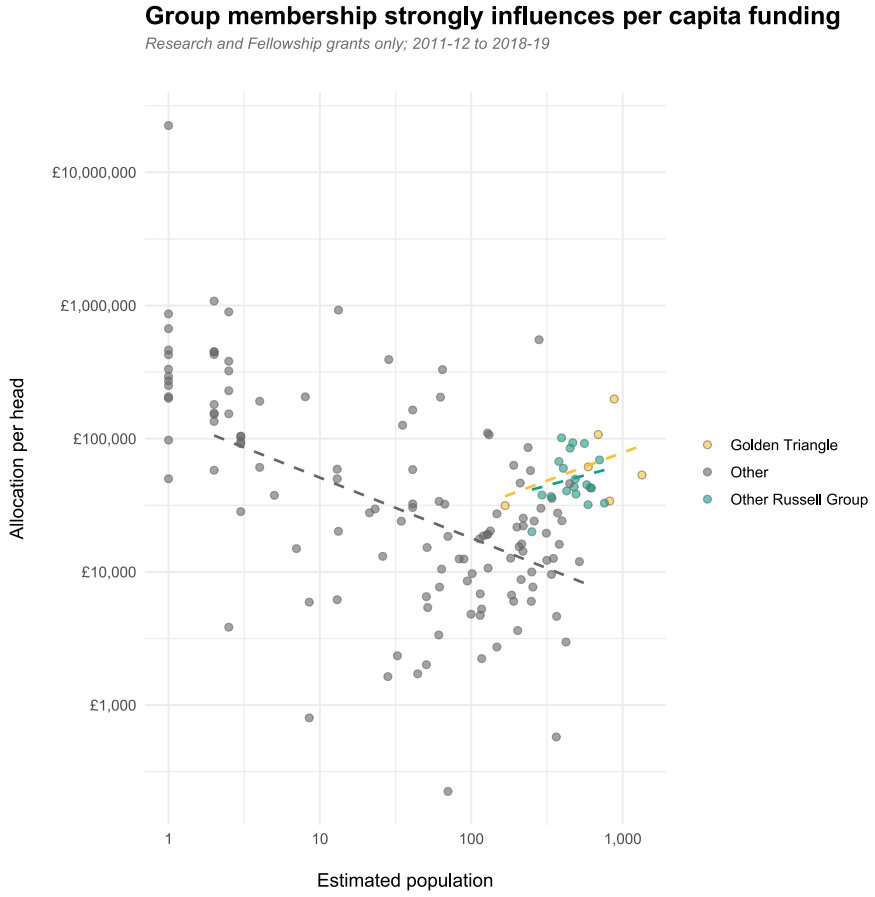

Figure 7.6: per capita (PI) value awarded against estimated applicant population (log scales,) 2011-12 to 2018-19, by organisation. Dashed lines are in-group regressions; group membership indicated by colour.

[37] Using the method of Chao found at http://chao.stat.nthu.edu.tw/wordpress/paper/1987_biometrics_43_P783.pdf. The estimate that results is a lower bound.

[38] The data also allows an estimate of the average time between applications for potential applicants. If the probability of a person applying is *p*, the probability of them not applying after *n* years is 1-(1-*p*)^*n*, which drops below 0.5 after about six to seven years – the average time between applications. More broadly, if about a third of applicants applied in the eight years covered by the data, the average time between applications is not less than five years. Both figures suggest a highly intermittent application pattern among potential applicants for ESRC funding.

The fact that the results of ESRC's processes reflect specific activities rather than a homogenous process is again apparent. There are some very large per capita allocations in the group of organisations which received just one award.

Even excluding these, overall there is for organisations which are members of neither the golden triangle nor the Russell Group more broadly a notable and significant ($p < .05$) negative association between the estimated number of applicants in an organisation and its per capita allocation[39]. For organisations within those groups, there is not a significant association between per capita allocation and estimated population[40]. Median per capita allocations were £57,000, £46,000 and £16,000 for the golden triangle, the broader Russell Group and other organisations respectively.

What does this mean? Organisations which have greater numbers of potential applicants tend to have lower levels of per capita funding, suggesting a sort of brake on total funding which might be awarded to an institution – unless that institution is a member of the Russell Group or within the golden triangle. Group members have per capita funding levels much higher than we would expect were they to behave in line with organisations which have comparable numbers of potential applicants. Whether this is a result of group membership (for example perception of greater competence) or is caused by something else which is associated with group membership (for example actual competence) cannot be determined from the data.

The propensity that each potential applicant has to actually apply will be a further influence on concentration of funding. For equivalent organisations, if the researchers in one are more likely to apply than the researchers in the other, the organisation with the more active application culture will receive more funding. Figure 7.7 shows the percentage of the estimated applicant pool that actually applied in the period 2011-12 to 2017-18:

[39] If data relating to organisations with smaller estimated populations (i.e. those which have probably been prompted to apply by a specific call) is excluded, the relationship becomes weaker, until with a lower threshold of about 20 potential PIs it disappears altogether.

[40] Although the regressions as plotted have positive slopes they appear to be driven entirely by the outliers at the bottom left extremity of each. (In the case of the golden triangle the outlier is Imperial College, while for the broader Russell Group it is QMUL.) An alternative interpretation of the figures is that there is in fact no relationship between potential PI population and per capita allocation, but that the golden triangle and broader Russell Group have higher per capita allocations which couple with their larger populations to produce a high degree of funding concentration. The question then becomes why those groups receive higher per-capita allocations.

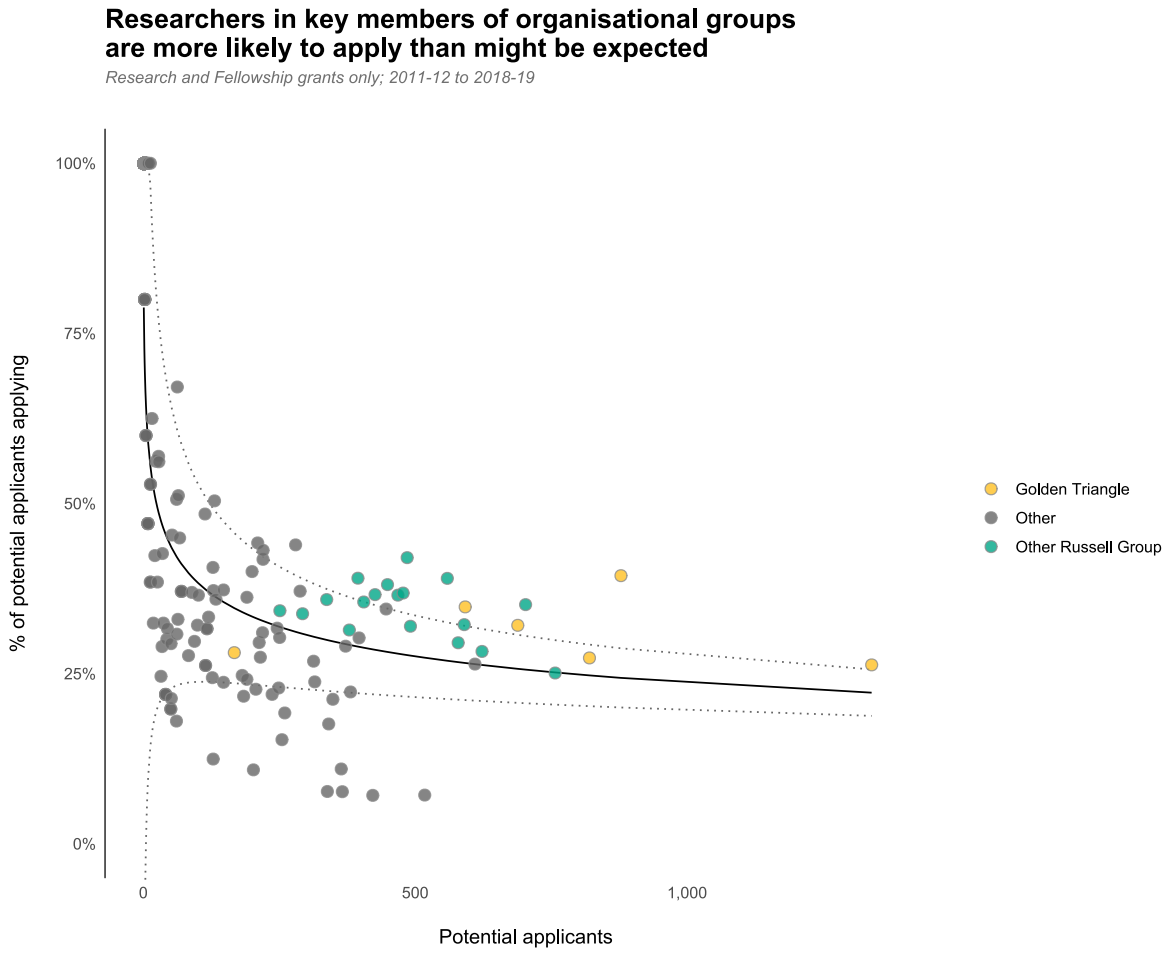

Figure 7.7: funnel plot of organisational application rates over the period 2011-12 to 2018-19 by organisational grouping, based on populations estimated using the Chao estimator. Funnel shows a 99+% control limit.

As Figure 7.7 incorporates uncertainty in the organisational population estimate it should be treated as indicative only[41]. The plotted line shows the expected application rate were there to be no effect of organisation on the rate. The same general conclusion as that in Figure 7.6 suggests itself: membership of the golden triangle or the broader Russell Group is associated with behaviour (in this case propensity to apply) that would not be expected if size on its own was the determinant. To some extent group members receive more funding because researchers in members of those groups are more likely to apply. Again, why this is the case and which way causality runs is an open question.

Summary

The facts relating to the outcomes experienced by the Russell Group or organisations which make up the golden triangle are quite simple: they tend to have more researchers who are potential ESRC applicants; those applicants are more likely to apply; when they apply they

[41] Some estimates seem a bit unlikely, for example Bangor University's population of 518 is very large. And it appears that 'Other' organisations behave entirely differently to those which are members of one of the two groups: there may in fact be two funnels.

ask for more; when their proposals are fundable they are more likely to be funded; and overall a proposal from one of them is more likely to be funded full stop.

In the absence of a reasonable comparator group, similar in every way except for group membership, it is impossible to tell whether these outcome differences are a result of membership, unreasonably associated with it, or the effect of some underlying (and possibly quite reasonable) difference in funding merit.

There is enough variability in outcomes to be confident in saying that proposals are not funded solely because of where they come from. It is possible though that when decisions as to which proposals exactly to fund are marginal (as they usually are) effects which might be described as ‘benefit of the doubt’ are felt and manifested. This may or may not be reasonable. It is also possible that excellence is not distributed evenly and that the results we see simply reflect the actual distribution of good ideas and projects. Experiments, natural or controlled, would be needed to establish the truth.

Annexes

A – characteristics of 50 leading organisations, 2011-12 to 2018-19 financial years

Organisation	Total applied for	Number of applications	Total awarded	Unique PIs	% responsive	% funded	% fundable	Estimated applicants	Group (as of 2019)	% population applying	*d* index 2016-17 to 2018-19
University College London	366100000	475	174400000	346	50	36	15	878	Golden Triangle	39	8
University of Essex	222300000	189	154900000	122	45	30	16	275	Other	44	6
London School of Economics & Pol Sci	162900000	290	74000000	221	40	37	12	689	Golden Triangle	32	6
University of Oxford	235000000	455	71500000	352	46	30	14	1339	Golden Triangle	26	8
University of Edinburgh	116100000	300	51500000	218	38	41	12	559	Other Russell Group	39	6
The University of Manchester	130000000	328	48800000	247	40	30	13	703	Other Russell Group	35	6
University of Glasgow	104800000	239	43200000	171	49	26	15	468	Other Russell Group	37	6
University of Southampton	100700000	216	40100000	154	53	28	15	395	Other Russell Group	39	5
Cardiff University	114300000	241	38100000	171	44	33	14	449	Other Russell Group	38	6
King's College London	103900000	265	36400000	206	45	31	14	592	Golden Triangle	35	6
University of Cambridge	139700000	286	28000000	224	45	31	18	820	Golden Triangle	27	7
University of Warwick	116600000	242	27400000	176	46	25	20	596	Other Russell Group	30	6
Lancaster University	71800000	198	26200000	161	46	27	12	610	Other	26	6
University of Exeter	98400000	224	26200000	171	49	35	8	579	Other Russell Group	30	5
Newcastle University	71500000	160	25500000	119	38	26	12	379	Other Russell Group	31	5
University of Sheffield	98900000	236	24900000	190	42	35	11	757	Other Russell Group	25	7
Queen's University of Belfast	67100000	191	24300000	144	43	28	9	405	Other Russell Group	36	5
University of Leeds	117800000	282	24000000	204	47	24	13	485	Other Russell Group	42	6
World Conservation Monitoring Cen WCMC	19300000	2	22500000	1	0	100	0	1	Other	100	1
Institute for Fiscal Studies	29800000	53	21300000	33	40	53	8	65	Other	51	2
University of Bristol	93600000	247	20800000	176	49	31	20	478	Other Russell Group	37	5
University of Sussex	99000000	216	20600000	154	44	26	13	447	Other	34	6
Coventry University	38500000	66	20300000	52	36	18	11	237	Other	22	4

Organisation	Total applied for	Number of applications	Total awarded	Unique PIs	% responsive	% funded	% fundable	Estimated applicants	Group (as of 2019)	% population applying	*d* index 2016-17 to 2018-19
University of York	102700000	201	18900000	157	54	30	17	491	Other Russell Group	32	5
University of Birmingham	95600000	255	18800000	190	44	30	15	590	Other Russell Group	32	6
Durham University	62400000	196	17300000	156	46	30	10	426	Other Russell Group	37	5
University of Surrey	42100000	114	14200000	78	50	25	10	246	Other	32	4
Swansea University	32100000	70	14100000	52	47	14	13	128	Other	41	2
London Sch of Hygiene and Trop Medicine	71500000	87	14000000	66	43	22	16	131	Other	50	4
Institute of Development Studies	45800000	74	12800000	42	22	26	16	63	Other	67	2
University of Nottingham	77200000	163	12400000	121	33	22	13	337	Other Russell Group	36	5
Liverpool School of Tropical Medicine	17900000	9	12200000	7	0	22	0	13	Other	53	2
School of Oriental & African Studies	54100000	70	12100000	60	31	27	10	341	Other	18	3
University of Aberdeen	29800000	89	12000000	69	56	29	12	191	Other	36	3
National Institute of Economic & Soc Res	13900000	27	11200000	16	26	59	11	29	Other	56	1
University of Liverpool	61600000	151	11100000	99	48	27	16	293	Other Russell Group	34	4
University of Leicester	56100000	132	10300000	108	46	24	14	372	Other	29	2
University of East Anglia	55900000	140	9800000	93	36	26	12	210	Other	44	4
Open University	61200000	154	9600000	120	43	21	16	397	Other	30	4
University of Kent	41300000	136	8700000	107	54	22	13	288	Other	37	4
International Institute for Env and Dev	15900000	10	6700000	9	0	30	10	41	Other	22	1
University of St Andrews	22100000	64	6300000	50	52	22	11	260	Other	19	2
Bangor University	22200000	48	6200000	37	52	21	6	518	Other	7	1
Loughborough University	36900000	103	6100000	85	39	21	11	381	Other	22	4
University of Reading	34200000	99	6100000	84	45	20	7	313	Other	27	3
University of Bath	44800000	129	5600000	95	55	19	19	220	Other	43	4
Imperial College London	33300000	60	5300000	47	48	23	15	167	Golden Triangle	28	3
Queen Mary University of London	40900000	112	5000000	86	55	16	24	251	Other Russell Group	34	4
University of Stirling	45400000	130	4900000	92	52	16	10	220	Other	42	3
National Centre for Social Research	12800000	26	4400000	15	27	62	8	35	Other	43	2

B – interpretation of funnel plots

Funnel plots of one kind or another have been used for a variety purposes for some time. Broadly speaking they are a means of assessing the reliability of a particular measurement in the context of uncertainty surrounding it. The funnel plots in this analysis are based on the approach originally set out by David Spiegelhalter[42] but they incorporate an adjustment for a relationship between success proportion and volume[43]:

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

Where *n* is the number of decisions made (or in a broader sense the number of trials) and the parameters *a* and *b* are derived from a nonlinear least squares regression model of the data. The parameter *b* determines the existence (or otherwise) of a success-volume relationship. If *b* = 1, *b*-1 = 0 and the success rate is constant (producing a traditional flat funnel plot.) If *b* > 1 there is a positive success-volume relationship; if *b* is < 1 the relationship is negative. In the text, a significant relationship is identified when the 95% confidence interval for *b* does not include 1. The parameter *a* allows an estimate of the success rate for a single trial (for example, for an organisation submitting only one proposal in a year) as that rate is simply 1-*e*^a.

The funnel plot limits are based on 2 and 3 standard errors, which correspond roughly to a 95% and 99+% control limit. The simple Wald interval is used to define the limits rather than a more complicated one because it imposes a further dose of caution on decisions about abnormality at low values of *n*.

[42] Spiegelhalter DJ. Funnel plots for comparing institutional performance. Statistics in Medicine 2005; 30;24(8):1185-1202; for a pdf see http://psmu.improvement.nhs.uk/psc-shared-library/measurement-evidence-base/16-funnel-plots-for-comparing-institutional-performance/file

[43] The derivation of this is at https://esrc.ukri.org/files/about-us/performance-information/the-relationship-between-decision-volume-and-success-rates/. It turns out that this is the Weibull cumulative distribution function