Quantitative and statistical approaches

Introduction

The role of quantitative research in coaching has traditionally been rather limited (Clutterbuck, 2013), perhaps because the primary aim of most researchers in the field has been to understand the interactions between indi-viduals as influenced by their perceptions and motivations. Quantitative approaches may also be considered inappropriate to describe the complexities of coaching relationships. However, there are many contexts in which researchers in coaching may wish or need to use a quantitative approach, and/ or gain numerical evidence (Fillery-Travis and Passmore, 2011). Typically, these have related to the measurement of impact, effect, or return on investment (ROI). Often, there is the further objective of generalisability, that is, deducing general principles that apply across a wide range of individuals and contexts. This chapter provides an overview of the main quantitative approaches to research in coaching, along with pointers to sources of more detailed information (see 'Further reading' at the end of the chapter). It also

gives an insight into the real issues that researchers encounter when undertak-ing quantitative research in coaching, and suggests ways of moving forward. Comment boxes are included to highlight specific issues that researchers in coaching may encounter when undertaking quantitative research.

Distinctive features of quantitative research

When we as research supervisors ask practitioners and students why they want to do quantitative research the reply is often less to do with methodo-logical consideration and more to do with the perceptions of validity and 'correctness' or 'truth'. Quantitative research can be identified as providing higher levels of evidence than its qualitative alternative although this is a contentious point (Corrie, 2010). It is indeed well suited to the generation of universal or generalisable findings that allow prediction of events and out-comes, and particularly useful for the demonstration of effectiveness and optimisation of practice. However, quantitative research is not foolproof, and can provide misleading results when based on an incorrect premise or where there is an error in interpretation or analysis. It is also not applica-ble to every research question. For example it is almost impossible to use if the phenomena of interest to us are the dilemmas, values, choices and relationships within practice unless it is present within a mixed-method design. It can also fail to take into account the nature of practice where innovations and improvisations are common and the outputs of quantita-tive methods, for example findings of theoretical significance, may be less informative than the pursuit of knowledge that is richly described and of direct practical application.

When considering whether to use quantitative methods the researcher must start at the beginning of the design process and consider how well it answers the question posed. Does the use of quantitative methods give the research integrity and coherence? Is it an effective strategy for this particular research question? Will the interpretation of data provide the basis for the output of strong evidence?

Box 7.1 Key concept – Strong evidence

Grant's (2017: 65) definition of evidence does not privilege one type of research approach over another but takes the more sophisticated view that different research designs have utility for addressing different research questions:

> Strong evidence can be understood as information and evidence from well-designed and peer-reviewed studies where the methodology is eminently suitable for the research questions being addressed and the results have been replicated in a range of populations where appropriate.

This is the heart of the quandary that researchers face when considering whether to use quantitative methods or not: What is their research question and what approach would be the most appropriate to address it? To explore this in more detail we need to look specifically at how researchers develop their research questions. This is by no means a trivial or straightforward task as the question must be specific, concise, and well defined, and by construct-ing it well the research question will lead directly to the most appropriate research approach. As suggested in Chapter 3 of this volume, it is the hub and anchor of all the activity within the research. It informs what methodology is appropriate and what data should be collected. The researcher will actively return to the research question during the research to check whether they are on the right track and therefore if it is ill-defined it can lead the research off track and lead to a confused outcome.

For practitioners particularly, the research question is not one developed through a disinterested curiosity but one that must lie at the heart of their practice. The researcher and the practitioner are two modes of working which cannot be completely separated. The beliefs, values, experience and knowl-edge about practice will influence how the practitioner views events, their role within them as well as the questions they have chosen to ask about their prac-tice. As researching practitioners are effectively putting their theoretical underpinning and practice forward for scrutiny this makes the process of research, in whatever paradigm or methodology they choose, a deeply per-sonal experience and one where reflexivity becomes an important considera-tion as illustrated in one practitioner's experience in Box 7.2.

Box 7.2 Doing research – Doctoral student 1

I worked in IT for years, and unlike many of my coaching colleagues, I like numbers, and so I always thought I would do quantitative research. But I soon released that the questions I was asking couldn't be answered that way. I've had to try very hard to learn how to use qualitative approaches like phenomenology.

The account in Box 7.2 suggests that it is important to start the research process with a clear question, and then to work out the correct methodolog-ical approach and design. At first sight this may seem at odds with the pos-itivistic tradition which has dominated quantitative research for so long as researchers strove for detached indifference to their research. An academic or professional researcher will be interested in the development of knowl-edge per se and will look to validate their work through analysis and peer review. In contrast a practitioner researcher will be interested in research as an improvement strategy for their practice where validity is sought through acceptance by clients and peers as well as practical applicability.

This difference in approach requires the practitioner to reflect upon their practice and positions their research within it.

Researchers who are contemplating a quantitative methodology may lean towards an objectivist ontology, and positivist epistemology. This would almost certainly be the case if they have a strong belief in the value of num-bers, and approach their research armed with a hypothesis, and needful of empirical data. However, many researchers within the field of coaching could be said to hold a pragmatic paradigm. This indicates that quantitative method-ology provides an avenue to answer some or part of the research question(s) posed, but is not necessarily at the heart of the research, and qualitative meth-odology may also play an important role (Guba and Lincoln, 1994). This would also be true of the critical realist paradigm, and practitioner researchers may often find that their paradigm has a 'critical' element in the sense that they approach research hoping to make changes (Fillery-Travis and Lane, 2008). In the context of coaching, this implies that the research undertaken will improve practice in some way, and lead to better client outcomes. (See also Chapters 2 and 10 for the discussion of pragmatic approaches.)

Quantitative research designs

As mentioned above, the nature of the research design, and the methodologi-cal approaches employed, should always reflect the research question posed. Quantitative approaches are good at describing and making inferences about data, but they cannot explain. So, a researcher who wants to understand a situation needs to take a qualitative path. It is possible, however, for the research design to incorporate both quantitative and qualitative elements, as indicated in the list of designs below.

Box 7.3 Key concepts – Exploratory and confirmatory designs

Quantitative methodological approaches tend to be either exploratory or con-firmatory (Atieno, 2009). Exploratory designs involve discovering and describing, and could address research questions such as 'What is coaching like in this organisation?' or 'What is the ROI resulting from a particular coaching practice in this organisation?' The quantitative approach to answering such questions tends to involve descriptive statistics, and data analysis might produce outputs such as bar charts (e.g. % different coaching practices), and means (e.g. ROI). It may also be appropriate to develop models based on the data obtained (Palmer, 2008).

Confirmatory designs, by contrast, are deductive in nature, and generally begin with a research question which is tied to a hypothesis. For example, the

(Continued)

question 'Will coaching change the job satisfaction of a sample of individuals?' is likely to be linked to the researcher's underlying belief that it will indeed change job satisfaction. This is often a reflection on ontology/epistemology, and is usually specific to a positivist paradigm. Put in more precise research terminology, there would be a null hypothesis (H_0) stating that coaching would make no difference to job satisfaction, and an alternate hypothesis (H_1) stating that it would make a significant difference. The methodological design in this instance is likely to involve a measurement of job satisfaction of the same individuals before and after coaching.

Despite the paucity of quantitative research in coaching, four main quanti-tative research designs have been utilised to date:

1. ***Mixed method***: a design with both quantitative and qualitative ele-ments which is often exploratory. Typically, the design consists of a survey distributed across an organisation followed by a small num-ber of interviews. However, in Gray et al. (2011), a series of case studies interviews was followed by a survey.
2. ***Case study***: this is dealt with in detail in Chapter 8, but it is worth noting here that some of the methods used in a case study design may have quantitative elements; for example, Andreanoff (2016) used a case study design, focusing on one organisation, in which there were quasi-experimental elements.
3. ***Quasi-experiment***: in this confirmatory design, there are at least two groups which are compared in some way. Often, one or more groups experience an intervention such as a particular coaching practice (these are the 'experimental' group(s)), whilst one group does not (this is the 'control' group). This is not a true experimental design because not all variables can be controlled, and the researcher has to be on the lookout for 'confounding' variables which may have as much, or more, effect than the intervention. A good illustration of this is the 'Hawthorne effect' (Wickström and Bendix, 2000). Variations on this design include comparison of groups with different coaching approaches, including uncoached (for example, Losch et al., 2016, looking at the effect of coaching on procrastination behaviours), or performance before and after coaching (for example, Hunt, 2010, who examined self-efficacy in female entrepreneurs).

4. ***Action research***: this type of design is discussed in detail in Chapter 9, but note that quantitative elements may be incorporated into action research designs to obtain numerical evidence of the efficacy of one or more implementation steps. For example, Olivero et al. (1997) used numerical indices of managerial competency within an action research design, and Leonard-Cross (2010) employed a quasi-experimental method within an action research cycle.

Quantitative methods

Quantitative research often involves direct measurement, and in the world of coaching this could involve measures such as sales over a particular time period, or the number of skills gained after a training course. There may be a requirement for a measure of more complex concepts related to impact or ROI, and a variety of scales exist within the world of coaching to quantify con-cepts such as empowerment and accountability (Hagen and Peterson, 2014).

One of the most widely used tools involved in the implementation of quan-titative research methods in coaching is the questionnaire (a survey involves the distribution of a large number of these). Questionnaires are frequently self-completion in nature, administered via hard copy or online. Response rates may be disappointingly low, especially if surveys are over-long or visually unappealing. Alternatively, questionnaires may be completed by an interviewer within a highly structured interview. This process can be time-consuming, but has the merit of ensuring that selected participants respond.

The questions in questionnaires need to be devised with care, bearing in mind that if the wrong question is asked, an opportunity to get the required information is missed. Questions may need to be framed in an indirect way, avoiding topics that participants might find uncomfortable or difficult to understand, and ensuring that wording is clear and unambiguous.

Questionnaire questions often aim to provide numerical values for concepts such as 'satisfaction' or 'confidence', and typically Likert scales will be used for this. For example, leadership potential could be gauged from 1 to 4 on a Likert scale. Clearly, such scales do not give absolute measures, just an index of them, and quantifying complex ideas in this empirical way may not always be appropriate.

Box 7.4 Key concept – Empiricism

Empiricism is an epistemology in which facts are thought to 'speak for themselves'. Its difficulty is that empiricists assume the world consists of facts; this is unrealistic. The real world is very complex (Olsen, 2004: 14).

The nature of quantitative data

From a positivist perspective, the best sorts of research data are 'scale' or 'interval/ratio' data. These are easily measurable numerical variables with absolute, known intervals (for example, the difference between 3 cm and 4 cm is 1 cm, which is exactly the same as the difference between 6 cm and 7 cm). Variables like this include length of service, sales income, profit, number of clients. However, the questionnaire questions based on Likert scales generate 'ordinal' or 'rank' data. These data reflect order, or relative magnitude, and do not give precise values (so, for example, we cannot be sure that the difference

between ‘not content’ and ‘quite content’ is the same as between ‘content’ and ‘very content’). A last type of data that may be involved in quantitative research is ‘categorical’, and includes categories such gender and job type. On an individual basis, such data are not quantitative at all, but summarising such data produces numbers (or ‘frequencies’) per category which can be analysed using quantitative approaches.

Quantitative data analysis

The approach to analysis depends on the research question and design, and also on the type of data generated. If research designs are exploratory, data analysis may focus on descriptive statistics, for example, mean, mode, range, standard deviation. Pictorial representations such as pie charts or bar charts may also be helpful to summarise data and show trends. However, this sort of analysis does not help deduce relationships or differences, and certainly does not indicate how significant they are. For this, inferential statistical tests are needed, and these will be particularly appropriate where research designs are confirmatory, and the objective is to test one or more hypotheses related to the research question. A range of common statistical tests is outlined in Table 7.1. These tests are ‘parametric’ in that they are based on an assumed set of pop-ulation parameters, one of which is that the data follow a ‘normal distribution’ (i.e. a distribution that looks like a bell-shaped curve with most values in the centre, and fewer at the edges). These tests work best with scale data. Ordinal data, as produced from questionnaires, are best analysed using ‘non-parametric’ variants of these tests, as listed in Table 7.2.

Table 7.1 A list of common parametric statistical tests (for scale data)

Type of research question	Statistical test/method	Example
Is there an association between two categorical variables? (e.g., gender, job type)	Chi squared test	Cunningham and Sagas (2003) used Chi-squared tests to compare duration of tenure of male and female sports coaches in the US
Is there a relationship between two scale variables? (e.g., between age and sales revenue)	Correlation and/or regression	Hwang et al. (2013) used correlation coefficients to assess relationships between paired combinations of emotional intelligence, coaching efficacy, and leadership style
Is there a difference (in a scale variable) between two separate groups? (e.g., number of customers logged by managers in two different regions)	Independent T test	Griffin et al. (2008) used independent T tests to compare the outcomes of coached and non-coached groups

Type of research question	Statistical test/method	Example
Is there a difference (in a scale variable) in the same individuals at two different times/conditions? (e.g., leadership rating before and after coaching)	Paired T test	Newnham-Kanas et al. (2008) compared variables including self-esteem pre- and post- one-on-one coaching using paired T tests
Is there a difference (in a scale variable) between more than two groups? (e.g., ROI from groups coached in 4 different ways)	ANOVA ('Analysis of Variance')	Fried and Irwin (2016) compared stress levels in university students in three different categories, pre-, mid-, and post-coaching intervention, using a one-way ANOVA
Is there an underlying variable influencing other measured variables? (e.g., one or two key variables influencing a range of variables assessed in a survey)	Factor analysis	Newsom and Dent (2011) used exploratory factor analysis to determine the main factors assessed by the Executive Coaching Work Behaviour Survey

The statistical test(s) used depend on the research question and the nature of the data generated, and should be selected at the stage of research design. It should be noted that the Chi squared test will handle data that are not nor-mally distributed, and also deals with 'categorical' variables.

Table 7.2 Parametric statistical tests and their non-parametric equivalents

Parametric statistical test	Non-parametric equivalent
Independent T test	Mann-Whitney test
Paired T test	Wilcoxon
Analysis of variance (ANOVA)	Kruskal-Wallis test
Pearson correlation	Spearman correlation

See also: https://research.phoenix.edu/blog/selecting-right-statistical-analysis-tool-your-research

The process of carrying out a statistical test involves the calculation of a par-ticular 'statistic' for the sample in question, and this is compared to a probability distribution – fortunately, both of these steps can be swiftly and painlessly car-ried out by software such as MS Excel or SPSS. A key output of a statistical test is a probability value, or 'p value': this ~~gives~~ the probability of the null hypoth-esis being correct. A small p value con~~tradict~~s the null hypothesis of a confirm-atory research design, and reinforces the alternate hypothesis, in that it indicates that there is a significant difference. In the example question posed earlier, 'Will

coaching change the job satisfaction of a sample of individuals?', if the p value obtained was less than 0.05 (or 5 per cent), this would indicate a significant effect of coaching on job satisfaction. It is important to note that statistical tests cannot prove cause and effect – they just imply that one variable has an effect on another.

Researchers may additionally need to report on values such as validity, reliability and effect size. Validity refers to whether a particular measure actu-ally assesses what it is meant to assess. For example, is a questionnaire ques-tion on 'length of lunch-break' an appropriate measure of job satisfaction? Validity is also important with respect to coaching scales (Hagen and Peterson, 2014). Reliability, by contrast, gives an indication of the repeatability of results. This is often discussed in the context of questionnaires: for example, would the same person answer the same question the same way on two dif-ferent occasions? Reliability can be gauged using a measure called 'Cronbach's Alpha', and this is discussed thoroughly by Field (2017). His book is also a good source for information on 'effect size', which may be required when looking at the relationship between specific scale variables. It can be thought of as another term for 'correlation', but there are other values that can indicate effect size, such as 'Cohen's d'.

Practical issues

It is not always easy to identify the exact knowledge you need (Knorr Cetina, 2001), or how to acquire it. It is not uncommon for doctoral researchers to take a considerable amount of time to find a research question that works for them, but until this is clear, it is not possible to decide on the best research design and methods, and the variables that need to be measured. This is certainly true of a complex topic such as coaching (Williams, 2014), and it is important not to oversimplify, and to be as specific as possible about the question that needs to be asked (Fillery-Travis and Lane, 2006) as suggested by one of our postgraduate students (Box 7.5).

Box 7.5 Doing research – Practical issue 1

I'd read loads of books, and I was really pleased that I was finding out more and more about my field. But in hindsight, my reading wasn't directed enough, and I was just muddying the waters in terms of working out what I was going to write in my research proposal. In the end, it boiled down to: 'What do I really need to know?' Once I'd sorted out my priorities, and realised what I needed for my coaching, I was a lot more focused about how to approach my literature review, and eventually my methods.

Another specific practical issue relates to the importance of selecting scales that are demonstrably valid and appropriate (Box 7.6).

Box 7.6 Doing research – Practical issue 2

I was working with a group of coaches, and we had set up a small research project focusing on leadership capabilities. We decided to interview ten managers in the company and chose Reed et al.'s (2011) scale to measure leadership potential. All went well until the seventh interview, which was with a senior accountant. He was adamant that the scale was inappropriate as a measure of leadership, and so we had to terminate his interview. This meant a loss of an individual from our sample, but the bigger issue was whether we should carry on using the Reed scale or find some other way to assess leadership.

Ethical issues

All research involving people should be planned carefully to ensure that any potential harm to participants, and researcher(s), is minimised, as discussed in Chapter 3. However, it is often thought that quantitative methods pose fewer ethical issues than more interpretive methods, and it is true that issues of anonymity and confidentiality are reduced by the use of tools such as online surveys. However, it is important to treat quantitative methods just like any others, and ensure, for example, that all participants consent to their partici-pation, even if online. It is always best to obtain written consent if possible, though in the case of surveys administered online it is common practice to explain to participants (for example, via an initial web page) that their contin-ued participation implies their consent.

Different coaching scenarios will have their own ethical issues in the con-text of quantitative research. Andreanoff (2016), for example, noted that in quasi-experimental designs, the selection of individuals for 'control' groups may pose issues, for such individuals may have been denied coaching inter-ventions, or may have been waiting for coaching for some time.

Quantitative research approaches may have many of the same ethical concerns as qualitative studies: in Box 7.7 a researcher deals with issues of anonymity.

Box 7.7 Doing research – Ethical issue

I had decided to get my scores on attainment, goal-seeking, and motivation via short interviews. I wasn't looking for any narrative, only scores, so I didn't need to use an audio recorder, and I was feeling quite chuffed about that because I knew that people sometimes object to being recorded. But what I hadn't bar-gained on was the room I was allocated: it was a small meeting room with glass partitions, and everyone in the open plan office could see who was in there. There was no privacy, and I knew no one would agree to going in there. We ended up using the office of one of the associate directors.

Sampling issues

A common aim of quantitative research is to find the characteristics of a pop-ulation, based on an investigated sample. So, for example, if it is found that a particular training course has raised the target-setting abilities of a sample of thirty employees, it might be assumed that the course would help all employ-ees. If there is a need to generalise in this way, it is best to use a large, ran-domly-selected sample. Small samples, and/or samples selected via 'convenience sampling' (such as selecting those who are easy to find) or 'snowball sampling' (asking a participant to pass on a request to their colleagues/friends, and so on) may result in biased outcomes. It is possible to use sample size estimators (for example, www.surveymonkey.com/mp/sample-size-calculator/) to find the minimum sample size for the research task based on variables including the required confidence interval and acceptable margin of error. However, it may not be so easy in practice to obtain the recommended number and, as sug-gested in Box 7.8, it is not always easy to obtain an adequate sample size.

Box 7.8 Doing research – Sampling issue

I was using skills levels as an estimate of ROI, and I thought it would be best to use an online survey because it wouldn't take too much of my time, and I could get a big sample. I sent the survey link out to all management personnel, but after a week I still only had five responses. So then I sent out a reminder, and I had a few more, but eventually I had to ask the CEO to email them all; she wasn't very happy about that. I got the data I needed in the end, but for a smaller firm, short interviews might have been a lot better.

Data storage issues

If personal data need to be retained, then data storage may have implications for confidentiality, and this ethical issue is discussed in Chapter 3. It may be helpful to upload quantitative data to data archives held within an organisation or in the public domain, so that other researchers may benefit from opportunities to carry out secondary analysis, so long as such storage conforms to the General Data Protection Regulation (GDPR) based on the Data Protection Act 2018. The storage of personal data in quantitative research may be an issue for some researchers as reported by the doctoral student in Box 7.9.

Box 7.9 Doing research – Data storage issue

If you're doing pre and post studies on the same people, you have to be able to identify them. This wasn't a problem for the coach I was working with as they were all her clients, but I had to make sure to explain on my application for ethical approval exactly why I needed personal contact details.

Analysis issues

Insufficient data is frequently an issue in quantitative research, especially as statistical tests have limited validity and reliability if sample sizes are small. It is therefore important to make contingency plans well in advance to help ensure that adequate data can be acquired, or alternative strategies employed.

Another common problem is not being clear about the correct statistical test to use. As mentioned above, this issue should be addressed at the stage of research design. Researchers can encounter difficulties when their initial plan-ning is weak. Fortunately, there was a work-around in the example in Box 7.10.

Box 7.10 Doing research – Analysis issue 1

We started up a coaching programme in an industry that was completely new to the concept, and one of the things we wanted to know about was the effect on coaches. At the end of the programme, we gave them some self-completion questionnaires with questions that focused on emotional intelligence. So that was fine for the end point, but we didn't have a baseline. Looking back, we should have planned for that, but it didn't occur to us at the time. As it happened, though, most of the coaches already had this information via feed-back we'd given them on a different project. So we used this in a 'post hoc' way – it wasn't ideal, but it allowed us to do the 'before' and after' comparison that was needed.

One doctoral supervisor explained why a statistical test which is appropri-ate for the data type as well as the research question should be selected (see Box 7.11).

Box 7.11 Doing research – Analysis issue 2

One doctoral candidate of mine wanted to compare the same coaching pro-gramme in two different public sector organisations. He wanted to know about people's satisfaction with the programme, and how they felt it had affected their ability to meet their goals. He utilised a quantitative design with a survey as the main method, and collected a lot of data, all of which consisted of responses on a six-point Likert scale. Given that this is ordinal data, I recom-mended that he use a Mann-Whitney test instead of a T-test.

The longer case study from the HR context (Box 7.12) illustrates a mixed methodological design. It also indicates the potential value of factor analysis when analysing questionnaire data.

Box 7.12 Case study – The HR professional also studying for a doctorate in professional practice

I'm a senior executive in the HR department of a major financial services firm, and also a candidate for a professional doctorate. Five years ago, my firm experienced a takeover, and shortly afterwards we had to downsize. The Directors wanted to make the change as positive as possible so just after the change we decided to invest in development activities, so that remaining staff knew that they were appreciated, and knew we wanted them to stick with us. From the HR point of view, we particularly wanted to prepare managers for potential leader-ship roles, and we decided to use coaching as our main development approach.

I also wanted to get data for my research that would help us assess the effectiveness of coaching and help us work out which aspects of the coaching interventions were most valuable. The firm was also keen to find the ROI, since a two-year coaching programme with sixty managers was a major expense, and so I needed to collect 'before and after' data for things like individual sales figures, and % retention of sales team members. This set me off on a quanti-tative pathway, with a survey right at the end to get coachees' assessments of the programme in terms of variables such as trust and satisfaction. I felt it was also important to understand their responses, so I added a qualitative element involving a small number of interviews. Overall, therefore, my research design was mixed methodology.

Analysing ROI data was quite straightforward, basically a case of subtracting before and after per individual, and then averaging. Fortunately, the survey went well also, and we had a great response rate. Again, it was fairly easy to find the modal response per question, but I went a step further and used factor analysis to find the questions that seemed to be significantly meaningful. Factor analy-sis was new to me: it is a statistical process based on correlation that basically weighs up which factors have the greatest effect, and so can point to underlying factors that could be responsible for results obtained. I used SPSS for this, and undertook a process called 'Principal Component Analysis' which confirms whether potential factors are as important as we might think. I expected that 'confidence' and 'encouragement' would figure prominently, and that proved to be the case. I had to bear in mind, though, that even with sixty managers, my sample size wasn't very big, and the margin of error would have been more acceptable if my sample size had been in excess of 100. The outcomes of thematic analysis of interview transcripts tied in well with what I found from factor analysis, though obviously there were a lot more details and nuance.

Evaluation

Quantitative approaches can be of great benefit to research in coaching. Quantitative methods such as surveys can be relatively simple to implement, particularly when administered online, and quantitative analysis, despite its

reputation, is usually far quicker and more straightforward than its qualitative counterpart. The numerical outcomes of quantitative analysis are relatively easy to judge and compare with others, and the large sample sizes involved tend to inspire confidence and can allow for generalisability. This means that clients and peers, particularly in the world of business, may be disproportion-ately impressed by quantitative research outputs.

However, whilst quantitative research can be very good at describing a sit-uation, it does not provide reasons, and for those we have to turn to the 'rich-ness' of qualitative research (Fillery-Travis and Cox, 2018). Also, generalised theories may oversimplify situations which are actually very complex (Olsen, 2004), and this is a particular hazard when it comes to the feelings and moti-vations of human beings. There are practical issues to bear in mind, also, such as the need for large sample sizes which may be difficult to achieve. There must also be careful consideration of possible confounding variables which can completely invalidate data which may have taken a great deal of time and effort to generate.

Overall, quantitative approaches have much to offer research in coaching, but they should be used thoughtfully, and in response to appropriate research questions. In many cases, a mixed methodological design will be the most useful approach and will help to ensure that the clarity of numbers is balanced the explanatory power of words.