

Chapter X

How do we know anything? Philosophical issues in the collection and interpretation of operational research data

Ben Hardy and Philip Stiles

X.1 Introduction

The economist John Kay had some sage advice for statisticians: whenever you see the product of serious analysis or calculation you should always ask yourself “what is the question to which this number is the answer?” (Kay 2011). This seems a fairly straightforward piece of common-sense guidance, but it is a much more sophisticated and disquieting observation. It hints at the complex processes that produce simple numbers and numbers’ tendency to give the illusion of a concrete answer.

The shift from operations research (OR) to *behavioural* operations research (BOR) (see Franco and Hämäläinen 2016) means that ever more evanescent concepts, such as product sales, employee satisfaction scores, and average call handling times, are being quantified. Many of these things are not discrete or concrete but are ephemeral and transient. Baked bean cans on a production line are discrete and concrete: perceptions of leadership are ephemeral and transient.

The key to improving measurement is twofold. One way is to improve our methods so that our counting is more accurate. The other is to examine the philosophical assumptions that underpin our research. BOR, as a field, has done this (see Becker 2016; White 2016; White et al. 2016), albeit tangentially by discussing research paradigms. This chapter aims to tackle the problem head on, by looking at three different factors which affect measurement and the conclusions we draw from any measurement. These are the nature of reality, what (and how) we can know about it, and how we generate knowledge.

To the researcher steeped in the natural sciences may view these questions in much the same way as the physicist Richard Feynman, who is reported to have said “philosophy is about as useful to scientists as ornithology is to birds; philosophy of science, as ornithology, is a civilizing hobby but of no real use to its objects” (Pernu 2008, p30). But the shift from OR to *BOR* also necessitates a shift from a natural science paradigm to a social science one, with the attendant need to consider underlying philosophical matters.

So, what if we don't? What if we ignore philosophy altogether? There are three potential problems that can arise from failing to pay heed to philosophical matters. The first is that we only capture a partial view of the area of study. The aphorism that ‘what is measured is what gets done’ may have a corollary that what is not measured does not get done. Some phenomena, such as trust in a supplier, are not readily captured and so may not be incorporated into models, despite being an important factor in determining transaction costs. The second problem is one of comparability between pieces of research. A lack of clarity about the philosophical underpinnings of a study – for example, how leadership was measured and how that measure pertains to a particular culture – may mean that attempts to compare studies, or synthesise them into a meta-analysis, leaves us comparing apples with oranges. The third problem is a fundamental one as we may not actually be measuring the thing that we think we are measuring (see Hardy and Ford 2014)

These three potential problems form an essay that we will use to assess the implication of philosophical concepts for *BOR*. In doing this we make three contributions. The first is to illustrate that the shift from OR to *BOR* requires a greater focus on philosophical issues than has hitherto been the case. The second is to enable researchers to identify where their research and proclivities might fall on a number of philosophical spectra. The third is to help researchers think about alternative methodological approaches which might, in turn, offer new understandings of *BOR* phenomena.

The rest of this chapter is laid out in three main parts. We begin by looking at the nature of reality and the assumptions we make about it – the ‘there’ that’s out there. Then we turn to how we might capture that reality – what we can know about the ‘there’. Having obtained our data we then will consider what it might or might not tell us about the world – the knowledge it produces. Finally, we trace the implications of this for behavioural operations research and how those conducting such research might ensure their work is both sound and useful.

X.2 What is the nature of the ‘there’ out there?

The study of the ‘there’ out there is known as *ontology*. Operations researchers are likely to be familiar with the term as it is widely used in computer science to describe “a formal, explicit specification of a shared conceptualization” (Gruber, 1993, cited in Breitman et al. 2007). As Gruber (1995) puts it “The term is borrowed from philosophy, where an Ontology is a systematic account of Existence. For AI systems, what "exists" is that which can be represented.” (Gruber, 1995, p 908). This is a very thin ontological view which captures what can exist in an AI system. Philosophers take a broader view and the assumptions they make about the world and nature of reality can vary greatly.

At one end is a world where there is an external reality which can be measured and about which we can develop rules. This is the view taken by disciplines such as mechanical engineering and computer science. At the other end of the spectrum is the sceptical view that nothing can be ascertained with certainty. Between these two poles there are a number of other interpretations of how ‘real’ the world is.

X.2.1 Realist perspectives

The natural sciences tend towards the realist end of the spectrum (Easterby-Smith et al. 2002). At the hard end of realism is “the view that entities exist independently of being perceived, or independently of our theories about them.” (Phillips 1987, p 205). This view

believes that there is a single state of the world which exists independently of us and will go on existing whether we are there or not.

Realism, however, has its problems. First, the history of science furnishes an abundance of empirically successful theories that were ‘abandoned one after another’ in the light of new information (Chakravartty, 2017). Second, there is often a gap between, say, the laws of nature and the knowledge that scientists have of these laws. Third, there is a representation problem – how do we come to know about the mind-independent world? - a problem we shall address later when we discuss epistemology.

These difficulties have led to a softening of this position to internal realism (Putnam 1987). Here our understanding of the world depends on conceptual schema which combines both fact and convention. So, when we see a baked bean tin, our perception is dependent on the language we use to describe it but not, necessarily, on our matching it exactly with other baked bean tins we have seen – or the elementary particles which make it up. There is a flexible vagueness to our precision. As Putnam puts it ‘*What is factual and what is conventional is a matter of degree. We cannot say, ‘These and these elements of the world are the raw facts, the rest is the result of convention.’*’ (Putnam, 1975).

There is no unique schema which captures the natural, mind-independent, properties of the world. It makes more sense to speak of the world having properties within a conceptual scheme – as BOR researchers do when they talk of paradigms. Internal realism suggests that there is a single external reality, but that it’s very hard to get direct access to it – not least because it is bound up with convention. Instead we have to gather indirect evidence and use that to infer the nature of reality.

These realist positions are plausible for some natural science disciplines, particularly those in which elements of a system can be isolated effectively. The laws or principles identified in realist ontologies are regarded as immutable and fixed.

For social science the debate takes a rather different form. The idea of hard realism is fanciful. The very nature of a social concept means that it needs people to exist. As a consequence it cannot, by definition, exist independent of people in the same way that a rolled steel joist can. Trust, for example, is not an objective entity which exists when there is no-one there. For trust to exist there must be some form of agreement between people on what trust is and how much of it there is.

The hardest realist position that the social scientist can adopt is internal realism, where there might be an entity out there e.g. the idea of trust, but it cannot be readily grasped. This ontological position permits scientific research on social phenomena.

X.2.2 Relativist perspectives

In this case there is not a single objective state of the world ‘out there’ but rather there are *states of the world* which exist in relation to one another. When considering the effectiveness of a service level agreement (SLA), for example, there may be multiple states of the world. The SLA could have been great for the vendor but hopeless for the buyer. Or it might have caused long-term damage at the expense of short-term success. Or it may have worked for both buyer and supplier but not for end-users or other stakeholders. In short, there are many states of the world which are relevant to the issue.

This may feel hopeless to the researcher of a scientific bent, but it is not. If we can agree on particular definitions and constructions, we are in with a chance of being able to compare like with like. And this is what researchers often do. We may use definitions of class or socioeconomic status (for example, advertisers may use ABC1, C2DE social classes) to establish who is being served by a product. But class is an artificial entity produced by human endeavour – as are the classification systems to measure it. We might agree to use these definitions of class, but they do not capture the multiple ways in which people’s social class

or economic status might be constructed. They are conventions, in Putnam's (1975) terminology; but in this case it's all convention and no 'fact'.

Of course, sometimes there is no *truth* out there, nor can we agree on a convention. All we can do is describe what we have seen and try to give some hint of the bigger picture. The Duke of Wellington summarised it thus when asked to describe the Battle of Waterloo.

'The history of a battle is not unlike the history of a ball. Some individuals may recollect all the little events of which the great result is the battle won or lost, but no individual can recollect the order in which, or the exact moment at which, they occurred, which makes all the difference as to their value or importance.....' (Macaulay, 1849, p180).

The history of the Battle of Waterloo is, therefore, a construction, it is not a set of facts nor an account which can be true. Much of the data acquired in operational research when enquiring into the history of an event or situation fits into this category. This ontological position – nominalism – is that there is no 'truth' – there is no one state of the world, everything is constructed and that the facts that we consider are human creations.

When considering the nature of reality in BOR, therefore, we have a continuum between a hard realist position and a hard nominalist. On the one hand there is a real, independent world out there and on the other there is a world that is solely constructed by the actions of humans – and a number of points in between these two poles.

So why does this matter? The problem is that the assumptions we make about the nature of reality affect the kinds of data we collect, what we can know about a phenomenon of interest and the usefulness of its measurements. The behavioural part of behavioural operations research means that we have to think carefully about the nature of things like trust, leadership or reputation. Assuming that something like trust is a concrete entity that can be

measured like, well, concrete, may lead us to draw precise but fatuous conclusions from our research.

X.2.3 Critical realism

Critical realism represents an attempt to navigate between the *Scylla* of the realism embodied in scientific and positivist positions and the *Charybdis* of relativism. It is something of an outlier in the philosophical literature in that it offers a structured ontological view. The key author, Roy Bhaskar, proposes a stratified account of reality. From an ontological point of view, he makes a distinction between the objects of the natural and social world – and hence the natural and social sciences. Critical realism proposes that there is a real, intransitive world out there; a view synonymous with natural realism. This real world, however, is not solely grasped through direct empirical observation, its nature can also be inferred from observations of other phenomena which imply that there is an underlying causal mechanism (Mingers 2000).

Critical realism treats social phenomena slightly differently. It doesn't regard them as existing as concrete entities but rather as something that only exists in that they have an effect. So, leadership does not exist unless it exerts an effect on other people. If there are no people or effect then there is no leadership. Social phenomena are dependent on people producing and reproducing them. They also tend to be localised in place and time (Mingers 2000).

The stratified nature of critical realist ontology has important implications for operations research. Taking our leadership example, we might observe someone responding to a charismatic leader. This response is based on their behaviour and their behaviour is based on a series of electrical impulses in their brain. These impulses are depending on a series of chemical reactions, and these chemical reactions are dependent on molecules and atoms. This layered ontology has two important implications for research. Firstly, given that we are

unlikely to get to the underlying 'real' world, are we using the correct level of analysis? This is rather like the difference between an analysis of events and a root cause analysis. The second is one of comparability. Are they using the same level of analysis as others? Different ontological layers may mean different pieces of research cannot be readily compared.

X.2.4 Ontological implications for BOR

The ontological view taken by the researcher has implications for the three tests which we set out in the introduction. The first of these is whether we capture the totality of the problem. A realist might be quite confident that this is possible. A subjectivist might believe that it is impossible to capture the totality and so we can only have a conditional fragment. A critical realist might believe that it depends on the depth of analysis.

The second test was whether comparisons can be made across different pieces of research. Again, a realist might feel this is easy. A subjectivist would think that it is impossible. And a critical realist might believe it depends on the level of analysis.

The third test is whether we are measuring what we think we're measuring. Here the realist comes up short because of the problem of apprehending social phenomena. Concepts like trust or leadership are not concrete entities and so their presence has to be inferred and the process of inference introduces its own set of inaccuracies. The interpretivist is pretty confident that are grasping what they think they are grasping but largely because their claims are modest. The critical realist is less certain as it will depend on the ontological stratum being addressed. This process of grasping the nature of the world is what we now turn to.

OR researchers adopting a realist perspective embrace the view that the structure and reality of phenomena are beyond their actions. Objective physical and social reality must be "discovered" by crafting precise measures that will detect and gauge those dimensions of reality that interest the researcher. Understanding phenomena is thus primarily a problem of modelling and measurement, of constructing an appropriate set of constructs and an accurate

set of instruments to capture the essence of the phenomenon. Much of OR, implicitly or explicitly, adopts this view. This is bound up with adoption of “traditional” scientific method and assumptions in the early days of OR just after World War II (Ackoff 1979) and subsequently in the nature of assumptions about modelling in OR (Wierzbicki 2007).

Ontologically, non-realist perspectives emphasise the importance of subjective meanings and social processes in the construction and reconstruction of reality (Morgan 1983). This tradition does not presume that organizational structure or social relations are objectively known and unproblematic, but attempts to understand how and why individuals, through their socialization into, interaction with, and participation in, a social world, give it a certain status and meaning. A number of researchers within the BOR tradition have taken this approach, particularly to better embrace the social aspects of model building and developing a less reductionist version of reality than many realist perspectives would provide (White 2016).

A third perspective in OR has seen scholars reject the “paradigm wars” (Kotiadis and Mingers 2014; Mingers 2000, 2003) and adopt critical realism as a way of existing in the space between realism and non-realism (Mingers 2015; Sayer 2000). The critical realist view attempts to re-invite a realist view of being in the ontological domain while accepting the relativism of knowledge as socially and historically conditioned in the epistemic-logical domain (Mingers 2000, 2003, 2015; Sayer 1992). According to Mingers, critical realism fits well with the debates on OR as an applied discipline (Mingers 2000, 2003). In particular, it has been suggested that critical realism could be the basis for understanding the social situatedness of OR (Keys 1997; Sayer 2000).

X.3 What can we know about what's out there?

If ontology is our assumptions about the 'there' out there then, when conducting research, we need to think about how we can know about what is out there. This is the field of *epistemology*, the study of knowledge. As with ontology, there are differing opinions, with a continuum between the view that we have direct unmediated access to the world and a more contingent view that our ability to apprehend the world is partial and conditional. In this section we will unpack these different views and then in the subsequent section look at another element of epistemology – the question of how we know whether things are true or not.

X.3.1 Positivism

Operational researchers are most likely to be familiar with positivism and its allied positions. Positivism grew out of the values of the Enlightenment and aimed to challenge the power of Popes and Princes by arguing that individuals could apprehend and test the reality around them and didn't have to rely on those in power to tell them what was true and what was false.

Positivism aims to use standardised methods of inquiry which rely on information observable through the human senses. Hypotheses should be developed inductively and then tested with the aim of being able to predict phenomena as well as the necessary and sufficient criteria for their existence. Research should be as value-free as possible, so not contaminated by the researcher or their views of the world.

In order to achieve this, positivism makes a number of assumptions which include the observer being independent of what is observed, that problems can be understood if broken down into parts and that causal relationships and fundamental laws can be produced that explain and predict human behaviour.

Experience teaches us that the assumptions underpinning positivism are not always tenable. Consultants appearing on a factory floor affects behaviour and performance – as we know from the Hawthorne studies. So the assumption of independence is not that sound. Similarly, some systems are so interconnected that they cannot be broken into parts without disturbing the system as a whole. For example, a star analyst at a bank may be good at their job but also rely on an interconnected web of relationships to perform (Groysberg et al. 2004) – and this may be almost impossible to identify and quantify .

X.3.2 Social constructionism

The opposing pole to strong positivism is social constructionism. This is the view that everything we know about the world is constructed by humans and not “the neutral discovery of an objective truth” (Castello and Botella 2007, p 263). When taking this view, researchers try to capture the process of construction as well as what is constructed. So *how* the phenomenon of interest was produced is evaluated, as well as the *what* of what was produced, and the *why* it came to be produced that way.

This leads to a very different set of research practices and, consequently, a very different sort of output to positivist approaches. In a social constructionist approach, the aim is to get a broad understanding of what is going on and what it means, rather than generate mechanistic causal relationships.

An example of where social constructionism might offer a different perspective would be in examining customer satisfaction. This is a subjective experience which has significant consequences for businesses. Parameters such as waiting time, churn rates and so forth may correlate with customer satisfaction but one person’s two-minute wait might have a very different impact to another. Even asking people directly how satisfied they are may not give a good measure. Satisfaction has several meanings, one is to fulfil an obligation (she has satisfied the examiners), another might be to be satiated (‘the meal left me completely

satisfied') or happiness ('I am happy with the service I received') (see Hardy & Ford 2014). Constructionist approaches would dig beneath the measurement and try to understand not just what satisfied people, but what that actually meant.

Social constructionism also allows researchers to participate in the research. More positivist approaches prescribe detachment and independence. Constructionist approaches, by contrast, acknowledge the researcher's presence and their influence (for good or ill) on the results. This is particularly important for those working in consulting who, by the nature of their work, cannot be detached and objective, as they are part of what they are studying.

Constructionist approaches are commonly associated with qualitative research, where researchers use verbal rather than numerical data to delineate phenomena. This has the advantage of allowing richer and more compelling descriptions (restaurant menus and sales brochures are seldom written in numbers) but it also means that the data produced are less easily analysed and manipulated. There is also a perception that qualitative data are less rigorous.

The whole continuum between positivism and social construction could, in some ways, be thought of as a continuum between the immutable laws that positivism seeks to uncover and the endless uniqueness that constructionists seek to describe.

X.3.3 Moderate positivist and constructionist positions

Between the extremes of strong positivism and constructionism are more moderate positions which sacrifice methodological purity in the interests of real-world practicality.

On the positivist side this may mean less of binary choice between truth and falsehood, a relaxing of assumptions about observer independence and an understanding of the contingent nature of the principles governing human behaviour. It may also involve embracing more qualitative approaches or examining single-cases which offer new insights

into phenomena. These weaker positivist approaches are less about identifying laws and more about identifying regularities which might illuminate different aspects of operations research.

At the interpretivist end this might mean relaxing assumptions about uniqueness and the importance of rich description and qualitative research, for example by embracing simple counting techniques which enable the reader '*to gain a sense of the flavour of the data as a whole.*' (Silverman 1993, p 163). The idea that everything is unique and valued may also be relaxed, acknowledging that some events and sources of data are more significant than others. This may help identify regularities which can then be used to inform practice.

In essence, these more moderate positions make different forms of research possible. Researchers are not tied solely to the principles of positivist scientific research, nor are they adrift in the endless variety of constructionist approaches. These more moderate epistemological positions also allow the mixing of different methods of research in order to provide converging lines of evidence which give the confidence that something robust and usable has been discovered.

X.3.4 Epistemological implications for BOR

As with ontology, we will examine the relationship between the different epistemological viewpoints and our three tests for research output. When looking at whether research captures all of the phenomena of interest, the positivist approach can fall short. Observer independence and attempting to isolate specific aspects of phenomena may mean, almost by definition, that only parts of the phenomenon are explored. Constructionist approaches are more likely to capture the totality of a phenomenon but, as the viewpoint suggests, this will only be one particular view and others may be just as relevant. The more intermediate positions allow these trade-offs to be more finely tuned to the objectives of the research.

When looking at comparability of research, the positivist approaches score well as they conform to similar standards. By the same token, constructionist approaches are, by definition, incomparable. That said, constructionists acknowledge this incomparability and see it as a virtue, rather than a deficiency. Again, the more moderate positions may allow comparability but may also produce such a variety of approaches that studies cannot be compared.

The question of whether the research actually measures what we think it measures is an important one. Positivist approaches may well not be measuring what they think they're measuring because they do not capture the totality of the phenomenon being measured, or are confounded by the dynamic and evolving nature of social phenomena. Constructionist approaches faithfully reproduce what they think they're reproducing because they are fully seized of the contingent nature of what can actually be grasped. More moderate positions walk a line between attempting to produce generalisable findings which can be transferred from context to context whilst not losing sight of the fact that what is being recorded is ephemeral and transient.

OR has its main roots in positivism; research conducted in the positivist tradition "has been the most widely used approach in BOR thus far," (Brocklesby, 2016, p 123). However, overreliance on this approach has brought criticism from within the OR discipline, as Linstone (1985, p 80) notes, 'reliance is placed on data and models, and combinations thereof, as the only legitimate modes of inquiry... in its most extreme form, modelling becomes an end rather than a means (the Pygmalion complex'). Quantitative analyses tend to drive out qualitative analyses.'

As a corrective to this dominance, a number of OR studies have sought to develop constructivist accounts. ¥ But more needs to be done. For example, Ormerod's (2014) analysis of OR case research published in the main OR journals shows that most pieces focus

on the technical aspects of the model used, less common ones report on why it was used, and, even less commonly, there are those that report on outcomes. Ormerod concludes that the lack of attention to social process, including the non-linear, iterative, and negotiated aspects of the model-use-in-practice requires a radical rethinking of how research in OR is conducted.

This discussion of ontology and epistemology might seem abstruse and barely relevant to the operations researcher, and there is a century of operations research to prove this point. But the shift to behavioural operations research means that ignoring ontology and epistemology is less of an option.

Having spent some time on ontology and epistemology we want to turn to a subsection of epistemology which gets relatively little academic interest but is fundamental to research. The problem of knowledge.

X.4 How do we know anything?

There's nothing new in wondering how we know things. Plato wrote about it and many other philosophers have subsequently. Knowledge, these philosophers observed, was justified, true belief (JTB as it is commonly abbreviated). We have a belief – that it's 14:00; that there are daffodils on the table; that the new R&D programme is working. It is true – it *is* 14:00; there *are* daffodils on the table; the new R&D programme *has* produced a number of prototypes. And there is a justification – the clock on the oven says 14:00; my wife commented on the lovely daffodils; we did change the R&D arrangements and appoint a new head of R&D. So far, so good, as JTB gives us knowledge. But is it enough?

The answer is, unfortunately, no. The problem is most commonly associated with the philosopher Edmund Gettier who suggested that JTB is not sufficient to produce knowledge

(Gettier 1963). One of the more accessible examples of what have become known as a Gettier problems, the following hypothetical case: Imagine that you are driving past a field and see what looks like a sheep. Based on that observation, you believe that there is a sheep in the field. Your belief turns out to be correct, as the field does contain a sheep, but it is actually out of sight on the other side of a hill. What you actually saw was a dog disguised as a sheep. Do you really have knowledge that there is a sheep in the field? You have a true belief that is justified by seeing what looks like a sheep in a place where sheep are normally found, but philosophers generally say that your belief does not constitute knowledge (Chisholm 1989). That is, you do not know that there is a sheep in the field if all you have seen is a dog in disguise, even if a sheep that you have not seen is really there. You've got the right result but the wrong mechanism. In essence, you're right but you've not got knowledge.

This is of fundamental importance for behavioural operations research and management research in general. People obsess with method – the dog identification business. But they are often in narrow disciplinary silos and so have little knowledge in the field surveying business. They find something that gives them JTB but it's not knowledge as the reason they are right lies elsewhere.

Picking up the example from earlier on, the new R&D programme is bearing fruit and we do have a new head of R&D. But it was an employee who has been with the firm for 30 years who happened to go to a conference and learn a new technique, coupled with some new equipment coming on the market and on-site childcare which all contributed incrementally to increased R&D output.

The problem here is that much management research focusses on complex methods using mathematical tools to establish whether the data support the conclusion (i.e. whether there is J for the TB). But do authors and reviewers have sufficiently broad knowledge to

question what was observed – the data and accuracy of their underlying assumptions – to establish that their findings are not just accidental, artifacts or luck?

X.5 A plea for philosophy to be considered

Considering philosophical problems in BOR may feel like a classic piece of ivory tower self-indulgence. The busy researcher under pressure to produce results may feel that there isn't time for this sort of thing. In the words of one manager "we don't need theories, we just need things that work". But this is at the heart of the problem. How do we know what things are and whether they will work?

When operations researchers turned to behavioural matters, it seemed a straightforward move. If problems of standardisation and decision rules can be adduced when making things, surely the same sorts of things can be done with people? But crossing the Rubicon into the social realm opens up a set of problems which complicate the research process. The nature of the social world is rather different and more evanescent. This means that existing approaches might work, but they might also lead us badly astray. We may not be addressing what we think we are addressing.

The question of what to do about this is one to which thick textbooks are devoted. But these textbooks may not be read by the non-social scientist – not least because they can seem quite impenetrable. But it is important to grasp that there are philosophical foundations to any research endeavour and that failing to attend to them – particularly in the behavioural domain, can lead to research that either does not properly capture the totality of the phenomenon of interest, is not comparable with other work or simply does not measure what the researcher believes it to measure. As the sociologist R H Tawney commented when talking about sociology – although he could equally have been talking about BOR - *"Sociology, like history, is department of knowledge which requires that facts should be*

counted and weighed, but which, if it omits to make allowance for the imponderables, is unlikely to weigh or even count them right” (Tawney, 1971, p 147).

This chapter is not intended to serve as a primer on philosophy, rather it is an attempt to illustrate to the behavioural operations researcher that there is a layer of complexity that can undermine their research. By thinking about what there is and how we can know about it, more effective research can be undertaken and conclusions be weighted with an appropriate degree of confidence. People don't behave like things and, whilst there are some principles which are relatively stable, there many are simply transient regularities. An understanding of the philosophical principles which undergird BOR helps ensure that any regularities observed are appropriately qualified and their transience acknowledged.

References

- Ackoff, R. L. (1979). The future of operational research is past. *Journal of the Operational Research Society*. <https://doi.org/10.1057/jors.1979.22>
- Becker, K. H. (2016). An outlook on behavioural OR – Three tasks, three pitfalls, one definition. *European Journal of Operational Research*, 249(3), 806–815.
<https://doi.org/https://doi.org/10.1016/j.ejor.2015.09.055>
- Breitman, K. K., Antonio, C., & Truszkowski, W. (2007). Ontology in Computer Science. In *Semantic Web: Concepts, Technologies and Applications* (pp. 17–34). London: Springer London. https://doi.org/10.1007/978-1-84628-710-7_2
- Brocklesby, J. (2016). The what, the why and the how of behavioural operational research—An invitation to potential sceptics. *European Journal of Operational Research*, 249(3), 796–805. <https://doi.org/https://doi.org/10.1016/j.ejor.2015.09.034>

- Castello, M., & Botella, L. (2007). Constructivism and Educational Psychology. In R. A. Kincheloe, Joe L., Horn (Ed.), *The Praeger Handbook of Education and Psychology, Volumen 1* (pp. 263–271). Praeger.
- Chakravartty, A. (2017). Scientific Realism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 201). Metaphysics Research Lab, Stanford University.
- Chisholm, R. M. (1989). *Theory of knowledge*. Prentice Hall. Retrieved from <https://books.google.co.uk/books?id=AOJOAQAAIAAJ>
- Easterby-Smith, M., Thorpe, R., & Lowe, A. (2002). *Management research: An Introduction* (2nd ed.). London: Sage.
- Franco, L. A., & Hämäläinen, R. P. (2016). Behavioural operational research: Returning to the roots of the OR profession. *European Journal of Operational Research*, 249(3), 791–795. <https://doi.org/https://doi.org/10.1016/j.ejor.2015.10.034>
- Gettier, E. L. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121–123. <https://doi.org/10.1093/analys/23.6.121>
- Groysberg, B., Nanda, A., & Nohria, N. (2004). The risky business of hiring stars. *Harvard Business Review*, 82(5), 92—100, 151. Retrieved from <http://europepmc.org/abstract/MED/15146739>
- Gruber, T. R. (1995). Toward principles for the design of ontologies used for knowledge sharing? *International Journal of Human-Computer Studies*, 43(5), 907–928. <https://doi.org/https://doi.org/10.1006/ijhc.1995.1081>
- Hardy, B., & Ford, L. R. (2014). It's Not Me, It's You: Miscomprehension in Surveys.

Organizational Research Methods, 17(2), 138–162.

<https://doi.org/10.1177/1094428113520185>

Kay, J. (2011). Sex, lies and pitfalls of overblown statistics. *Financial Times*. Retrieved from <https://www.ft.com/content/dfe55458-ccb6-11e0-b923-00144feabdc0>

Keys, P. (1997). Approaches to understanding the process of OR: Review, critique and extension. *Omega*, 25(1), 1–13. [https://doi.org/https://doi.org/10.1016/S0305-0483\(96\)00049-7](https://doi.org/https://doi.org/10.1016/S0305-0483(96)00049-7)

Kotiadis, K., & Mingers, J. (2014). Combining problem structuring methods with simulation: The philosophical and practical challenges. In *Discrete-Event Simulation and System Dynamics for Management Decision Making*. <https://doi.org/10.1002/9781118762745.ch04>

Linstone, H. A. (1985). Multiple Perspectives: Overcoming the Weaknesses of MS/OR. *INFORMS Journal on Applied Analytics*, 15(4), 77–85. <https://doi.org/10.1287/inte.15.4.77>

Macaulay, T. B. (1849). *The History of England: From the Accession of James the Second*. B. Tauchnitz. Retrieved from <https://books.google.co.uk/books?id=uFoPAAAAQAAJ>

Mingers, J. (2000). The Contribution of Critical Realism as an Underpinning Philosophy for OR/MS and Systems. *The Journal of the Operational Research Society*, 51(11), 1256–1270. Retrieved from <http://www.jstor.org/stable/254211>

Mingers, J. (2003). A classification of the philosophical assumptions of management science methods. *Journal of the Operational Research Society*, 54(6), 559–570. <https://doi.org/10.1057/palgrave.jors.2601436>

- Mingers, John. (2015). Helping business schools engage with real problems: The contribution of critical realism and systems thinking. *European Journal of Operational Research*.
<https://doi.org/10.1016/j.ejor.2014.10.058>
- Morgan, G. (1983). *Beyond method : strategies for social research*. Beverly Hills, Calif.: Sage Publications.
- Ormerod, R. J. (2014). The mangle of OR practice: Towards more informative case studies of “technical” projects. *Journal of the Operational Research Society*.
<https://doi.org/10.1057/jors.2013.78>
- Pernu, T. K. (2008). Philosophy and the Front Line of Science. *The Quarterly Review of Biology*, 83(1), 29–36. <https://doi.org/doi:10.1086/529560>
- Phillips, D. C. (1987). *Philosophy, science, and social inquiry: contemporary methodological controversies in social science and related applied fields of research*. Pergamon Press.
Retrieved from <https://books.google.co.uk/books?id=O-iAAAAAIAAJ>
- Putnam, H. (1975). *Mathematics, Matter and Method: Volume 1, Philosophical Papers*. Cambridge University Press. Retrieved from
<https://books.google.co.uk/books?id=HjI9AAAAIAAJ>
- Putnam, Hilary. (1987). *The Many Faces of Realism*. Open Court.
- Sayer, A. (1992). *Method in social science : a realistic approach* (2nd ed.). London: Routledge.
- Sayer, A. (2000). *Realism and Social Science*. London.
<https://doi.org/10.4135/9781446218730>

Silverman, D. (1993). *Interpreting qualitative data : methods for analysing talk, text and interaction*. London: Sage Publications.

Tawney, R. H. (1971). *The attack, and other papers*. Freeport, N.Y.: Books for Libraries Press.

White, L. (2016). Behavioural operational research: Towards a framework for understanding behaviour in OR interventions. *European Journal of Operational Research*, 249(3), 827–841. <https://doi.org/https://doi.org/10.1016/j.ejor.2015.07.032>

White, L., Burger, K., & Yearworth, M. (2016). Understanding behaviour in problem structuring methods interventions with activity theory. *European Journal of Operational Research*, 249(3), 983–1004. <https://doi.org/https://doi.org/10.1016/j.ejor.2015.07.044>

Wierzbicki, A. P. (2007). Modelling as a way of organising knowledge. *European Journal of Operational Research*, 176(1), 610–635. <https://doi.org/https://doi.org/10.1016/j.ejor.2005.08.018>