


Revisiting Realist Theory of Science: A practical guide



By Tom Fryer & Cristián Navarrete

"This is a wonderful route map for beginners trying to make sense of Roy Bhaskar's *A Realist Theory of Science* – the most important foundational work on critical realism. Bhaskar's book can be a tough read for social science researchers and it's not always clear why they should stick at it. But this short, entertaining (yes, really!) book makes the core ideas accessible and shows how they are relevant to social science. Once you've read this, Bhaskar's own work will be much easier to get to grips with."

Dr Dave Elder-Vass, Loughborough University

"In their latest work, Tom Fryer and Cristian Navarrete undertake the formidable task of unpacking the Realist Theory of Science—a subject as vital as it is vexing. With the finesse of master storytellers, they weave a narrative that not only clarifies but captivates, turning potential estrangement into intellectual engagement. Coming to critical realism late in my PhD, without a formal background in philosophy or sociology, I recognized CR as the approach I needed, yet felt daunted by the dense thicket of theory that seemed to guard its riches. Fryer and Navarrete act as skilled guides who offer the keys to unlock this gate, with an approach rich with examples and devoid of unnecessary jargon. This book ensures that what could alienate instead enlightens."

Sandy Rao, PhD Candidate, University of Calgary

Biographies

Tom Fryer is a Lecturer in Education at University of Manchester. His work focuses on opportunities to transform inequitable education systems, and has recently investigated access to university, graduate outcomes, and the impact of elite private education. He also has a particular interest in evaluation and research methods, and their underlying philosophical foundations.

Cristián Navarrete is a sociologist and MA in Sociology graduate from Pontificia Universidad Católica de Chile. His research draws on studies of science, technology, and medicine to address antimicrobial resistance from a sociohistorical perspective. He has conducted research on the gendered economies of scientific uncertainty, the concept of emergence, scientific realism, and has recently been working on problematic nature of "the future" as a research object within sociology. He is also interested in questions about the nature of society and scientific explanation.

Revisiting Realist Theory of Science: A Practical Guide

1st edition

Tom Fryer and Cristián Navarrete, 2024

Sharing

Open access: some rights reserved

Anyone can download, save, perform or distribute this work in any format, including translation, without written permission. This is subject to the terms of the Creative Commons License Deed: Attribution—NonCommercial—NoDerivatives 2.0 UK: England and Wales. Please ask the author for permission to use this work for purposes other than those covered by the license.

Cover Image: Joanna Kozak

Thanks

Thanks sections of books are hard to write so let's just put it like this: thanks to all of those that haven't rolled their eyes when we mention 'ontology'; thanks even more so to those of you that have helped us craft this text into something that is more readable, more robust, and more useful. Emphasis on 'more' – you can only imagine where the starting point was...

Contents

Introduction	6
Chapter 1: What is ontology and epistemology and why do we need them?	11
A framework	13
The so-called paradigms	17
Positivism/empiricism.....	18
Constructionism	20
New realist positions.....	23
Chapter 2: What is critical realism and how does it reach its conclusions?.....	27
Chapter 3A: General ontological conclusions from critical realism	29
Feature 1: Experiment is important to science	30
Feature 2: We apply knowledge outside of experiments.....	34
Feature 3: When science identifies a cause, this becomes something to be explained	37
Chapter 3B: Critical realist conclusions about the social world	42
Feature 4: Intentional action by scientists is needed	42
Feature 5: Scientific experiments are impacted by resources	44
Feature 6: In turn, resources are impacted by knowledge production.....	46
Other stuff you'll see in critical realism	49
Three domains of reality.....	49
Transitive and intransitive	51
Chapter 4: What are the implications for your work?.....	53
Section A: Concrete implications	55
Feature 1: Experiences, events and causes	55
Feature 2: Powers and necessity	58
Feature 3: The world is stratified, albeit in complex ways	59
Feature 4: Agency and intentional action.....	60
Features 5 and 6: Social structure and agency exist relationally	61
Bonus feature: Transitive and intransitive objects of science.....	62
Section B: Implications for methodology	65
Implications for quantitative methods	65
Implications for qualitative research.....	70
Implications for mixed methods research	73
A quick goodbye.....	75
Appendix: An annotated bibliography on critical realism's main critiques and contending positions.....	76

1. Critical realism is not alone nor the best	76
2. Naturalist critiques.....	77
3. Critical realism is not well grounded	77
4. Actor Network theory	77
5. Debates on cultural sociology.....	78
6. Emergence	78
References	80

Introduction

A Realist Theory of Science is an incredibly valuable book.

Revisiting Realist Theory of Science, dreamt up by yours truly (Tom and Cristián), is also a book.

In RTS, originally published in 1975, Roy Bhaskar manages to simultaneously strike a hammer blow to positivism, begin a rebuttal of social constructionism, and outline a nuanced description of the basic features of the world. That's pretty good going.

When I (Tom) first read RTS—or at least after I'd read Andrew Collier's brilliant *Critical Realism: An Introduction to Roy Bhaskar's Philosophy*, before returning to give RTS a second shot—I was blown away. I was in the early stages of my PhD, largely flip-flopping between different research questions and designs, and into that caffeinated confusion came RTS. It was exactly what I'd been looking for.

Sometimes you just read something and it resonates with you. It clarifies some vague and uncomfortable feeling that you've got that you've not quite been able to put into words. And more than that, it gives you some of the tools to move past the thing that was making you uncomfortable in the first place. That was RTS for me.

Let me be a bit more specific. I'd had a sneaky suspicion that ontology and epistemology were super important for research, but the way these were taught in research methods classes didn't seem to do them justice.

It's pretty common to hear something like: "Here's positivism and here's constructionism. Look they're different. That's funny. Now, pick one and make sure you stick with it through your PhD. See you later, ta ra".

At the other end of the spectrum, some of my other training sessions were more detailed, but quickly descended into the world's worst game of bingo. Interpretivism, tick. Post-positivism, tick. Hermeneutics, tick. Pragmatism, critical realism, constructivism, tick, tick, tick.

Does bingo even involve ticking? Not a clue.

In the face of this huge range of terms, it's understandable why lots of people either switch off or get overwhelmed. And let's face it, neither of these responses is a recipe for leaving with the energy and desire to think more about ontology and epistemology and how it impacts your research.

I really do think this is an issue. Or to bring Cristián back into this, we really think this is an issue. And it's one we hope to help address in this book.

We think there's better ways to teach ontology and epistemology than: "Here's the options, pick one and be consistent. Good lass, now off you pop." Don't get us wrong, we're firm believers that consistency is important (as long as this isn't blind faith to a position), but

there's more to it than that. Just like our research aims to use the best explanations and theories for whatever we're studying, the same applies to ontology and epistemology. Ideally, we should reflect on the strengths and limitations of our philosophical theories, the ways they have been critiqued by others, and then take an informed decision on which position (or combination of positions) to adopt. That's much more demanding than just being consistent.

Throughout this book we'll explain why we think ontology and epistemology are essential for everyone (well, all researchers at least). They have implications for your research questions, literature review, methodology, analysis, findings and conclusions. So, pretty much everything really.

To indulge in some shameless speculation, we think that if more social scientists had a grounding in ontology and epistemology there'd be tonnes of knock-on effects. Sure, a decent and thoughtful philosophical position doesn't guarantee that you'll produce good work, but adopting a dodgy position from the outset sure as hell makes it harder to do decent research.

This is exactly the journey we've been on—learning about these philosophies (and specifically critical realism) and thinking through the implications for our own work. For Tom this is educational research, and for Cristián this is health and sociology of science. This book is an attempt to introduce some of this thinking about ontology and epistemology, hopefully in a vaguely accessible way.

This accessibility is really important to us. Let's face it, philosophy can be a pretty daunting discipline for many applied social scientists. It's just not feasible for everyone to take 6 to 12 months out of their PhD to read and engage with this stuff. Sure, go for it, if that's what floats your boat—it'll be a fun conversation with your supervisor (cough, cough)—but for many this really isn't possible (or what you want to do, in any shape or form).

We think there's a real space for work that tries to translate the insights from philosophy of science, without necessarily getting bogged down in all the details.

We've already been working in this area a bit, grappling with how to produce these more accessible or 'bridging' works. I (Tom) wrote *A short guide to ontology and epistemology* as a way to get people started on this area, and I (Cristián) produced a Spanish-language version. We then collaborated to write a new and better version of the guide, which is available [here](#).

It's always amazing to hear people have read our guide and developed a bit of an interest in critical realism, but we've often struggled with how to help people take the next step in their learning journey. How could we help people explore things further? There's some great books out there, like those by Berth Danermark and colleagues (2019) or Doug Porpora (2015), but we always found ourselves coming back to the advice: "Well, if you really want to get all this stuff, as in like *really get it*, then you probably need to read RTS". There's just something fundamental about the way Bhaskar reaches his conclusions in RTS

(the method is as powerful as the conclusions) that means reading RTS is essential to really get what critical realism is on about.

But, and it's a big but, we always made this recommendation with a slightly heavy heart. RTS is a tough read. It's primarily aimed at a philosophy of science audience, and this means that it's engaged in a range of discipline-specific debates, the same type of niche debates you find in all areas of academia. This means it's not always the easiest text for practical social science researchers to engage with and learn from. That's why a core aim of this book is to communicate critical realism's insights and lessons to a non-specialist audience.

To put this another way, *Revisiting RTS* was born from the belief that RTS has amazing insights that all social science researchers could benefit from, but there's space for communicating this in a different way.

Now, onto the obligatory explanation of why we chose the title *Revisiting RTS*. We normally skip these parts of an introduction, so feel free to do the same (no offence taken) and rejoin us for the next section, but the title *does* capture some of the scope of what we're trying to do, and importantly, what we're not doing.

It's *Revisiting RTS* because it is very much a *revisit*. This means a couple of things. First, it's written for a new audience (applied researchers), and so it's a new set of people that we hope will get to learn from Bhaskar's work.

It's also revisited, rather than say *A Dummies Guide to Bhaskar*, because the original was published almost 50 years ago. Over this time, critical realism has evolved and changed, as has the wider field of philosophy of science. Critical realism is faced with new ideas and arguments that simply weren't around when Bhaskar was writing. This means there's value in revisiting and drawing on newer developments when telling the story of RTS.

Plus, there's also things we'd change in RTS. We didn't want to just blindly reproduce everything in an uncritical way, so there's things that we revisit and update. We're mostly going to stick to Bhaskar's central conclusions, but sometimes we'll introduce new things and refine other bits and bobs. We'll always try to point this out, just so you can put your extra-sceptical hats on for those sections.

Okay, if that's what we're doing, let's clarify what we're not doing.

One. Although the book is called *Revisiting RTS* and clearly prioritises the work of Bhaskar, we're not aiming to present him as the one and only founder of critical realism. His philosophical work does map out many of the basic characteristics of the world, but there's a range of ways critical realism's conclusions have been developed in conversation with others. In particular, we're thinking of Margaret Archer and her work on structure and agency, which goes far beyond anything in RTS. While it's not our intention to subsume Archer's work into Bhaskar through our process of revisiting RTS, we have sometimes decided to incorporate aspects of her work. We hope this adds value for you, our dear reader, while avoiding giving you the impression that this was all Bhaskar.

Two, and the flipside of One, we're equally not trying to summarise all of critical realism's contributions. For us, critical realism is an odd combination of meta-theory (which is more or less founded by Bhaskar) and social theory (which is more or less founded by Archer). In line with the original RTS, this book focuses squarely on the former. We aim to repeat the abductive process that Bhaskar engages in in RTS (don't worry if this is meaningless to you now!), and although this doesn't capture every critical realist conclusion, we do manage to establish an account of causation and a description of the basic ontological features of the world, which isn't bad going for a partial account. What we won't do in this book is engage with some of the specific social theories proposed by prominent critical realists, such as Archer's work on reflexivity, which we see as more of a social theory.

Caveat Three. There's times we might be a little flippant. There'll be times when we fail to engage with ongoing debates in critical realism in the level of detail that some people might want. We'll always try to point you in the direction of these debates, so that you know they exist, but our main intention is to provide an accessible account of critical realism and help you think through the ways you could practically apply this to your research. We're proud of this aim. It's important to who we are as researchers. But this does mean that sometimes we gloss over things for the sake of clarity. It's great if you want this extra detail, but as a head's up, this won't be the book to deliver it to you.

Four. For anyone who's more involved with critical realism, it's worth mentioning that this work is very much based on RTS and 'first wave' critical realism. We're not going to engage with Bhaskar's thinking in dialectical critical realism and metaReality. This could be an important project, but not one we undertake here.

So, to sum up, we're revisiting RTS with the aim of producing a more accessible and updated account, albeit one that doesn't summarise all of critical realism nor diminish the importance of other theorists. We see our revisit as a chance to look again at a brilliant and insightful text with fresh eyes. Our aim is to clarify, to improve, to sign-post other developments, and to leave you better prepared to think about the philosophical underpinnings of your own research.

We've organised things into four chapters:

1. What is ontology and epistemology and why do we need them?
2. What is critical realism and how does it reach its conclusions?
3. What are the main critical realist conclusions from RTS?
4. What are the implications for our work?

If you're familiar with RTS, you'll see that our organisation doesn't stick to its structure. This is because we're not aiming to give a commentary that follows chapter by chapter. We don't think this is the most effective way of introducing and building on the theorising from RTS.

Also, you might look at the first chapter and think it's not for you. The first chapter focuses on explaining ontology and epistemology and providing a simple framework to understand the main families of philosophical theories in this area: positivism, constructionism and other new realist positions (we see critical realism falling into the latter category). We've included this chapter to give readers a helping hand into the topic. We hope this is useful for

people who're less familiar with ontology and epistemology, but please do skip over Chapter 1 if this isn't useful for you.

It's the second and third chapters where we get to the heart of critical realism and its key ideas. The second chapter maps out the method Bhaskar uses in RTS to reach his conclusions, and the third (long) chapter maps out what these conclusions actually are.

We've then ended the book with a chapter focussed on methodology and the implications critical realism has for your research. In some ways this reflects our own bias – we're health (Cristián) and education (Tom) researchers. This means we're primarily interested in critical realism because it helps us do better research. Sure, sometimes we end up down an interesting philosophical rabbit hole, but ultimately what keeps us going is the practical usefulness of critical realism. This meant we didn't want to stop with an abstract discussion of ontology and epistemology, but instead we wanted to help you unpack the implications for your practical research. Again, we hope this is useful to some of you.

So, without further ado, let's get cracking.

Chapter 1: What is ontology and epistemology and why do we need them?

Okay, what's the point of thinking about ontology and epistemology?

First things first. What even are ontology and epistemology?

Generally speaking, ontology is the philosophical study of being. It looks at what stuff is in the world and what that stuff is like. It includes questions like: what's in the world or what's real?¹

The other key term is epistemology. Generally, epistemology refers to the study of knowledge. Instead of questions about what the world's like—that's ontology—epistemology is more interested in how we can produce knowledge about the world. Are there good and bad ways of producing knowledge? How reliable are our conclusions?

Now, it's pretty easy to get ontology and epistemology mixed up. To paraphrase a commentator on Twitter: 'If you know what ontology and epistemology are without Googling, then there's a 100% chance you're a douche'.

Ouch. That stings.

But it doesn't have to be this way. Here's a silly way to remember the difference:

- Ontology sounds like 'on toe logy', or the study of *what you just dropped on your toe*. Now, if you just dropped a hammer on your toe, we guarantee you're going to be thinking about reality. You're going to be thinking about real hammers and real pain, in the real world. There's no way you're going to be in the mind-frame to ask: "How do I produce knowledge about this hammer?" You'll be pretty focused on its reality. That's ontology.
- Epistemology sounds a bit like 'epic stem ology', or the study of *epic stems*. Imagine your mate, Dave, who is a plant scientist, comes up to you and says: "Hey buddy, look at the epic stem on this plant, how cool". We're guessing your first reaction will be: Is that really an epic stem? How does Dave know that's an epic stem? Why does Dave have some right weird opinions? You can see these are all questions about *knowledge*. That's epistemology.

A tad daft, we know. But hopefully it helps you remember.

Now we know the basics, let's have a go at convincing you it's actually important to think about ontology and epistemology.

Ontology is important because all research sets out to find things. Some research sets out to find universal laws of nature. Others embark on a journey to collect stories. And a few clever

¹ In anthropology and some areas of sociology it's common to talk of the 'ontological turn'. This is normally used to describe the idea that we need to treat the practices of other people as real, not just as beliefs. Here, we're using 'ontology' in a different way. We're using it in its philosophical sense (which we explain later), see Graeber (2017) and Turska & Ludwig (2023).

cookies set out to understand the nuanced action of causal mechanisms (more on this later).

Now, having a bad theory of ontology means that you risk setting off into the world to try and find things that don't exist. In this case, your research would be doomed to fail from the start.

Think about this example from Aldi (Tom's local supermarket):

I'm having a mid-life crisis and want to buy a Ferrari. I decide to go to Aldi to buy one. As much as I might want to buy a Ferrari, and as much as I might delude myself that Aldi actually sells them, this has absolutely no impact on the reality of the world. Aldi just doesn't sell Ferraris. It's silly to look for Ferraris in Aldi.

Research that uses a bad ontological assumption is as daft as the person having a mid-life crisis and looking for a Ferrari in Aldi. You might set out to find something in the world, but if you're wrong about the nature of the world, then you're never going to find it. It's a wild goose chase.

Actually, it's more like a wild unicorn chase, because geese actually do exist.

There's another way that a bad theory of ontology can set you off on the wrong path. Instead of looking for something that doesn't exist, this time our dodgy theory encourages us to ignore big chunks of the world. Again, let's go to Aldi:

I love carrots. Carrot soup, carrot mash, roasted up with a bit of olive oil and rosemary. Yummy, yummy, yummy, in my tummy, tummy, tummy. I buy all my carrots from Aldi. In fact, it's the only thing that I buy from Aldi.

This is all perfectly fine, if a bit weird. But if you go and tell your friend that *Aldi only sells carrots* then this would be a mistake. You'd be missing out on the big picture and all the other things that Aldi sells. They don't sell everything (no Ferraris for example), but it's definitely more than just carrots.

Now, this deluded carrot-loving fiend is a bit like researchers that only set out to find stories. There's nothing wrong with carrots and there's nothing wrong with collecting stories. It can actually be a super important kind of research, engaging with people that academic research has tended to ignore in the past. But this does become a problem when we assume there's nothing else in the world other than stories. Just like Aldi sells more than carrots, the world consists of more than stories.

So, to sum up, research needs a decent ontology, so that: 1) you don't go looking for something that doesn't exist, or 2) you don't ignore the full reality of the world. If you set off with a bad theory of ontology, you really are risking that you'll end up with a naff project.

Convinced? Maybe a little?

Key points to remember:

- *Ontology or 'on toe logy' is the study of being and reality, like the reality of dropping a hammer on your toe.*
- *Epistemology or 'epic stem logy' is the study of knowledge, or how Dave knows if that stem really is epic or not.*
- *You should think about your ontology and epistemology from the outset of your research design otherwise you risk searching for stuff that doesn't exist or missing out on important parts of the world.*

A framework

You might be minding your own business, eating a sandwich, dozing off in a lecture, in an intense meeting with your supervisor, or maybe even down the pub, but if you're a researcher then at some point you're going to get asked:

'What philosophical paradigm are you going to use?'

'What's your ontological and epistemological position?'

Panic stations.

What the hell are you going to say?

You've probably heard a ridiculous number of terms used to describe different philosophical positions. Some of these probably ring a bell: subjectivism, constructivism, interpretivism, phenomenology, hermeneutics, positivism, objectivism, post-positivism, post-modernism, critical realism, critical theory, post-colonialism, pragmatism, structuralism... The list could go on and on. But which one are you going to pick?

Let us tell you a secret. The person asking you these questions probably also hasn't got the foggiest what a philosophical paradigm is. Like you, they keep getting all the terms mixed up. Like you, they think that ontology and epistemology sound like a hospital department specialised on bones and stem cells (maybe that's just us...). Like you, they've probably been put off from looking into philosophy because too much of it seems dense for the sake of being dense, complex for the sake of being complex, and discombobulated for the sake of being discombobulated.

The second thing to note is that most of these complicated terms are bollocks.

Let's be precise with our 'bollocks'.

By bollocks, we don't mean useless. The terms do often refer to a particular research tradition, so in that sense they're helpful(ish) in placing someone within a particular body of work. But the terms really are pretty useless for *understanding and categorising different positions and theories*.

The terms are often a quick label to throw at different scholars. But there's a real danger that these labels lure us into thinking we know everything about a particular position, stopping us from engaging with it in a more detailed way.

In case it's not obvious from our tone, we think this superficial labelling can be a bit of a problem. It's a quick way to say: "Oh this person is X, but I'm Y, so I can just ignore what they say...". This labelling can limit the range of thinkers we engage with and learn from, and that's a problem.

But, getting back to the core issue, there's *tonnes* of different labels to describe philosophical positions. These are often used in different ways by different people, so it becomes almost impossible to work out what anything means.

We think there's a chance to look at this differently.

There's a chance to develop a consistent set of terms, so we all know what we're talking about, giving us a framework to navigate any conversation about ontology and epistemology. As an added bonus, this framework could also encourage us to engage with the nuance of what different theories or people have to say, rather than just throwing a quick label at them.

To develop this framework, we need to chuck the baby out with the bathwater. Actually, we need to rip out the bath and install a shower—who has time for a bath nowadays?²

We're proposing to go back to the basics of what makes a philosophical position and build up our concepts from there.

To be clear, in this chapter we're not going to make an argument about the basic characteristics of the world, i.e. we're not going to present a case for a particular ontological perspective. Instead, we're going to try and give you a simple framework that you can follow to navigate conversations with people about ontology and epistemology. We hope this helps you have the foundation (and confidence) to engage with and learn from different people and theories.

So, for any philosophical position we need to think about two things:

- Ontology
- Epistemology

Luckily, we already know what these are. Ontology is about stuff (dropping a hammer on your toe) and epistemology is about knowledge (how do you know that's an epic stem, mate?).

² It may interest a small and potentially negligible subsection of readers to learn that this book was written during a bathroom renovation, in which a perfectly functional shower was replaced with a bath precisely because it's hard to bathe a toddler in a bucket. It's always good to know that one of the authors of the book you're reading is an absolute hypocrite.

When it comes to ontology, you usually hear people talk about ‘realism’ and ‘irrealism’. However, we’re going to adapt this a bit and instead encourage you to think of ‘realism’ and ‘scepticism’. So, it’s realism/scepticism not realism/irrealism for us.

We’ve used these two terms to try and help us to be a bit more nuanced. It makes us ask: what does this position think is real and what’s it more sceptical about?

The term ‘irrealism’ just seems too prone to label-throwing: “That’s an irrealist position, so it’s clearly nonsense and I don’t need to engage with it”. In contrast, we reckon the term ‘scepticism’ is a bit more friendly—it kind of invites the follow-on question of: “What’s this position sceptical about then?” instead of just shutting down the conversation.

Another benefit of realism/scepticism compared with realism/irrealism is that *most people are realist about some things and sceptical about others*. Maybe you believe in atoms but not in God, or maybe you believe people are real but are sceptical about the reality of social structures.

The same applies to our philosophical and social theories—each theory tends to assume something is real. This means it’s not all that useful to label them as ‘realist’ or ‘irrealist’. It’s far better to ask what the theory thinks is real and what it’s sceptical about.

Take the example of Bruno Latour. It’s more useful to ask what things Latour is a realist about, and what he’s sceptical about, rather than rushing to apply a pre-made label to categorise his position. This pushes us to engage with the nuance of his ideas, which admittedly does mean we actually need to read his work. This is a healthier attitude and far better than simply writing people off by sticking an ‘irrealist’ label on them.

We’ll return to Latour a quite a few times over the course of this book, as Cristián is pretty convinced that people are too quick to label him an irrealist and then stop engaging with his work.

Okay, so that’s ontology, the first part of our philosophical position.

Moving onto the second part, epistemology, it’s also common for research methods textbooks to break this into two extremes: objectivism versus subjectivism.

It might not come as much of a surprise to say we’re not sure these are the most useful terms. After all, who doesn’t want *objective* knowledge?

We think it’s better to think about epistemology in terms of how *direct* or *indirect* our knowledge of the world is. Or to phrase this another way—is our knowledge of the world direct (unmediated) or is our knowledge mediated by other things (e.g. scientific instruments, concepts and theories, interests and biases)?

Some people adopt an epistemology that says we have direct knowledge of the world. From this perspective, there aren’t big barriers to producing knowledge—we can just sit down in our ergonomically designed chairs, observe the world, and produce truthful knowledge.

These claims of a direct epistemology have been widely critiqued, with many people arguing there's a bunch of things that do mediate knowledge production. These things could include any combination of scientific instruments, concepts and theories, individual biases and interests, and a range of other things.

If that's all a bit dense, let's try to explain this with an example. It's always easier to use a real-world example to chat about this kind of stuff.

Imagine that you're a researcher looking at how refugees are treated in the UK. One approach might assume we can just go out into the world, observe how refugees are treated, then write up our findings. That's all there is to it.

However, others might say that's way too simple. Who we consider to be a 'refugee' is theory-dependent, not 'neutral' and 'objective'. While some theories might make a big distinction between 'economic migrants' and 'refugees' (e.g. the former are migrating for economic benefits and the latter are forced to flee for political reasons), other theories might question this simple binary and argue a refugee is anyone who has fled their home country, for whatever reason.

In this way, the theory we use influences how we define 'refugee' which could have super important consequences for our research. Imagine if the UK treats those fleeing for political reasons reasonably well but is pretty inhumane to those fleeing due to economic factors. If we only look at the former, this would miss the inhumane treatment of the latter.

This argument could be used to support one form of indirect epistemology, i.e. that there's no way to produce knowledge without the mediation of theories and concepts.

Another example of an indirect epistemology, albeit one that focuses on how researchers' biases, interests and social positions can impact knowledge, is seen in feminist standpoint theory.

Proponents of this theory make the case that male-dominated scientific research has had various impacts on knowledge production. Research has done less to focus on problems that impact women (e.g. how menstrual cycles influence exercise (Carmichael et al, 2021)), this research often treats a male body as the default (e.g. crash test dummies shaped like men have led to higher rates of injuries for women (Criado-Perez, 2019)), and 'masculine' styles of thinking are often prioritised (e.g. abstract and acontextual reasoning (Anderson, 2024)). Proponents of feminist standpoint theory may argue we can only understand these three features of the world if we accept that researchers' biases, interests and social positions impact the knowledge they produce.

So again, this is an argument for a specific form of indirect epistemology.

It's also worth saying that an indirect epistemological position can hold that one or more things impact knowledge production. It's perfectly possible to accept that both theories *and* researcher biases impact knowledge.

Now, just like we argued it's helpful to move beyond the realist/irrealist labels and shift to more nuanced questions about *what is assumed to be real by a given position*, it's the same situation for epistemology. Rather than labelling a position as 'objectivist' or 'subjectivist', we think it's better to ask how a theory assumes we produce knowledge about the world. Is this direct or is this mediated by other things? And, if the latter, what are these things that do the mediating?

Although this is more complicated than just whazzing labels at different theories, we hope this invites you to consider the nuance of each position (what are people actually arguing?), rather than jumping the gun and labelling things without this deeper engagement.

To sum up, all philosophical positions can be boiled down to two fundamental questions:

- What is the position realist about?
- How does the position assume knowledge production is mediated, if at all?

This framework (i.e. the two questions above) tries to balance simplicity and nuance. It encourages you to only look for two things (ontology and epistemology), but by using *questions* rather than *abstract labels* this pushes you to engage with others' (and your own) philosophical positions in more detail.

We hope this framework gives you a basic map to navigate any discussion of philosophical positions that you encounter.

Key points to remember

- *A philosophical position is made of ontology and epistemology.*
- *There are two fundamental questions: 1) what does a position assume to be real? and 2) how does a position assume knowledge production is mediated?*

The so-called paradigms

Okay, so far, we've seen how you only really need to answer a couple of questions to start thinking about ontology and epistemology. We've been banging on about how you can use this framework to explore different theorists and positions in detail, rather than jumping to label them. However, this doesn't mean that there's no value in sketching out some broad families of philosophical positions (sometimes called paradigms) and giving you a kind of top-level overview of what they say.

Before we begin describing these different paradigms, hopefully it goes without saying that we need to take these with a pinch of salt. All the paradigms are simplified labels that try to describe a broad set of ideas. When you encounter a particular thinker, we'd always encourage you to move beyond these labels and see what they actually argue. But, like we said, there's some value in having a sense of the main paradigms and how they differ, and that's why we've included this section.

We're going to introduce three different positions: positivism, constructionism and new realisms.

Positivism/empiricism

In the philosophy of science, there's a diverse range of theories that are often called *logical positivism* or *logical empiricism*. Although it's debatable if this distinction is all that meaningful, the first of these is usually used to talk about an earlier and more extreme version of empiricism, mostly related to something called the Vienna Circle (Uebel, 2013). The second, logical empiricism, is normally used to talk about a later, more moderate and wider set of ideas, associated with thinkers like Carl Hempel, Hans Reichenbach and Kurt Lewin.

We've just thrown some jargon and names at you, so let's start to unpack this.

It probably makes sense to start by explaining what 'empiricism' means. This is normally summarised as the idea that *the only source of knowledge is experience*. Empiricism can be traced back at least as far as the classical works of John Locke, George Berkeley, and David Hume in the seventeenth and the eighteenth century.

A couple of dudes called Carl and Paul, but better known as Hempel and Oppenheim (1948), built on these thinkers to develop the Covering Law Theory of Explanation, which became one of the key approaches to producing explanations and knowledge in empiricism. The main idea is that *to explain something is to show how to infer it using a logical argument, where at least one statement uses a natural law*.

Without getting too bogged down in the details, from this perspective, the aim of science is to track and anticipate patterns in our *experiences* (remember, it is empiricism). It's by anticipating patterns in our experiences that we can start to see the natural laws that connect these experiences together (where a law of nature is just some kind of regularity or basic pattern of events). Once we've identified these natural laws, we can then use them to predict stuff. For example, there could be a natural law (or some people might say a 'statistical law') that states: gender predicts 12.5% of the variation in GCSE maths scores in England.

Now, the positivism that you'll encounter in social science is its own special breed. It's got some similarities with the philosophical theories above, but it's probably fair to say that it's often more extreme. In social science, it's more common to see 'positivists' make a strong claim that researchers can have direct access to a reality that exists independently of us: "Of course objective reality exists, dummy, we just need to develop ways to eliminate our bias".

This is actually a pretty big shift from logical empiricism in philosophy of science, where philosophers tend to dismiss talk of an independent objective reality as a meaningless metaphysical question.

In this book, we focus on the positivism you'll find in social science. This tries to find patterns between events with the ultimate goal of finding laws. Remember, a law just means a pattern between two events, like 'whenever we see the first event, then the second event will follow', e.g. whenever demand exceeds supply, prices will rise.

Now, there's a few issues with all of this, which mainly relate to positivism's concepts of experience and regularity. First, positivism takes a narrow view of reality, saying we can only know something to be real if we experience it, and that we need consistent experience to identify a causal law. The second issue is that positivism denies the many ways knowledge is mediated, only really accepting mediation in the form of the physical changing of human experience (e.g. using a microscope to help us see small stuff). In this way, positivism is interested in the reliability of scientific instruments, but puts less emphasis on other things that mediate knowledge (e.g. theories and concepts).

To better understand these issues, let's look at an example: gender and maths test scores.

Okay, so, maths scores might be influenced by a tonne of things, from teaching quality to confidence, from educational background to whether you consistently go to school hungry. Even if gender is one of the things that influences your score, it could interact with all these other factors in complex ways. Gender might interact with confidence, if girls are discouraged from grappling with problems: "Don't worry about it, Love, maths just isn't your thing". Gender might also interact with parental educational background, if students from less privileged backgrounds are more likely to have mothers who weren't encouraged to study maths, cementing the idea that it's not for girls. Gender could have similar interactions with teaching quality (e.g. if the teacher focuses mainly on boys), hunger (e.g. being hungry might be a tipping point for some girls who were already discouraged) and any other factor. All this might vary from person to person and from context to context.

The Covering Law Model of Explanation, and its emphasis on finding laws that connect experiences, can push us to search for simple relationships that simply don't exist. The world isn't built in a way that 'whenever X happens then Y happens next'. So, because things just aren't as regular as positivism would have us believe (or at least the naïve version of positivism we've described so far), it's an incorrect ontological assumption that there's a natural law connecting gender and maths scores. In this way, positivism can end up pushing science to search for unicorns.

Any positivistically inclined readers might be frowning and think we've not given a fair description of positivism. It's certainly the case that many positivists would accept that the world rarely has simple and consistent patterns of events, instead there's lots of things at play that make the pattern of events more complex. But what these more nuanced positivists would argue is that amongst this complexity, we do have statistical regularities. Here you could imagine a regression model that uses gender, parental education, teaching quality and socio-economic background to predict maths scores.

Now, even if we concede the point that some statistical regularities do exist at a given place and point in time (and depending what you chuck in your regression model), it's still the case that the mere finding of a regularity tells us nothing about *why these events are linked*. Finding a statistical regularity between two events, doesn't explain *why* those experiences are connected. To understand causation, and the reason why two events are linked, we need to go further than just measuring statistical associations. We need to understand not just the statistical regularity between gender and maths scores in a given context, but also *why* and *how* this relation comes to exist.

In this way, research that describes a pattern of events can be useful. As a bit of a spoiler alert, in Chapter 4 we unpack how quantitative methods can be used in appropriate ways. We see this as a process of adapting methods that have traditionally been grounded in positivism. It's worth flagging this now, just so no-one thinks we're arguing all statistical work should get in the bin.

But it's important to recognise this shouldn't usually be the end point of our work. We need to go a step further, developing an explanation of why these events relate to one another in the first place.

So, at this point, we've basically been arguing that there's a risk positivism's search for patterns of events and natural laws pushes us away from the goal of understanding why two events are related. It encourages us to stop at the statistical regularity and hinders the move to answer *why* and *how* questions.

Before we wrap up positivism, it's worth mentioning another problem, one that comes from its epistemology. This epistemology assumes that all knowledge is built from human experiences. The problem with this is that it doesn't recognise that knowledge production is mediated by concepts and theories, i.e. that our observations are theory-dependent.

Think back to the refugee and economic migrant example. By ignoring the ways that research is theory-dependent—e.g. drawing on a theory that defines a 'refugee'—positivism can misrepresent knowledge production. The theories we use don't just enhance our observations, but they modify or condition them. Any adequate epistemological position has to acknowledge this.

One consequence of this mistake is that positivism doesn't always recognise the need for reflection in the research process. We need to think about the theories we're using, where they came from, who they benefit, and who they exclude. This is our responsibility as researchers.

Constructionism

Let's move onto our second paradigm, constructionism. As with positivism, this is far from a unified position in the philosophy of science.

Ian Hacking (1999) says there's an important difference between two types of constructionist claims: 1) that which argues a great deal of our world is socially constructed; and 2) that which argues a particular thing is socially constructed, e.g. the social construction of gender.

But, what exactly is a social construction?

Dave Elder-Vass (2012) says that this is essentially the idea that 'the ways...we collectively think and communicate about the world affect the way that the world is'. So, saying something is socially constructed involves the claim that *a particular thing exists because of the way we collectively think and communicate*.

Or put a different way, if humans were all wiped out by an asteroid tomorrow, socially constructed things would cease to exist. There'd still be rocks (not a social construct) but there'd be no citizenship (a social construct).

Returning to Hacking, he says it's also useful to think about the aims of a researcher when they say something is socially constructed. According to him, people are trying to make the point that 'something' doesn't need to have existed, or that 'something' doesn't need to have existed in its current form. In this way, the claim something is socially constructed is often linked to the idea that it isn't necessary or inevitable.

Some claims go a step further to argue that *not only is something unnecessary, but it's also pretty problematic*. In this case, the claim that something is socially constructed opens the door to the argument: "We'd be better off if we constructed this 'something' in a different way".

Take the example of someone claiming contemporary masculinity in Chile is a social construction. This involves the claim that Chilean masculinity could have been different to how it is now, i.e. there's no reason why it couldn't have been constructed in a different way. More than this, the claim also links to the idea that some elements of Chilean masculinity are undesirable (think of toxic masculinity), and we'd be in a better position if we changed the way it's constructed.

When we think about Hacking's two types of construction, it's most common to see a claim that a particular thing has a constructed reality (Hacking's second type of claim), rather than a stronger claim about the world in general (the first type of claim).

Now, given that some things definitely *are* constructed (think of obvious stuff, like our gender roles), this means a lot of 'constructionist' research focused on these topics is great.

However, sometimes constructionism (especially Hacking's first type of constructionism) does run into issues. In particular, it adopts too narrow a perspective on the world by downplaying the role that things play in producing phenomena. Sure, many of these things may be socially constructed, but 1) once constructed they take on an existence that, to some extent, escapes the control of the individuals that constructed them, and 2) they're often constructed on the basis of something real (that our mind doesn't control) that does limit and influence the social construction. We can't always construct whatever we want.

An example of 1) would be the construction of a 'good life' in Christianity, which when shared escaped the control of the original prophets and writers. An example of 2) would be that while the category of 'dog' is socially constructed, there's something real that prevents this category including both poodles and jellyfish.

This means constructionism doesn't adopt a false ontology (and in that sense it's a step forward from naïve positivism's search for simple natural laws). Instead, the issue is more

that it's just too narrow in its outlook. It focuses too much on meaning, discourse and ideas, and can neglect all the other things that are in the world.

At this point, you might be wondering how it's possible to justify the idea that constructionism takes too narrow an ontological perspective.

One way to do this is to show you the importance of considering stuff beyond meaning, discourse and ideas, that have tangible impacts on the world. We're going to do this by showing you that social structures are real and have a causal impact. Think about the example of university access, which we've adapted from Margaret Archer (2003):

When a person decides whether or not to go to university, they are impacted by various social structures that shape the situation in which they make their decision, and the costs associated with the various options. If a person is from a less privileged background, has caring responsibilities or has been out of education for a while, these all influence the situation in which the person makes a decision, making it harder to go to university because the opportunity cost is higher.

In this example, real social structures impact the world and a person's decision of whether to attend university or not. The wealth you have access to at a particular point in time exists independently of your mind, i.e. your access to finances is not fully determined by how you think about money.³ Instead, there's a real relation in the world (i.e. a relation between people and wealth), and this has impacts on the decisions and actions people take. This is why it's important for social science to study these social structures and their impacts in the world.

Now, there might be a number of self-professed constructionists that would agree it's important to study social structures. This is probably a good time to clarify that in our characterisation many people who describe their work as constructionist would actually be closer to the 'new realist' position. Remember, our aim has been to provide a broad understanding of different paradigms to help you have a conversation with people. Unfortunately, this doesn't get around the problem that people use the same terms to mean very different things!

Okay, back to constructionism. The epistemological position is also important to consider. In this, there's a divergence between extreme and less extreme versions of constructionism. Extreme constructionism argues that knowledge production is *theory-determined*, not just *theory-dependent*. This denies that the world influences our theories, and ends up in a position of complete relativism, where there's no way of saying one explanation is better or worse than another.

Imagine you go to the pub and the bloke at the bar starts telling you about how babies are brought by a flying stork after being made in a secret lake. That's one theory you could hold about human reproduction. But just because this fella believes the theory doesn't make the

³ Notice that we have consciously opted to not talk about mind independence. This is because, following Chang (2022), we think the concept isn't very clear. It's better to talk about 'mind control' and 'mind framing'.

world that way. His theory doesn't *determine* whether babies *really are* formed in a secret lake. Our observations might be dependent on our theories, but our theories don't fully determine what reality is like.

Now, there's clearly not that many people that believe this more extreme version of constructionism. So why have we bothered including this?

Although few people would sign up to these ideas, we do think constructionism influences research in more subtle ways.

One of these influences is seen in the idea that all research can do is construct stories and worldviews.

This is crazy.

Sure, it's important to ask people's opinions about things, but research shouldn't just turn these into stories. A study of poverty doesn't have to end when we've got stories about poverty. Instead, we could go further and identify the causes of why people are in poverty and how these barriers could be transformed.

Let's drum this home with another example.

Imagine a study that analyses the ways refugees are represented in discourse—whether linking refugees to diseases or the ways in which a binary is created between refugees and economic migrants. This is good and important research. The problem is when we assume the analysis has to stop here. We could still ask questions like:

- What impact does this discourse have in the world?
- In what ways do people suffer because of this discourse?
- To what extent do different people adopt this discourse and how does this influence their actions?
- What other 'non-social' elements impact these discourses?

Hopefully these questions show that the fact there is a discourse does absolutely nothing to show the *impact* of this discourse in the world. But we can (and should) study the impact of this discourse on policy, media discussion, and public perceptions. The issue is that more extreme forms of constructionism deny this is possible, focusing only on constructions, which discourages us from undertaking this important further work.

Again, we want to stress that many self-professed 'constructionists' would agree this extra research is important. That's why we'd say these scholars are better labelled 'new realist' in our framework.

New realist positions

There have been a number of developments in the philosophy of science in the second part of the 21st century. Many of these developments involve scholars that are keen to distance themselves from positivism, but equally don't want to embrace constructionism. For example, these people might critique positivism's idea of causation as event regularities, but

they don't think that the ontological scepticism of constructionism provides an answer. Instead, many of these people have sought to maintain a realist stance, alongside their rejection of positivism.

We think these various positions can be grouped together, and we're going to call them 'new realist' positions.

These new realist positions build on both positivism and constructionism, acknowledging that the world is 'real', but that knowledge production is fallible and theory-dependent (not theory-determined). And, while they're happy to say that meaning and discourse are important, they make a strong case that these aren't the only things that exist.

You might have noticed that we're talking about new realist positions – plural. That's because we think the middle ground between positivism and constructionism actually has a bunch of different (and somewhat complementary) positions. As mentioned, these are often united by a rejection of positivism's epistemology and its focus on natural laws, and by a belief that constructionism doesn't provide an appropriate alternative due to its overly sceptical ontology.

To be a bit more specific about the types of approaches we're talking about within the 'new realist' category, we'd include much of pragmatism. In this we're similar to Hasok Chang, a leading pragmatist scholar, who wrote a book *Realism for realistic people*, that explicitly argues pragmatism should be seen as a realist theory. And we'd also include other realisms such as speculative realism, perspectival realism, new materialism and actor-network theory. The list could go on and on, but we hope some of that name dropping is useful.

At this point, it might be useful to take the example of Bruno Latour to flesh out the new realist position, as we see it. Latour has consistently said he's a realist, but people are still keen to label him as irrealist, constructivist or even postmodernist. This is where labels fall short. We've argued it's far better to look at his position and ask two main questions (what is Latour realist about, and how does Latour think knowledge production is mediated?). If you do this, we think you find Latour advocates a realist position, albeit one that accepts a relatively extreme version of knowledge mediation.

Let's take the example of a test for a urinary infection. Latour argues this test isn't a simple observation of the world. Instead, a doctor needs to take a urine sample and send this onto a laboratory. In the lab a small sample is taken, then microscopes are used to assess the type of microorganisms that are growing. These results are then put into words, before being printed off in a report. In this way, there's several steps from the test to the results.

Now, Latour has no quibble that microorganisms are real, and in this he's a realist, albeit not one focused on natural laws and a narrow conception of experience. But, he would argue we only know that microorganisms are real because of the long process outlined above. To study things in the world, we need to make a large effort to interact with things and bring them into relation with ourselves.

In this way, Latour is arguing for a relatively extreme version of knowledge mediation. This mediation does lead Latour to argue it makes little sense to talk about microbes as 'independent' of us, but we see this as an epistemological rather than an ontological argument. In other words, this seems to be more related to Latour's idea of how we come to gain knowledge about microbes, rather than a claim about what microbes actually are.

So, for us, Latour is like many new realists in abandoning positivism's definitions of objective reality, natural laws, and experience, while not turning to constructionism's scepticism as a solution.

Now, this book is focussed on one of these new realisms: critical realism. We think this is a particularly useful position, in no small part because of the powerful method through which Bhaskar reaches his ontological conclusions in RTS. There'll be much more on this in the next two chapters.

By locating critical realism as one of several new realist positions we're trying to stress the point that we see the tradition as part of a broader family. It's a family that we critical realists could probably do more to engage with and learn from than we do at the moment. But that would be a whole new project of its own, one that builds on existing engagements such as Tim Rutzou and Dave Elder-Vass's (2019) work with Gilles Deleuze and Manuel DeLanda; Hasok Chang's (2022) engagement with pragmatism and Annemarie Mol; or David Graeber's (2015) debate with Viveiros de Castro's position. These all share a realist position but in different ways, and this pluralism is an invitation to work out 'realism' beyond disagreements.

You'll see the impact of our positioning of critical realism as one realist position among many when we turn to Chapter 4, which begins to explore methodology and methods. In this we draw on a range of work, including some new realist work outside of critical realism, that we believe are compatible with the ontological conclusions we outline in this book. We do this precisely because we think there's considerable overlap between various new realist positions.

Anyhow, with this caveat aside, let's end this chapter with a quick rant about people who say philosophical questions of ontology and epistemology don't matter for researchers.

You'll often hear people say things like "It's okay to adopt whichever philosophical position you want. All you need to do is make sure it's consistent". If we were betting men, we'd guess that most methods training in universities tend to adopt this position. In some ways this makes sense. If you have a class full of researchers from very different disciplines, then it might be particularly challenging to make a strong case for one philosophical position (or even one 'family' of positions). But, in other ways, it's utter rubbish.

I hope you can see why. Without a decent theory of ontology and epistemology you could go looking for stuff that doesn't exist, or you could end up accidentally ignoring a big chunk of the world. Being consistent doesn't make these problems go away.

What we need is an appropriate and robust philosophical position, not just one that we apply consistently. We think critical realism provides a set of tools to help us move towards this goal. It might not get us there alone, at least not without engaging with other new realist positions, but it's certainly a big help on the journey.

Key points to remember:

- *Positivism goes looking for universal laws that don't exist or at best explain little. It gives a very shallow perspective on experience.*
- *Constructionism encourages researchers to only consider meaning and discourse. We must go further to look at causes, social structures and their impacts.*
- *You can't avoid ontology and epistemology in your research, so don't just ignore it.*

Chapter 2: What is critical realism and how does it reach its conclusions?

Okay, so what is this critical realism we've been talking about? And why is it a powerful approach to derive an ontological and epistemological position?

Let's start with the second question.

Critical realism offers a robust theory of ontology and epistemology because it reaches its conclusions through a powerful method called abductive reasoning. You'll also see this called retroductive reasoning by some critical realists, but we're going to use 'abductive' as this is the more general term used in philosophy of science.⁴

Don't worry, abductive reasoning isn't as complicated as it sounds. It's basically a kind of logic that looks for the best explanation. Say you just put some chocolate next to your baby, and then a minute later the chocolate has gone, and Little Billy has chocolate smeared around his face. What's the best explanation for this? We'd be justified in concluding that 'Little Billy ate the chocolate'.

In RTS, Bhaskar used a similar kind of reasoning, sometimes called transcendental reasoning, to make the case for what the ontology of the world must be like. He asked the slightly odd question of: *what must the ontology of the world be like in order for scientific experimental practice to be possible?* Or, phrased a different way: if scientists think that experiments are worth doing, i.e. if scientific experiments are intelligible as a practice, what does this tell us about the ontology of the world?⁵

A key part of Bhaskar's approach is to take *scientific experiments* as an appropriate way to produce knowledge from the outset. Instead of trying to prove the validity of knowledge produced within scientific experiments, he assumes this knowledge is valid. After all, if this knowledge wasn't valid, why would scientists bother with experiments in the first place? He then moves on to ask what the world must be like in order for this to be the case.

You might be thinking this is a bit controversial. Can Bhaskar really just assume scientific experiments produce knowledge and use this to make an argument about the basic ontological characteristics of the world?

We don't think this is too big of a deal. In fact, chances are there's a bunch of ways that you live your life that assumes scientific experiments produce knowledge. To give just one

⁴ See Bridget Ritz (2020) for a detailed discussion of the difference between Peircean abduction and critical realist retroduction. In brief, the first refers to the process of identifying a range of possible facts that if true could explain an observed phenomenon, while the latter refers to a later phase of research where one reaches a conclusion about *what must be the case* for the phenomenon to happen as it did.

⁵ For those of you that're interested, there's a fair amount of debate about the exact question and type of reasoning that's used by Bhaskar. Some versions are framed in terms of the *best explanation* for how experiments produce reliable knowledge, pushing us to an empirical debate about the best explanations or theories. But, other framings are more *transcendental* and less in need of empirical corroboration. For more on these issues, see Tuukka Kaidesoja (2013) or Matti Sarkia & Kaidesoja (2023).

example, if you're happy to rely on drugs prescribed by your doctor, then you're probably okay with the idea that scientific experiments can produce knowledge.

But, there definitely are ways Bhaskar's abductive argument could be faulty.

A key way critical realist conclusions could be wrong is if there's some dodgy reasoning in the argument about what the world must be like, for scientific experiment to be possible or intelligible. In other words, maybe Bhaskar has proposed an ontological feature of the world, but this feature isn't actually required in order to understand why scientific experiments are possible. This would mean the ontological conclusion is unsupported by the abductive reasoning process.

Now, it's definitely plausible that critical realism has made mistakes like this. But this wouldn't make critical realism useless. Correcting this type of mistake would only make the tradition a better account of the world. Critical realism's a bit like the hulk getting angrier, the more you correct (or add to) the abductive argument the stronger it gets...

That doesn't work perfectly as an analogy, but you get the idea.

For anyone that's interested, it's this analogy that inspired the cover image for our book.

As an aside, another method Bhaskar uses, mainly to assess and critique rival philosophical positions, is *immanent critique*. This method is used to uncover rival's weaknesses and inconsistencies. The method begins by accepting the assumptions of a given philosophical position, and then by a process of extension reveals the perverse conclusions that result. For example, positivism adheres to the empiricist claim that for a statement to have meaning, it has to be tested against experience. However, we have no way of distinguishing between which experiences are more or less useful for knowledge generation, which ends up with the idea that *all observations of the world are equal in their ability to produce knowledge*. When a student sits in a dull lecture and counts the number of times they can spin a pen around their fingers, this involves an empirical observation just as much as the scientist who measures the result of an experiment. But clearly there's something different about these two experiences. The observation in the scientific experiment is more likely to lead to knowledge production than the pen-spinning. But positivism provides no basis for understanding this difference, so in this way, immanent critique has revealed positivism to be an inadequate account of science.

That's the end of this short chapter. We thought it was worth making this a standalone section just to make sure it didn't get lost elsewhere. The abductive method is so central to critical realism and why its conclusions are robust, that we couldn't take this risk.

This chapter also highlights the ways you can get involved in developing critical realism. If you think Bhaskar (or we) make a mistake in the reasoning we'll outline in Chapter 3, then you could help to develop and improve things.

Consider yourself invited.

Chapter 3A: General ontological conclusions from critical realism

For those that read the first few chapters, we hope they were useful for you. And for those of you that are joining us here, welcome.

In this chapter we're going to follow a step-by-step process to explore Bhaskar's abductive reasoning in RTS. Remember, the main abductive question he's asking is: 'What must the basic characteristics of the world be, in order for scientific experiment to be possible?'. It's from this question that Bhaskar's able to reach his ontological conclusions, which are central to critical realism.

Our reading of Bhaskar is that he doesn't simply ask one big question about what the world must be like for science to be possible. Instead, he seems to pick out a few different features of science, and then asks what the world must be like for each of these features to be possible.

So, what're these features of science that Bhaskar identifies?

We think there's at least four main features in RTS, which Bhaskar uses in his abductive reasoning. These are:

1. Experiment is important to science.
2. Knowledge generated from experiments can be applied outside of experimental situations.
3. When science finds a cause, this itself becomes something to be explained.
4. Intentional action by scientists is needed.

In RTS, Bhaskar also has a couple of other features: a) that scientists need training to produce knowledge, and b) that scientific knowledge changes over time.

We've chosen not to use these latter two features in this book. Ultimately, we think it's possible to reach the same conclusions as Bhaskar, and also extend these somewhat, if we adopt two different features of science.

Although we'll make it clear what we're adding/tweaking/updating in the relevant sections, at this point it might be useful to know that the two other features of science we include in this chapter are:

5. Scientific experiments are influenced by resources.
6. In turn, resources are impacted by knowledge production.

Also, as a quick head's up, this is going to be a long chapter. In many ways it's the heart of the book, as it's the place where we really look at RTS and *revisit* it.

So you don't get lost along the way, here's a roadmap of what we're going to do:

- This chapter is structured around the six features listed above. For each, we ask an abductive question (What must the world be like for scientific experiment to have this particular feature?) and use this to derive conclusions about the basic ontological characteristics of the world. Unsurprisingly, we'll just progress one by one through each of these six features.

- To make life a bit easier, we've actually divided things in two. Chapter 3A focuses on the general ontological conclusions that critical realism provides us. This covers Features 1-3.
- Chapter 3B focuses on critical realist conclusions about the social world, which deals with Features 4-6.
- The next chapter (i.e. Chapter 4) takes these ontological conclusions and starts to think about the practical consequences for your research. While this won't give you a fully fleshed out research design, it'll give you a damn good start on this journey. Our hope is that Chapter 4's practical focus will help you put some of the abstract ideas from Chapter 3 into practice.

Okay, so onto the first feature of science, and what we can learn about the world from this.

Feature 1: Experiment is important to science

Bhaskar's first feature of science is that experiments are important.

If this statement rubs anyone up the wrong way from the outset "Ahh, I hate people that say randomised control trials are the only way to produce knowledge" then 1) we feel you, this is clearly a problematic opinion, and 2) Bhaskar is mainly talking about the natural sciences here (like physics, chemistry and biology), so it's less controversial than you might initially think.

As a side note, the reason why Bhaskar's focussing on the natural sciences is because (at least in part), if he can show that positivism is bollocks *even in the natural sciences*, then this is even more of a powerful critique than showing positivism doesn't work for the social sciences.

So, when it comes to the natural sciences, Bhaskar is saying that experiments are important to knowledge production. That's his first feature of science.

For many of you this might feel pretty obvious. We've all sat in science classes as kids, we've all been taught that experiments are important to science, and most of us will have popped on a lab coat and messed about with the odd pipette and Bunsen burner. Experiments are important to science. Simple.

The next step in Bhaskar's reasoning is to make this feature of science a bit more specific. He says experiments are important to science because observations under experimental conditions are often more useful than observations under non-experimental conditions.

Again, hopefully that's fairly obvious stuff. Not all observations are equal.

For those of you that need more convincing, think about this: a chemistry researcher doesn't just plod around the world making random observations in order to write their papers. Sure, they might have an idea about chemistry when they're out for a walk, but it's their observations under experimental conditions that are more likely to generate knowledge about chemistry. This doesn't mean everyday observations are useless (there's always things to be observed on any good walk), but it's observations under experimental conditions that are particularly important to scientific knowledge production.

We then turn to the abductive question for this feature of science: if there's a difference between observations under experimental and non-experimental conditions, what must the world be like for this to be so?

The first conclusion Bhaskar reaches is that we can only account for the difference between experimental and non-experimental observation if there's a distinction between *open* and *closed* systems.

Scientific experiments are *closed* in the sense that, while there are many things acting (such as humans and technical equipment), they are designed so that these factors don't interfere with the things we are interested in. And, because it's possible to isolate this thing, we can study and predict its impact on the world.

This isn't normally true of non-experimental observations. Instead, non-experimental observations tend to happen in *open* systems that involve a bunch of different things acting at the same time. All these things can interact with each other, making stuff really messy and complex. One thing might inhibit, enhance or reverse the impact of another thing. All this complexity is why non-experimental observations are different to observations under experimental conditions.

Okay, so we need the idea of open and closed systems in order to understand why scientific experiment is important to science.

But Bhaskar doesn't stop there. He thinks there's other important conclusions that we can reach from this feature of science. He says that to account for this difference between experimental and non-experimental observations the world must have non-Humean causation (don't worry we'll explain this later). More specific than this, he says that the non-Humean form of causation that the world must have involves three key concepts:

- Experiences - these are the perceptions of things by agents, e.g. I experience eating chips by the seaside.
- Events - these are the things that happen in the world, some of which are perceived by agents, e.g. a bird poops on my head.
- Causes⁶ - these are the things that produce events, e.g. financial precarity constrains a person's decision to take an unpaid internship.

Armed with these three concepts we're in a position to account for the difference between experimental and non-experimental observation.

Experimental observations take place in closed environments, meaning that only one cause acts to bring about an event. This means experimental observations are particularly well placed to understand the nature and functioning of a given cause because the environment

⁶ Bhaskar uses the term 'causal mechanism', whereas we've used the term 'cause'. This change is in response to an ambiguity, noted by Ruth Groff (2017), that mechanisms can be used to refer to a cause (the propertied thing) but also to the propertied thing's power. A third alternative, which was pointed out by Dave Elder-Vass in an earlier draft, is that mechanisms refer to the processes that produce the powers of things. To avoid all this conflation, we use the terms *propertied things* and their *powers*, with only the former acting as causes. Hopefully, this is more precise and helps avoid ambiguity.

around the cause has been controlled, helping us to remove (or at least reduce) its complex interactions with other things. In contrast, in non-experimental contexts, events are produced in open systems where many causes can act and interact in complex ways.

If this doesn't seem all that revolutionary, it's probably worth telling you a little more about Humean causation, just so that we can get a better idea of what Bhaskar is arguing against.

Humean causation gets its name from David Hume. Central to Hume's account is the idea of *constant conjunctions of events*. We see a leg swing towards a football and then the ball moves away. These two events are a constant conjunction in the sense that the first event (the swinging of the leg) comes before the second event (the ball moving away). Swing of a leg, the ball moves away. Swing of a leg, the ball moves away.

So far so good.

Now, Hume goes on to say that there's nothing that really connects the first and the second event. So, there's nothing that really connects the swinging of the leg and the ball moving away. The only connection between the two events is a *feeling of anticipation* that an observer has when watching the world. This feeling of anticipation is then projected onto the world, and that's all there is to causation.

So, although a Humean would call the first event (the leg swinging) a cause, they think that there's nothing about the swinging of the leg that actually *brings about* the movement of the ball.

Ruth Groff (2021) explains that this Humean position ends up assuming that the world is a little bit like a child's flipbook or a reel of film. Even though it might look like the images in the flipbook are moving and doing things, this is an illusion created by our mind. The flipbook is really a series of still images, but our mind creates an illusion that the things are moving. In the same way, a Humean views the world as a really passive place and thinks it's only our interpretations and anticipations that give the world movement and action. That's why this position is sometimes called passivist.

If you think this is a bit bonkers, don't worry, you're not alone.

In fact, there's a bunch of ways this Humean position on causation has been critiqued. Most traditions within philosophy of science have some way of saying that this must be a load of nonsense. Because we're writing a book called *Revisiting RTS* we're obviously going to focus on Bhaskar's critique, but it's worth flagging that this is far from the only way to criticise Humean causation.

Okay, so what's Bhaskar's critique?

Bhaskar critiques the Humean position of causation using immanent critique. Remember, immanent critique is basically a form of argument that: 1) accepts the basic premises of a theory, and then 2) uses these premises to assess if we reach any crazy conclusions by following these premises through to their logical end points.

The main premises of the Humean account are that causation consists of constant conjunctions of events and the only necessary connection between events is a feeling of anticipation that's projected onto the world. If we accept these premises, then one absolutely batshit conclusion we reach is that there's no difference between observations under experimental and non-experimental conditions (i.e. Bhaskar's first feature of science).

See how this all interconnects?

There's no way to explain why experiments are important to science if we start from the premise that the only thing connecting events is a feeling of anticipation. This isn't some illusion projected onto the world, but something real about the way the world is. It's only if we accept that closed systems really do create a real context in which one thing can act, that we can begin to understand why experiments are important to science. This all has absolutely nothing to do with the expectation of observers.

So, Humean causation is unable to explain the difference between experimental and non-experimental observations, and this shows it's a load of tosh. It's clearly not the way the world works. Experiments are important to science, and our approach to causation needs to be able to explain this. That's some pretty compelling evidence that the Humean position is fundamentally flawed.

Linking this back to the discussion from previous chapters, a Humean approach to causation is closely tied to positivism. So, you can think of the weaknesses of Humean causation as yet another nail in the coffin for positivism.

Okay, now we've got a better idea of what theory of causation Bhaskar is rejecting, let's revisit his three concepts of experiences, events and causes and unpack what his alternative looks like.

Bhaskar proposes a powers-based approach to causation that's anti-passivist, i.e. it doesn't see the world as passive like a child's flipbook. This anti-passivist approach says the world contains active things, things that are able to do stuff, and things that are able to act. From this perspective, causation isn't the projection of anticipation onto the world, but the action of *propertied things*.⁷

Propertied things have different powers (which they have because of their internal structure) and it's the action of these propertied things that makes the world active.

⁷ We follow Groff (2020) in emphasising *propertied things* rather than *things and their powers*. The latter formulation implies that we can divorce things and their powers, but this doesn't make much sense. Groff (2020, p. 36) explains this citing an example from E.J. Lowe (2012, p.247) about why some objects roll downhill and others do not: "[It makes little] sense to suppose that it's sphericity or cylindricality 'drags along' the object's other [properties] with it, and thereby makes the object as a whole move. To think in such terms is illicitly to hypostatize [properties], treating them as substances within an object, rather than just a particular 'ways' an object is".

There's a fair amount going on in that sentence, from 'powers' to 'internal structure', from 'propertied things' to 'activity', so this all needs some unpacking.

Actually, on second thoughts, let's leave you hanging for a bit.

Sorry.

It actually makes more sense to unpack this all in Feature 2, after we've introduced a couple of other things.

At this point, it probably makes more sense to just take a minute to clarify how critical realism's approach to causation differs from positivism and constructionism. But we promise we'll come back to *propertied things* soon.

Now, we've already mentioned that positivism adopts a Humean approach to causation. This means it's unable to explain why experimental observations are important to science. More than this, it's flawed approach to causation pushes researchers to focus on patterns of events rather than uncovering the actual causes that underlie these patterns. We've also seen that we need to know not only that two events have a particular pattern, but we also need to understand *why* they have this pattern to have an adequate account of causation.

On the other hand, constructionism also has an issue accounting for the difference between experimental and non-experimental observations. Although it rightly rejects positivism's Humean approach, it throws the baby out with the bathwater. As we mentioned, constructionism can, at least implicitly, deny the reality of a range of propertied things in the world and the way they can act as causes, focussing only on things like meaning and discourse. However, Bhaskar makes it clear that we need the concepts of experiences, events and causes to account for the difference between experimental and non-experimental observation. To ignore causes means that constructionism can't account for this feature of scientific experiment, and this is a pretty clear signal that there's a problem with the theory.

Feature 2: We apply knowledge outside of experiments

Moving onto the second feature of scientific experiment: that the knowledge we produce can be applied outside of experiments.

Again, hopefully this is fairly intuitive.

If a medical trial shows a drug effectively cures a disease, then we're justified in applying this knowledge outside the experimental context. Sure thing, we don't know for certain that the drug will be effective, but it's definitely reasonable to think that it will be. This is even more reasonable if we've some idea of how the drug impacts the body, i.e. if we've some explanation of why the drug has the impact it does.

Or, to put this another way, we can apply knowledge from experimental conditions to the rest of the world, albeit in ways that are rarely fully predictable, and this unpredictability stems from the fact that more than one thing is acting.

Okay, so with this feature of science, we can again ask the abductive question: given this feature, what must the world be like?

Bhaskar argues that to account for this feature we need a *theory of natural necessity*. This theory of natural necessity aims to explain both the *applicability* and *unpredictability* of our causal knowledge in open systems.

At its most basic, the theory says that: 1) entities have powers, i.e. they are propertied things, and 2) entities have these powers because of the way they are, i.e. they have these powers because of their internal structure.

Let's explain this by returning to the medical trial example. The drug has the power to cure a particular disease because of its internal structure, i.e. its chemical composition. This chemical composition isn't going to spontaneously change as we move from experimental to non-experimental contexts. The chemical composition won't change just because someone in a lab coat is taking observations. It's this stability in internal structure that explains why the drug's powers stay the same in and out of experimental conditions. It's this that justifies our claim that the drug is likely to be effective outside of experiments.

To re-iterate, there's no guarantee that the drug will work. We live in an open world in which many propertied things act, so there's always a chance that something else acts to hinder or reverse the drug's power. Maybe the drug is inhibited by spinach, but this wasn't found in the trial because patients were kept on a strict diet of cheese toasties.

But this spinach-based caveat aside, it's still reasonable to apply our knowledge outside of experimental contexts. This applicability is because the drug's powers stay the same (as its internal structure stays the same), even if its impact on *events* (i.e. whether we're cured or not) isn't 100% predictable in an open system.

One word of warning on Bhaskar's theory of natural necessity: don't interpret the term 'necessity' as an argument for a deterministic worldview. Bhaskar isn't saying that an event is *necessary* once we know a particular cause is acting. Instead, he's saying that things with a certain structure *necessarily* have certain powers, but whether or not those things exercise their powers in the world, and what impact they have, is dependent on a bunch of other things.

All of this is fundamental enough to critical realism that it's worth taking a second example to drum things home. We'll do this by adapting (or bastardising, depending on your perspective), Andrew Collier's (1994) example of cricket.

Now, dear Reader, we should probably warn you that we know absolutely jack shit about cricket. Why anyone would spend time watching something that lasts for five days and normally ends in a draw is beyond us.

Anyhow, we thought we'd write the next couple of paragraphs in the persona of two people that love cricket. There's no reason for this beyond the fact that this was entertaining for us. Apologies in advance.

Shirley is an amateur cricket fan. Like any civilised woman with an ounce of cultural decency she's partial to the occasional game in her spare time. Let's face it, who could think of anything better to do than stand on a slightly waterlogged field for hours on end with the occasional walk to pick up a ball. Absolute perfection. We'd challenge anyone to think of a more productive and enjoyable way to spend your Sundays.

Okay, now our Shirley has the power to score a six. For any non-cricket fans, 1) we're impressed you got this far through the book you uncultured philistine, and 2) a six is that thrilling moment when you hit a ball over a boundary rope. It's a feeling of elation that surpasses the birth of your first-born son.

Yes, we did mean to gender that last sentence.

The fact that Shirley's got the power to score a six does beg the question of why she's got this power. What is it about her internal structure that results in this power? Whenever we say something has a particular power, this automatically begs the question of why this is so. In the case of Shirley, the answer might lie in her muscle strength or hand-eye coordination. We're not biologists, so don't ask us for the details. We're just two people with an unrivalled love for cricket.

Now, just because someone has a power doesn't mean that this power is always *exercised*. Shirley has the power to score a six even if she's sat in the audience eating a sandwich with some kind of posh European cheese on it—you know, the kind that has holes in it. Her power doesn't spontaneously come in and out of existence as she steps on and off the field. It's always there, it's just that sometimes she's not using it.

Don't get us wrong, we don't mean Shirley's power is timeless. If she lives to be 100, it's pretty unlikely she'll be able to score a six. This makes sense. Shirley's unlikely to have this power because her internal structure will have changed with age. However, at least in the short-term, Shirley's power to score a six continues to exist even when it isn't being exercised on the field. It's because of the persistence of her internal structure that this power continues to exist. This is what justifies the *applicability* of knowledge.

This example also shows us why knowledge applied in open systems is somewhat unpredictable. The fact that we've identified a power doesn't guarantee predictable outcomes. Sure, Shirley has the power to score a six, but she might be caught out by a fast fielder or be bamboozled by a semi-professional spin bowler whose balletic movement is a joy for all to behold. Truly orgasmic.

That's a sentence we never thought we'd write in a book about research methods, and it's probably a sign we've gone too far in our cricket-loving personas. Let's return to reality.

Now, the drug and cricket examples both summarise the key points of Bhaskar's theory of natural necessity. They show that: a) propertied things and their powers continue to exist even when inactive, and b) propertied things and their powers are not tied to events in predictable ways in open systems. That's Bhaskar's theory of natural necessity in a nutshell.

Bhaskar's theory is really quite a long way from Humean causation. We're not going out and looking for constant conjunctions of events. These regularities are actually pretty rare because the world is full of active things doing stuff, which makes the events we see pretty complex. This messiness can't really be boiled down to natural laws.

Instead, in an open system, if we know a propertied thing is acting, all we can say is that it might *tend* to bring about a particular event. We're looking for these tendencies, rather than predictable natural laws.⁸

So, to sum up, Bhaskar's theory of natural necessity accounts for why knowledge is applicable, albeit unpredictably, outside of scientific experiments. We've seen that experiments can reveal insights into a propertied thing's powers, but in open systems these powers may be unexercised, or exercised but not result in predictable outcomes.

Feature 3: When science identifies a cause, this becomes something to be explained

Onto the third feature of science (hopefully you're getting into the swing of this by now), Bhaskar says that when we uncover a cause this itself becomes something to be explained. We've already seen this in the example of Shirley. Our explanation that she has the power to score a six, automatically led to the question of why she has this power.

Bhaskar himself uses a slightly less cricket-focussed example:

Imagine a scientist is interested in explaining why hydrogen gas is released when you add sodium to hydrochloric acid. The explanation lies in the electronic structure of the atoms involved, i.e. sodium has a greater tendency to give up an electron, enabling it to form a full outer shell, than hydrogen. This explanation about electronic structure of atoms now itself becomes something to be explained. This further explanation might involve some theory of electrons. And then this theory about electrons would also need to be explained through yet another theory, maybe with some talk of further sub-atomic particles.

You can see how this process could continue into deeper and deeper levels of reality with no necessary limit.

Okay, so, if this is a feature of science, what must the world be like to account for this?

The main ontological conclusion that Bhaskar draws from this feature is that the world must have some kind of *depth*. This is sometimes called 'depth realism'. At its most basic, Bhaskar says that the world must be stratified; there must be emergent entities, rather than the world

⁸ Just as a head's up, if you go on to read other critical realist work, the term 'tendency' crops up a lot. One issue is that it's used by different people in different ways. Sometimes it's used to refer to the power itself, like "This object has this tendency", or sometimes it's used to refer to instances when powers are likely to bring about events, like "This power tends to result in this outcome". We think this is confusing because it uses the same word to describe 1) the powers themselves and 2) the empirical results of propertied entities acting. We think it makes more sense to stick to calling things 'propertied entities' and avoid using the term 'tendency' as a synonym. This means that we use tendency as a verb (i.e. this powerful thing *tends* to result in this event) rather than as a noun (i.e. this tendency does this). We hope this side-steps a potential confusion.

being flat. If the world was flat, you'd discover a cause and that would be that. But in our world, science seems to function in a way that has depth: our theories of chemistry are rooted in physics, and physics is rooted in sub-atomic physics.

A supplementary bit of evidence for this stratification is that different academic disciplines (such as physics, chemistry, biology, psychology and social science) more or less correspond to emergent entities at different 'levels'. Each discipline focuses on entities that share some similarity in terms of their level of emergence. To deny this, you'd have to argue that the divisions between disciplinary subject areas are *completely* arbitrary and have absolutely nothing to do with the way the world actually is. That seems implausible to us.

Bhaskar then goes on to add a little more detail to the idea that the world is stratified, saying that emergent entities are 'rooted in, but emergent from' their parts. This means that higher level entities are influenced by the powers of their lower-level parts, but that these higher-level entities possess their own *powers*.

If that sounds like gibberish, then let's think of some examples.

First, both Tom and Cristián can jump. As in like bounce up in the air a bit. If we want to explain the reason why we're able to jump, or maybe even who can jump higher, then this is going to involve some type of *physiological* explanation about our muscle capacity, composition and strength. It might be the ratio of fast to slow twitch fibres or the strength of the muscles that's the key factor, and that's something physiological that can't be reduced to biochemistry or physics.

Now, this doesn't mean that the height we can jump is *unrelated* to biochemistry or physics. The strength of gravity or the stability of our cell membranes are both going to have an impact. Stronger gravity would mean smaller jumps, and less stable cell membranes would mean one jump and splat.

In this way, the physiological explanation is *emergent from but rooted in* biochemistry and physics. But, given we're both human and living on Earth, the key to any explanation of who can jump higher is clearly going to be physiological (e.g. muscle strength and composition) rather than at a deeper level.

A second example, and we're unapologetic about sticking to natural science here, as we'll really bang on about social science in Features 4, 5 and 6, is evolution by natural selection. There's no way that evolution by natural selection (a biological process) can be reduced to chemical principles. For example, natural selection explains how a particular gene, say one that enables increased hand dexterity, can spread within a population of homo sapiens. You can't explain this in terms of chemistry alone. There's nothing chemical about the gene that makes it spread through the population—it's not like it can diffuse better in and out of people—that's really not how human reproduction works. Instead, it's because of the selective advantage the gene provides and the tendency for these individuals to be able to reproduce more successfully that leads to the gene spreading within the population.

Just as with our jumping example, although evolution by natural selection can't be fully explained by chemistry, it does remain *rooted in* chemistry. To give just one example, the chemical structure of DNA influences its stability. If DNA was less stable it would be more susceptible to mutation, and less likely to copy genes from parents to offspring. If we take this to an extreme and make DNA really unstable, this would pretty much make evolution by natural selection impossible, as you'd never be handing over the same genes to your offspring. In that way, evolution by natural selection does depend on a certain level of chemical stability of DNA, but this chemistry doesn't explain all of evolution by natural selection.

If that's making absolutely no sense, we do touch on emergence again in relation to social structures, so that might be more your cup of tea. The key lesson here is just that the world isn't flat. Things have powers, which can't all be attributed to their lower-level parts.

At this point, we could end this feature of science and move onto Feature 4. However, we think that Bhaskar's talk of emergence is a bit vague and leaves a lot of stuff up in the air.

Pun absolutely intended.

We're not alone in this. A lot of people, including some other critical realists like Dave Elder-Vass, have said that although Bhaskar gives us some intuitions about emergence and depth/stratification, we can't really get a systematic account of emergence from RTS. Sure, Bhaskar says things are *rooted in but emergent from*, but this leaves a tonne of important questions like: how and when do things emerge, in what ways are they emergent, and in what way are they still rooted?

Without answers to these questions there's a risk that critical realists use the concepts of emergence and stratification in contradictory ways, and we wouldn't have a way to say which (if any of these) were problematic.

If none of this interests you, then do feel free to skip on through to Feature 4 in Chapter 3B. But if you're up for a bit of emergence chat, then do stick around for the next couple of pages.

Currently, there's different positions on emergence in critical realism. However, the majority of people hold some kind of compositional version. In a nutshell, a compositional account says that a property is emergent if it would not exist without some sort of *relation* between the parts it is composed of.

Let's look at an example.

Schools have the emergent capacity to expel students due to the relations between students, teachers, administrators, norms and rules. If these parts weren't related in a specific way, then the school wouldn't have that power. This doesn't mean the school is a spooky entity with weird and unexplainable causal powers. Quite the opposite. Critical realism actually argues that causal powers, while irreducible to their parts, are explained in terms of their composition and interactions (see also Elder-Vass, 2010). In other words, it's precisely the nature of students, teachers, administrators, norms and rules, and the ways that these are

related to one another, that can help us to understand why a school has the power to expel students.

We believe this compositional theory of emergence has strengths, but also falls short for many interesting phenomena. This is mainly because some things are non-compositional in a strict sense. We're going to make the case that critical realism could learn from a more fleshed out theory of emergence called contextual emergence. This theory has been developed over the last decade or so. If you're interested, some of the big names here are Robert Bishop and Michael Silberstein.

We think contextual emergence is compatible with, but also extends and develops, Bhaskar's conclusions in RTS as well as other critical realists — such as Dave Elder-Vass and Tony Lawson. It gives us more detail and nuance about the type of stratification and emergence that we see in the world.⁹ At its most basic, the idea of contextual emergence is that many phenomena emerge thanks to *contexts*, not only due to their compositional parts.

Contextual emergence basically says that entities emerge when lower-level conditions are combined with higher-level features. The lower-level conditions are the parts and how those parts are related. The higher-level conditions, often called the *concrete context* (hence the name contextual emergence), are the other things needed for an entity to emerge even though these don't *compose* the entity, i.e. they're not a lower-level part.

Let's try to make this more concrete with an example.

Addiction.

Addiction is an emergent state of individuals. It only emerges when there's a series of higher-level conditions that are present. You need higher level features like poverty, family dysfunction, isolation and loneliness, or emotional abuse, rather than simply particular neural relations or responses to dopamine (i.e. the lower-level conditions). It's only when these lower-level conditions are combined with higher-level features that addiction emerges.

Because contextual emergence involves both lower-level conditions *and* context, this means that it's a non-compositional account. Poverty doesn't *compose* addiction in the same way as neural relations or cellular responses to dopamine. This means that addiction, and emergence in general, depends on lower-level conditions (the component parts) *and also* the higher-level conditions (which are not components).


Okay, so we've just thrown some pretty abstract ideas at you.

But this does have some pretty important consequences for the way we view the world.

Contextual emergence, like the standard critical realist worldview, definitely sees the world as having some kind of depth. However, it's a more complex view than the simple

⁹ For a detailed discussion see Navarrete & Fryer (2024).

compositional account of stratification that you tend to see outlined in introductory works on critical realism. See the diagram below for an illustration of this simplified view.



Contextual emergence says this diagram is wrong.

As emergence needs lower-level conditions *and* higher-level context, we can't just view the world as a hierarchy of compositional levels. That would ignore the role of higher-level features in emergence.

Abandoning this simple approach to emergence and stratification has a number of benefits.

First, people often get in a muddle with the levels in the diagram above. Dave Elder-Vass (2005) argues that the levels are a way to group entities with some similarity, but we should remember that it is *things* that are emergent rather than *levels*. In other words, levels are an analytical rather than ontological category. Elder-Vass (2010) also claims that we should picture emergence in terms of a branching tree rather than the lego-like or pyramid-like form depicted in the figure above.

Second, there's some dispute in critical realism about the order that the levels should appear. Andrew Collier (1994) places the psychological above the social, whereas most other people reverse this. We think these debates happen precisely because *context* is key for emergence. It's because many psychological phenomena require both lower and higher-level conditions that makes them difficult to place in a strict pre-defined compositional hierarchy. We saw exactly this with the example of addiction.

Third, there's a whole bunch of stuff that just doesn't fit into this neat and simple diagram. Think of like the environment, or technology. Where do they fit? The simple answer is they don't. That's another bit of evidence that we should give up on the diagram, and the simple version of stratification and emergence that underlies it.

For these various reasons, we think critical realism could learn a lot from contextual emergence. It gives us a robust theory of stratification that better captures the way the world is.

Chapter 3B: Critical realist conclusions about the social world

Feature 4: Intentional action by scientists is needed

The next three features of science enable us to make conclusions about agency, social structure, and their relation. This means we're getting pretty solidly into the conclusions that are most relevant to social science.

But it's probably worth saying from the outset that critical realism isn't going to give you enough detail that you can avoid reading more about structure-agency debates and how this plays out in your particular field.

Sorry, that's just some reading you'll need to do.

But critical realism will help you navigate these debates, so don't be too disheartened.

Anyhow, back to what critical realism has to say. The fourth feature of scientific experiment that Bhaskar picks out in RTS is that experiments rely on *intentional action by scientists*. Nature doesn't really present us with closed environments, or at least it doesn't do this all that often, and that's the reason why researchers have to work to create experimental conditions.

Think about something as extreme as the Large Hadron Collider, that huge 27-kilometre tunnel that cost over £6 billion to study sub-atomic particles. This is a pretty cut and dry case of intentional action. People wanted to build a big ass tunnel, they got together and got funding to build a big ass tunnel, and now they use this big ass tunnel for science.

There's just no way to explain the building of such a humongous structure without accepting some type of intentional action. The tunnel didn't just spring into being; it was designed, built and developed by people who acted (at least partially) in intentional ways.

So, what must the world be like for this to be a feature of science?

At its most simple, this feature of science is only possible if some type of individual agency is an ontological characteristic of the world. Humans must have the capability to act in intentional ways.

This has important consequences for those of us working in applied social science. It means that our social theories must allow for (at least the possibility of) human agency.

Again, don't get us wrong, critical realism isn't going to tell you how or to what extent, human agency manifests itself in your specific research area. Again, sorry, that's just some reading and/or research you need to do. But critical realism will give you a solid starting point: *you should assume that people have agency and that they're able to undertake intentional action, unless you have a good reason to think this agency is inhibited or blocked by something.*

So, because agency is a characteristic of the world, if you want to make an argument that doesn't involve human agency, the burden of proof is on you to show why intentional action

has been inhibited or blocked by some other factor. In the absence of this, your explanation should involve some reference to people's intentions and reasons for their action.

Bhaskar doesn't stop here. There's more he has to say about agency and intentional action.

First, he explains that intentional action is only possible if *reasons can be causes*. In other words, intentional action only makes sense if people can act on the basis of reasons.

What's an example of this?

Example 1. Alex goes to get an ice-cream. Alex does this because they think ice-cream is delicious and they're just kinda in the mood for a sugar-fix.

Example 2. Alex chooses to take a particular job because they'll be able to make a contribution to society. Alex might reason that they should take the job, even though the work is poorly paid, as they have the chance to make a difference, which outweighs the low salary.

In both examples, Alex acts on the basis of reasons. Alex's reasons are a cause of their action.

Now, this definitely doesn't mean that we always act on the basis of reasons. Alex (as with any mortal) might open up the fridge and think "Why the hell have I just opened this... what was I looking for?".

It also doesn't mean that the reasons for our action are always true. There's all kinds of reasons for action that might turn out to be false. Alex might think: "I'll invest in that pyramid scheme to get rich". Or, if Alex is a fan of certain tabloids, they might think: "I read in the newspaper that wine causes cancer, best not to do any drinking". And a week later: "I read in the newspaper wine helps avoid cancer, best get bladdered".

Reasons can be wrong, contradictory, or absent, but still serve as causes of our action. For our purposes, what matters is only that reasons can be causes of our action.

A second, and more specific thing, that Bhaskar adds to his theory is that agency is an emergent power of individuals.

Oh damn it, back to emergence again.

All that Bhaskar means by this is that *agency can't be analysed in purely neurological terms*. Sure thing, agency might be rooted in the neurological processes of our brain, but agency is emergent from all these neurons firing and whatnot.

Think about it this way, when Alex acted on the basis of a reason (e.g. taking a job because of the contribution they could make to society), this reason involves *meaning*. It involves the meaning of the phrase 'contribution to society' and Alex's specific take on and valuation of this. Now, all of this is a social phenomenon, influenced by Alex's upbringing, education and life experiences, and there's no way this can be reduced to neurons firing.

Now, it's fair to say Bhaskar is pretty vague on the way that agency is emergent. In part, this is understandable. It's the role of science (especially a combination of neuroscience, biology and psychology) to explain the exact way in which individual agency emerges from neurological phenomena. This isn't something we're able to work out by logical reasoning, so it's a scientific question, rather than a philosophical one. That's why it's fine that Bhaskar's a bit vague on some of the details.

But, to sum up, what Bhaskar is able to say is that agency is emergent, that it can't be reduced to biology, and that this is a legitimate and necessary concept to use in our social science theorising. Unless you've got a good reason to ignore agency in your specific context, it's essential that this features as part of your social science explanations.

Seeing as we've just defended Bhaskar, let's slightly change tack and say there's definitely a bunch of ways his theory of agency could be developed.

We're actually going to try and do some of this development in Features 5 and 6 below.

In this, we build on the work of Margaret Archer (2000), who reached the same conclusion that agency is an emergent power. However, as a sociologist, Archer offers loads more nuance about how agency functions, and how it links to social structure, than we find in RTS.

Now, it's true that a lot of Archer's theory is derived using evidence and data, making it more social science than philosophy. But we think there's aspects of Archer's work that can be derived through the same abductive processes Bhaskar uses in RTS.

Specifically, we think we can use two new features of science, which Bhaskar does not discuss in RTS, in a way that allows us to develop a more robust theory of structure, agency and their relation. We see this as a way to extract the philosophical elements of Archer's theorising and use them to complement the conclusions found within RTS.

Feature 5: Scientific experiments are impacted by resources

The fifth feature of science, as identified by us (so put on your sceptical hats), is that scientific experiments are impacted by resources.

What does this feature of science mean?

It means that the actions of scientists are influenced by the resources they have access to. To make an obvious point, scientific experiments happen in the world, they happen in existing organisations, and these organisations have different levels of access to resources. If we compare a physics laboratory at an elite university in the UK to a less affluent one in Bangladesh then it's likely these labs are going to have access to pretty different resources, whether financial or human.

This different access to resources could have a bunch of impacts on the work that happens in the labs. We're just brainstorming here, but it's easy to imagine that:

- In the Bangladeshi lab, some experiments will be impossible because the equipment is too expensive. This clearly limits the type of knowledge that can be produced.

- Even if an experiment isn't so expensive that it's impossible, an expensive experiment is *riskier* in the Bangladeshi laboratory than the UK one. The Bangladeshi scientists might be less able to risk large amounts of capital on an experiment that has little guarantee of producing useful results. This risk may discourage the Bangladeshi lab from doing risky work, which again impacts knowledge production.
- Also, the UK laboratory might be able to offer compensation packages to incentivise people to come and work there, but the Bangladeshi lab might be less able to do this. Again, this has implications for the knowledge produced.

Resources impact knowledge production.

So, back to our standard abductive question, if this is a feature of science, what must the world be like?

This feature of science only makes sense if our actions take place in contexts that are causally impacted by social structures. Put another way, there's no way of understanding the differences in research output between the two laboratories in psychological terms alone. It's not the psychology of the Bangladeshi researchers that prevents them from under-taking super expensive experiments. Instead, there's an objective constraint that influences and restricts their action. This constraint is the causal impact of a social structure in the world. In this way, social structures set the *contexts* in which people act, making some actions more or less possible, more or less risky, more or less likely to succeed.

Now, we know that the term 'social structure' is thrown around a lot and it's not always clear what people are jabbering on about. So, what can critical realist philosophy say about social structures?

First, social structures are real. We can't understand different knowledge production in the two labs in psychological terms alone, so there must be *something* that's real and impacting the situation.

Second, these social structures involve a relation. It's a relation between a person (or a group) and something else (like, resources or ideas). In the case of the labs, it's the relationship between the scientists and resources.

Some people are a tad sceptical that relations are real, so let's very quickly justify their existence. Think about it this way, the relationship between the two physics laboratories (e.g. their access to resources) exists independently of our knowledge of this inequality. It's not as if our knowledge of this inequality causes the inequality to exist—it exists whether or not we know it.

Third, social structures set the context in which people act. They make some actions more or less risky, more or less likely to lead to the desired outcome.

Beyond these three basic aspects of social structure, we're starting to move onto a more social scientific theory that's based on data and evidence rather than philosophical abductive reasoning.

The three points above outline a ‘minimal theory of structure’, which is all that’s really possible using the abductive method in RTS.

But, before we move on, it’s worth quickly flagging one influential definition of social structure within critical realist work. As much as anything we’re doing this to point you in the direction of other work that’s happening and where you could go next with your reading.

Here we’re talking about the work of Doug Porpora (2015), who’s a leading critical realist, and author of a great book called *Reconstructing Sociology*. We’d definitely recommend this as follow-on reading.

Porpora defines a social structure as: ‘(material) relations between social positions and social constructs’.

In this definition you’ll see *relations* play a key role. According to Porpora, a social structure is just a relation involving social positions and constructs. The researcher in the Bangladeshi lab (i.e. an actor with a social position) has access to a particular level of financial resources (i.e. one type of social construct), and this may be different to a researcher in a UK lab (i.e. an actor with a different social position).

There’s obviously a lot more to be unpacked around what a social position is, what a social construct means, and what makes the relation material or ideal, but that’s a job for Porpora. Apologies, and happy future reading.

It’s also worth a quick comment to say that we’re glossing over *culture*. Porpora builds on Margaret Archer’s (1996, 2000, 2003) work in this area to say that we can have cultural structures, as well as social structures. A cultural structure is the relationship between ideas, but there’s an active debate in this area, with a challenge to this conceptualisation by Dave Elder-Vass (2012). If you’re interested in these questions, *Reconceptualising Sociology* by Porpora (2015) and *The Causal Power of Social Structures* by Elder-Vass (2010) are great resources.

Feature 6: In turn, resources are impacted by knowledge production

Okay, so far, we’ve made the case that intentional agency is possible (Feature 4) and social structures are real and have causal impacts on the world (Feature 5).

With the sixth (and final, thank God) feature of science we can begin to think about the *relationship* between agency and structure. You can think about Features 4 and 5 as showing agency and structure are ontological characteristics of the world, and the sixth feature as beginning to sketch out some of the ways they’re related.

Okay so the sixth feature of science is that not only is knowledge impacted by resources, but *resources are in turn impacted by knowledge production*. We’ve already made the case that knowledge production is impacted by resources (this is Feature 5), but here we’re arguing that over time the knowledge produced also impacts access to resources.

Let's revisit our Bangladeshi lab. Imagine that the scientists decide to undertake less risky research, which then yields little ground-breaking work. When these scientists apply for future funding, they're at a disadvantage compared with researchers at wealthier labs that have been able to demonstrate a track record of innovation. This means that by undertaking less risky experiments the Bangladeshi scientists are then less able to gain access to resources in the future, partly as a result of the outputs they achieved in the past. In this way, the social structure is reproduced over time through the actions of the individual scientists.

Now, it doesn't have to be this way. The Bangladeshi lab is not guaranteed to stay in this disadvantaged state. It's possible that the Bangladeshi lab is able to use its resources to produce world-leading research. The reputational gains from this work would then increase the opportunity for the lab to gain future funding. In this case, the action of scientists would have led to a change in the social structure—the social structure is modified, at least to an extent, as the lab changed their level of access to resources.

This is all intended to show that the knowledge produced in the labs impacts their future access to resources, and this happens in a way that either reproduces or transforms the social structure.

So, for the sixth and final time, we can ask the abductive question: what must the world be like for this to be a feature of science?

We think that it's reasonable to conclude that this feature of science only makes sense if there's a relationship between agency and structure that broadly resembles the three-step process outlined by Margaret Archer (2003). These three steps are:

- Social structures shape the contexts in which people act. These contexts serve to constrain and enable certain forms of action.
- When people come to act in a given context, they have a particular set of concerns or things that they value.
- People then act in these contexts, and these actions either reproduce or transform their social position and relations, as well as potentially impacting wider social structures.

You can clearly see that social structures and agency exist in a dynamic relation. Social structures set the context in which people act, people then act, and this action leads to the reproduction or transformation of social structures. This in turn conditions future action, and the cycle repeats again and again and again.

Although this three-step model might appear relatively simple there's a bunch of stuff that it helps to clarify when it comes to our practical social science work. We'll explore this more in the next chapter, but there's a broader point about causation we can make. And this point returns us to Features 1 and 2, and the conclusions about natural necessity.

The eagle-eyed among you may have noticed that when we've spoken about the causal impact of social structures, we've used the terms 'constraints' and 'enablements'. In this we're following Archer. These terms stress that the existence of a social structure doesn't

guarantee or fully determine a particular outcome. This is because people have some agency to decide what to do in their context.

In our example, we saw how the scientists in the Bangladeshi laboratory can choose to take a risk and undertake an expensive experiment (assuming we're not in the extreme case where it's simply too expensive to perform the experiment), even if a social structure means that this action is particularly risky. The social structure determines how risky the action is for the lab, but it doesn't force the scientists into performing a certain type of experiment. They could choose to take the risk.

This is all pretty similar to what we were talking about in Feature 2. Remember, propertied things only *tend* to bring about events, rather than guaranteeing outcomes. It's pretty much the same here.

The only difference is that we're using the terms 'constraints' and 'enablements', to emphasise the role of agency. Often people have some choice in what to do, in how to bring about what they value, and in how much they are willing to sacrifice to try and bring about this goal. It's this intentional action, this reflexivity, which is present in the social sciences (but not, or at least less so, in the natural sciences) that makes us use slightly different terms. But, in both cases, the presence of a cause doesn't necessarily guarantee a particular event.

So, to sum up, social structures set the context in which people act, opening some doors and making others more challenging to open. But, it's up to actors to decide what they actually want to *try* to do in this context. And this action, in turn, has the potential to impact social structures, either reproducing or transforming them.

As we've mentioned, Archer offers tonnes more detail about the relationship between structure and agency. For example, she describes the ways different people develop different *concerns* about the world, linking this with their mode of reflexivity. However, much of this comes from empirical research (largely from in-depth interviews), and so moves beyond the 'minimal theory of agency' that we're looking to outline in this book.

Again, that's another one for your reading list. We'd particularly recommend Margaret Archer's (2003) book *Structure, Agency and the Internal Conversation*.

Okay, let's pause for a minute to take a breather as that's been quite the journey through research philosophy.

Let's try to pull all of this together and summarise where we've got to. The table below tries to map out the six features of science, and the linked ontological conclusions.

Basic ontological characteristics of the world, and the relevant features of science

Feature of scientific activity	Ontological conclusions about the world
1. Experiments are important to science.	• Experiments are <i>closed</i>, and the world is <i>open</i>. • The world is structured. It contains experiences, events and causes. • We should explore the causes that underlie events.

2. Knowledge is applied outside of experimental situations.	• Causes only <i>tend</i> to bring about particular events outside of experimental conditions. • Theory of natural necessity: propertied things have powers and act to bring about events.
3. When science identifies a cause, this itself becomes something to be explained	• The world is stratified. It has emergence and depth, albeit complex.
4. Intentional actions by scientists are needed	• Agency is possible, as humans can act on the basis of reasons.
5. Scientific experiments are impacted by resources	• Social structures are real.
6. In turn, resources are impacted by knowledge production	• Social structures constrain and enable people's action. • Social structures and agency interact: ○ Structures set the context. ○ People then act in these contexts on the basis of their concerns. ○ This action reproduces or transforms the structures, conditioning future action.

Other stuff you'll see in critical realism

There's also a bunch of other stuff that you might come across in critical realism that we've not included in the table above.

As this is a weird hybrid between an introductory text and one that tackles RTS in particular, we thought that it was worth ending this chapter by mentioning some of these other ideas. Although some of these concepts are widely used in critical realist work, we either: 1) quibble with their exact formulation (e.g. the three domains of reality), or 2) don't think they've got huge implications for our work (e.g. the intransitive/transitive distinction).

For those of you that aren't planning to read more about critical realism, feel free to skip this section. But if you're planning to engage with critical realist work, then here's some of the other common ideas and concepts that people use, and our take on them.

We'll talk about:

- 1) the three domains of reality – adapting a piece that we wrote on this topic.¹⁰
- 2) The idea of transitive and intransitive objects of science.

Three domains of reality

Now, you might have heard critical realism talk about the three domains of reality. In some senses, this is one of critical realism's most central conclusions. It's one of the main conclusions that Bhaskar highlights in his intro to RTS, and it's also one of the key ideas that's been picked up and applied in critical realist work.

¹⁰ The full text of this can be found: Elder-Vass, D., Fryer, T., Groff, R. P., Navarrete, C., & Nellhaus, T. (2023). Does critical realism need the concept of three domains of reality? A roundtable. *Journal of Critical Realism*, 22(2), 222–239. <https://doi.org/10.1080/14767430.2023.2180965>

The three domains are often described in the table below.

You'll see they're linked to the concepts of experiences, events and causes, that we've already discussed in this book.

Defining the three domains of reality, adapted from Bhaskar (2008)

	Domain of real	Domain of actual	Domain of empirical
Causes	✓		
Events	✓	✓	
Experiences	✓	✓	✓

Essentially the domain of the real contains all three, the actual contains events and experiences, and the empirical contains experiences.

Now, even though you see the three domains of reality mentioned all the time in critical realist inspired social science, we actually think we'd be in a better place if we stopped talking about the three domains of reality. We make this argument because we think the three domains of reality are redundant and somewhat confusing.

Argument 1: The three domains of reality are redundant

If you already have the concepts of *experiences* (the perceptions of things by agents), *events* (the things that are being perceived) and *causes* (the things that tend to produce the events),¹¹ then the three domains of reality are redundant. They don't add much to our understanding of the world.

If you know the difference between experiences, events and causes then you're already in a pretty good position to go and do research in a vaguely coherent way. The concepts allow you to understand what your data is (often agents' experiences of events), and what you should be aiming to do (often offering explanations of how causes bring about particular events and experiences). It stops you falling into the trap of focusing on what you can measure and observe (events and experiences) and keeps us on the task of finding the underlying causes.

¹¹ Our argument should not be interpreted as a claim that reality only consists of experiences, events and causal mechanisms, or the claim that these are the only ontological concepts we need. A robust ontology will involve several other concepts, such as *powers* or *emergence*, to name but two.

For a researcher that understands the differences between *experiences*, *events* and *causes* we can't see any extra insights that come from adding the three domains of reality. They're redundant.

Argument 2: The three domains of reality are too confusing

It's also much simpler to just talk about experiences, events and causes rather than the three domains of reality. After all, you can't explain the three domains of reality without introducing experiences, events and causes—that's exactly what the table above is trying to do.

This is anecdotal, but in our experience, people tend to find experiences, events and causes reasonably easy to grasp. The idea that there are *experiences* (the perceptions of things by agents), *events* (the things that are being perceived) and *causes* (the things that tend to produce the events) is often accepted as pretty intuitive.

However, it's the opposite for the three domains of reality. Often when we speak to people about critical realism, they mention the three domains as something they're struggling to grasp. This conceptual challenge applies to both the nature of the domains ("Why are the domains overlapping?"), as well as how to apply this in their research ("How do I know which domain I'm looking at? If I interview someone about their experience, is this in the empirical, the actual or the real?").

Maybe the people we speak to are confused because we're not keen on the three domains of reality and do a bad job of explaining them. But maybe, just maybe, these people's intuitive dislike of the domains is telling. Maybe the three domains of reality really are too confusing and ultimately add little to our understanding.

If the domains of reality are both redundant and confusing, and if people find experiences, events and causes relatively straight forward, then critical realists should stop talking about the three domains. We should always strive to give the most accessible accounts of critical realism, enabling more people to benefit from its insights. Not to do this just puts up barriers. It's gatekeeping.

Maybe that's a bit harsh. We know there's plenty of people that find the three domains a useful way to grasp some of the points Bhaskar is making in RTS. We'll definitely concede that the three domains can be a useful short-hand for when critical realists talk to one another—but we do think that there's simpler and more accessible ways of talking through these issues.

Again, we've included this section only because many critical realists do talk about the three domains, so it felt wrong to ignore them completely.

Transitive and intransitive

Another set of concepts you'll see in critical realism are the transitive and intransitive objects of science. We see this as less important, so here's a brief overview, in case this crops up:

Bhaskar argues there are two different objects of scientific practice:

1. Intransitive object – these are the *real objects* that exist independently of the researcher, e.g. the real electron structure of sodium.
2. Transitive object – these are the *scientific theories* that describe the intransitive objects. Scientists take past scientific theories and extend, refine or replace them in ways that better correspond to the intransitive objects of the world.

That's pretty much it.

Intransitive object = stuff in the world; transitive object = our theories about this stuff.

Don't be too put off by the jargon.

Chapter 4: What are the implications for your work?

Okay, first of all, well done for getting through that mammoth discussion in the previous chapter.

It took us quite a while to explore the various abductive questions that feature in RTS, but we've been left with a relatively detailed set of ontological conclusions about the world.

At this point, you might well be thinking that it's all well and good having these conclusions, but what impact does this actually have on your research?

That's exactly the question we aim to tackle in this chapter.

Just to whet your appetite for the methodological feast that's to come, a phrase that no-one has ever used before, here's a quick summary of some of the main ways these critical realist conclusions influence your research design:

- It influences what you seek and helps you to avoid looking for things that don't exist. For example, you should ask research questions that seek detailed descriptions and nuanced causal explanations, rather than ones that seek natural laws.
- It influences how you see the strengths and limitations of quantitative and qualitative research. This helps you perform a literature review and evaluate what we know on the basis of previous work. It also helps you design your own research, adopting methods that are consistent with critical realism and able to answer your questions.
- It influences the theories you use. Some theories make problematic philosophical assumptions, and so you should reject (or adapt) these.
- And finally, it influences the conclusions you make, helping you think through how transferable your findings are.

That pretty much covers all research design, which hopefully gives you a sense of critical realism's potential. If you want more, a recent paper by Adam Nichol, Catherine Hastings and Dave Elder-Vass (2023) goes into these links in more detail.

Before we get cracking, a couple of quick caveats. One positive, one negative (ish).

First, it's worth stressing that critical realism has *real implications* for your research.

By this, we mean that you can't just pick and choose whether you want to follow critical realism's conclusions or not. Bhaskar's abductive arguments give us a really powerful method to develop a theory about the basic ontological characteristics of the world. It's the power of this method that underlies the confidence we can have in critical realism's conclusions, and the reason why we *should* take these conclusions seriously and adapt our research to be compatible.

Sure thing, we're not saying: "Follow us, we've got the perfect theory". That's one of the beauties of critical realism. Yes, there's ways the abductive reasoning could be dodgy, but this means that you can get involved to develop, refine, reverse or update the conclusions.

A second caveat.

We don't want to over-egg the potential of critical realism. Let's not set our expectations too high for what it brings to our research.

Yes, it's powerful and it has implications for pretty much all areas of research design, but it isn't going to tell you exactly what to do.

It'll definitely help you to think through some tricky decisions, and help you narrow down the range of possible solutions, but this won't lead you to the one perfect research design. Instead, it's on you to make decisions about what's best for you and your project, with some guidance from critical realism.

Think about it this way, if critical realism was a theory of cooking, it would tell you: 1) don't put your hand on the stove, and 2) food tends to taste good when it balances salt, fat and acid.

Now, this is all well and good, but this 'theory of cooking' isn't going to tell you what to eat tonight, and it's not going to tell you how to make a great spaghetti bolognese. That's on you.

So, you get a pretty good helping hand from critical realism, but that's all it is, a helping hand.

We've tried to organise this chapter in a way that builds on Chapter 3, but equally offers clear advice about methodology. To do this, we decided to break things into two sections.

The first looks at the conclusions from the last chapter, and their implications for research. The second combines all these implications together and thinks more broadly about how to undertake critical realist work using different methodologies (i.e. quantitative, qualitative, and mixed methods). You can think of the first section as giving the concrete and specific implications, and the second as the broader insights for research design.

Here's the structure we'll follow:

- Section A: Concrete implications
 - Feature 1: Experiences, events and causes
 - Feature 2: Powers and necessity
 - Feature 3: Stratification
 - Feature 4: Agency
 - Features 5 and 6: Social structure, and the relation with agency.
 - Bonus: intransitive and transitive objects
- Section B: Implications for methodology
 - Implications for quantitative research
 - Implications for qualitative research
 - Implications for mixed methods research

Section A: Concrete implications

Feature 1: Experiences, events and causes

The conclusion that the world is structured into experiences, events and causes is likely to have the biggest contribution to your research design.

It really can be revolutionary.

It's got the potential to bring clarity, or at least the beginnings of clarity, to any muddled field of study.

To be more specific, the three concepts help you understand what your data is (often agents' experiences of events or measurements of events), and they help you to avoid falling into the trap of focusing only on what you can measure (events and experiences) rather than the underlying causes.

With this alone, you're already in a pretty good position to go and do research in a vaguely coherent way.

Let's share an example from our own research (in this case Tom's work).

I'm interested in the impact of higher education on graduates. Or at least that's what I did during my PhD. In reality, I like dabbling in all kinds of different areas related to educational inequality in a way that doesn't form an especially coherent research programme. But I digress already. So, after wading through RTS, failing, getting some help from Andrew Collier, and then going back to RTS again, I came to realise that my whole research area was muddled.

It was muddled mainly because people confused events and causes. Sure, this wasn't the only problem, but this mixing up of events and causes made everything infinitely more complicated.

To be more specific, both policymakers and researchers often looked at things like 'Do graduates have a job?' or 'What type of salary do graduates get?' and called these 'graduate outcomes'. They even used these things to make conclusions about the 'quality' or 'value' of higher education.

But, whether or not a graduate has a job is an event. Like many events in open systems, they're caused by a huge range of things, like: where the graduate lives, the discrimination they face in the labour market, prior social connections, the symbolic value of their degree, their employability etc etc etc.

Critical realism helped me to see that these so-called 'graduate outcomes' were actually complex events caused by a bunch of different stuff, many of which have diddly squat to do with higher education and its quality.

This means you can't just use these events to make conclusions about the causal impact of higher education.

Instead, if we're interested in higher education and its impact, we need to look beyond events and consider causes. And, more specifically, we need to unpack the causal impact(s) of higher education on these events, and how this varies for different students, in different universities, with different experiences of higher education.

Now, although this might seem obvious, for me, this was an absolute game changer. The three concepts allowed me to put my finger on exactly what had been bothering me about my research area (mixing up events and causes) and offered me a way to get out of this trap.

Returning back to critical realism's implications for your research, the three concepts also help clarify our research aims.

In particular, critical realism helps us see that causal knowledge requires us to go beyond a focus on events and experiences to consider the underlying causes. This means our projects are justified in seeking to develop explanations of how particular causes act to bring about the events we're interested in. Critical realism justifies asking questions that explore underlying causes.

Take the example of gender and maths scores. It might make sense for some research questions to look at events and experiences (e.g. How do girls experience maths classrooms in English state schools?). However, unless there's a reason to stop here, critical realism nudges us to go further and explore the relevant causes. This might mean a question like: 'What are the causes that influence maths scores, and how do these causal mechanisms operate?'. Or if that's too broad, we could look at the impact of one of these causes, say the impact of gender identity on teacher-student interactions and the link with maths attainment.

So, we've seen critical realism justifies and defends the use of causal research questions, but you might also see some people claim that *the ultimate aim of research is to produce causal explanations*. The line of reasoning goes: Bhaskar showed us that it's possible to move beyond a focus on events to understand underlying causes, so this must be our *ultimate aim*.

Now, we don't 100% agree, even though we've come out with similar statements in the past. Guilty as charged.

That being said, we don't 100% disagree either.

We're completely on board with idea that Bhaskar shows us the importance of causal knowledge, and the ways this focuses on underlying causes rather than events.

Our quibble is whether this is the *ultimate aim* of research.

First, we think there's plenty of times when causal knowledge isn't the ultimate aim of a particular research project.

Take the example of an under-explored research area. In this case, it might make sense to focus on describing people's experiences or describing the events that happen. This might result in nuanced and detailed work that tries to understand the new area, and maybe even begins to identify some interesting patterns or features.

As a general rule, we need a decent understanding of phenomena, before starting thinking about underlying causes. Plus, if you've got limited time and budget, then it's totally fine for research not to focus on producing causal knowledge.

But it's not just in under-explored research areas that description can be important. In contexts where there's lots of evidence, the main aim of a research project may be to re-frame the problem. Again, this type of work wouldn't have causal knowledge as its ultimate aim.

Take the example of international student experiences in UK higher education. There's tonnes of research here, but a lot of work (especially the older stuff) was framed around a deficit approach where international students were seen as lacking something (lacking skills, lacking cultural knowledge etc). Now, it's an important contribution to come along and say: "Hang on a minute, is this deficit framing really the best and most ethical way to talk about international students? Does this not just place all of the burden on international students rather than our own institutions, programmes and pedagogies?". A project might then aim to re-describe and re-frame the research area, moving away from a deficit approach and posing a new set of questions for future work.

So, at least for individual research projects, we think that causal knowledge doesn't have to be the ultimate endpoint.

However, there's also a second reason why we quibble with the idea that *causal knowledge is the ultimate end point of research*.

We're not sure that the simple binary between *descriptive* and *causal* knowledge holds.

Sure, there's a difference, but there's also areas of overlap.

A critical realist view of causal knowledge emphasises that knowledge should outline how a cause operates to bring about events. But this outline involves description: we describe how the cause operates, and we describe how the events are related. In this way, we reckon critical realism actually makes a strong case for valuing description in research—this description is essential when producing explanations about causal mechanisms. We're not alone in making this argument, Doug Porpora (2015) has said much the same. So, any talk of causal knowledge shouldn't be seen as downgrading or belittling descriptive work; they are deeply complementary, especially when we accept non-Humean forms of causation.

All that being said, there's no denying that in a lot of research contexts, there's very good reason to ask causal research questions. Critical realism does justify this focus on understanding underlying causes rather than just focussing on event regularities.

It can also make sense to focus on producing causal knowledge because: 1) this is the part of research that's often most transferable and applicable, and 2) if we want our research to have an impact in the world, contributing to flourishing (unless you're a heartless bastard), then it's often causal knowledge that we need. To promote flourishing we need to know the barriers people face, as well as the things that could facilitate this flourishing.

Again, remember that all this causal knowledge builds and relies on a deep understanding of the research area from prior descriptive work.

Okay, we've jabbered on for a while, so let's state this clearly and make your note-taking a bit easier:

"Causal knowledge can be an important research aim, and critical realism helps to justify this. However, descriptive work remains vital, a valid endpoint in its own right, and a core component of any causal work."

That's our take on the idea that the ultimate aim of research is to produce causal knowledge. And to be honest, we think many critical realists would agree.

To sum up, from this ontological conclusion we've seen: 1) the clarity that you can get from applying the concepts of events and causes, 2) the nudge to ask a causal research question in your work, and 3) the value of description.

Feature 2: Powers and necessity

If the three concepts of experiences, events and causal mechanisms bring clarity, then Bhaskar's theory of natural necessity turns this into super-duper clarity.

Academic prose at its finest.

Okay, now remember Bhaskar's theory of natural necessity is all about critical realism's approach to causation. The most basic implication for your research is that you should aim to produce causal knowledge that's compatible with this theory.

This helps you avoid using theories of causation, such as a naïve Humean approach, which sets you up to fail from the start. Universal natural laws don't exist in social science, so there's no point setting out to find them.

For example, there's simply no such thing as the one universal impact of gender on maths scores. It's a mistake to ask this question in the first place.

Now, don't get us wrong, we're not saying that there's no point in looking for gender variations in maths attainment and how this varies for different sub-populations. Quite the opposite, we think this makes a lot of sense. This attempt to describe patterns in events is often an essential starting point for research. But this focus on *detecting patterns of events* isn't the same as seeking the *universal impact of gender on maths scores*. Although the methods might be similar, the understanding of what the research is doing is radically different.

We'll talk more about quantitative research methods in the next section, so stay tuned.

So, rather than using Humean approaches, critical realism shows you should adopt an approach that unpacks the causes underlying events. This aims to understand why we see a given pattern of events, rather than just describing the patterns themselves.

Again, we'll discuss this later, as we talk about the ways qualitative and mixed methods approaches can be particularly powerful ways to produce this causal knowledge. That's all coming up below.

Feature 3: The world is stratified, albeit in complex ways

Bhaskar makes the case that the world is stratified, with entities emergent from but rooted in their lower-level parts. We built on this to argue contextual emergence could add to and build on Bhaskar's initial ideas. This theory maintains a sense of depth, but results in a more complex version of stratification.

Now, even if you don't buy the ideas of contextual emergence (which just so you know deeply hurts, maybe we need to go watch some cricket to feel better about our lives), the main takeaways are the same.

First, you'll be glad to hear that social science is a justified field of study.

Hurray. Let's raise a glass to that.

It's because our world has emergence, and specifically because social structures are emergent from individuals, that we can't just use neuroscience to understand everything in the social world. Sure, brains are important (remember, *emergent from but rooted in*), but neuroscience won't tell us why someone chooses to go to university or not. We can only explain this if we understand: 1) the objective constraints people face in attending university, and 2) peoples' own concerns and aims.

Emergence allows us to rebut claims like: "To understand society, we only need to look at individual behaviour" or "The social world is just neurons firing".

Because critical realism shows us the world consists of real things with real emergent powers, we can't just focus on the lower-level parts and treat the higher-level as redundant.

The fact that the world is stratified has a second implication, namely that interdisciplinary work is important in many contexts. Lots of events result from complex interactions between propertied things at different levels of emergence.

Take the example of the climate crisis. This is impacted by things such as individual human behaviour, national governments, and industry, as well as the chemical properties of carbon dioxide, the stability of peat fields, and the fluctuations of weather systems, all of which operate at wildly different levels of emergence and require very different methods.

In a research area like this, interdisciplinary work is incredibly valuable.

Feature 4: Agency and intentional action

We live in a world where intentional action and individual agency is possible, as Bhaskar showed through his fourth feature of science.

A kinda obvious consequence is that the theories and descriptions you use should allow for this. You shouldn't use a theory (or an interpretation of a theory) that ignores agency, unless you've got a really good reason.

This mistake does crop up from time to time. We're thinking of some research that applies Bourdieu's concept of habitus but makes this so concrete that people are stripped of their agency and unable to undertake intentional action.

There's also some slightly less obvious implications.

First, the fact that agency is possible is one reason why we need to take people's perspectives and worldviews seriously. Because people can act on the basis of reasons, critical realist social science needs to listen to people. If we ignore people, and just look at structural factors, perhaps in some mistaken attempt to be objective, then we miss a key part of the explanation of why people do what they do. (Unless of course you're studying some real dystopian shit in which people have pretty much no agency).

In this, critical realism emphasises the value and necessity of understanding people's experiences and worldviews. These are absolutely central to any social science explanation.

A second, and somewhat similar implication, is that people's understandings of the world can be a hugely important resource for social science. People's ideas often take the form of causal statements, like: "I think that my rising energy bill is caused by X". Bhaskar calls these ideas 'proto-scientific theories'.

According to Bhaskar these proto-scientific theories give our research a useful starting point. They can give us the first ideas about what might explain a phenomenon.

For example, when Alex chooses which job to take, they have some understanding of the objective constraints on their decision (e.g. will they be able to pay their bills) and other causal factors (e.g. are they likely to suffer socially if they take this job). There's no guarantee that Alex is aware of all the causes of their action, nor that the things they believe are true, but this is a pretty good starting point.

For fear of stating the obvious, this is a uniquely social science phenomenon. There's no equivalent in the natural sciences. It's not like sodium atoms understand why they tend to give up electrons.

As we said, these proto-scientific theories could be wrong. That's why they're only a possible starting point for social science, rather than an end. You don't have to accept people's arguments as true, instead you should subject them to scrutiny and critique.

Related to this, a third implication is that if social science can take peoples' understandings and show some of them to be mistaken, this makes social science emancipatory. It's emancipatory in the sense that it helps us move beyond *false belief*.

Andrew Collier (1994) sums this argument up as:

"Given that (other things being equal) it is better to believe what is true than what is false, it is also better (other things being equal) that institutions that cause false beliefs should be replaced by, or transformed into, those that cause true ones." (p.172)

Let's make this a bit more concrete with an example.

Take the idea that contemporary *British society is meritocratic*. Social science has revealed the various ways that success is influenced by social background rather than merit. This means that someone justifying a particular privilege (say a tax break), or justifying the lack of welfare for others, using a purely meritocratic argument is false, e.g. "They don't deserve that because we live in a meritocracy". Because social science is able to show this idea to be false, this provides support for transforming any systems that are premised on this mistaken assumption.

Now, of course, not all research needs to focus on political topics and false belief. But if your work does, critical realism gives you an argument for why this is an entirely appropriate role for social science. Don't let anyone tell you that research has no relevance to politics.

Features 5 and 6: Social structure and agency exist relationally

Okay, for the final set of implications, we're going to combine the insights that come from Features 5 and 6. It's from these two features that we worked out social structures are real, and that they exist in a relational way with agency.

This relation was all captured by Archer's (2003) three-step process:

- Social structures shape the contexts in which people act. These contexts serve to constrain and enable certain forms of action.
- When people come to act in a given context, they have a particular set of concerns or things that they value.
- People's action can either reproduce or transform their social position and relations, as well as potentially impacting wider social structures.

Again, the most basic implication is that social science should consider agency, social structures and their relation. Sometimes this might involve a really concrete application of Archer's three steps, and in others this might be more implicit, say by ensuring you stress both social structures and agency in your descriptions (we'll give an example of this in a sec).

Some social theories don't do this, either leaving no space for agency (e.g. some interpretations of Durkheim) or ignoring social structures (e.g. some versions of rational

choice theory). You shouldn't use these theories, or at least you shouldn't use them without adapting them first, otherwise you'll be making an ontological mistake.

Now, critical realism won't tell you exactly which theory to use, but it will help you to cull those that make these ontological mistakes. At the very least, this narrows stuff down for you.

We promised you a practical example, so here we go:

I (Tom) began a qualitative analysis that was interested in how higher education influences students' political attitudes. I started off by describing the way that studying Sociology appeared to influence students. My descriptions kept talking about 'Studying sociology did this' and 'Studying sociology did that'.

It took me a while to realise that in my descriptions I was only focussing on social structures. I'd completely neglected the students, their concerns, and the reasons they did stuff within the context of studying Sociology. My bad...

So, I went back to my data (with Archer's three steps) and modified my descriptions to emphasise both: the structural constraints and the concerns and aims of the students. This was super useful. It allowed me to develop a much richer and more nuanced argument, looking at how students developed projects in contexts that weren't of their own choosing.

Now, maybe you'd have done this automatically. If so, give yourself a pat on the back. If not, then Archer's three steps really can help to keep you on the right path.

Bonus feature: Transitive and intransitive objects of science

Okay, the final thing we'll talk about in this section is the distinction between *intransitive* and *transitive* objects of science, or between the world and our knowledge of it. Basically, this distinction helps us to better understand the aims of research. Research should seek theories (transitive objects) that better correspond to the real world (intransitive objects).

This neat division is more complex in social science. People like Ian Hacking (1999) have argued that *our theories can influence the world*. He gives the example of 'teenage pregnancy', where a theoretical concept (i.e. the transitive object) has influenced the real world (i.e. the intransitive object), as people now use the category of 'teenage parent' to identify or label others.

However, even with this complication, the distinction between intransitive and transitive objects of science is still useful. Just because there's a potential interaction between the transitive and intransitive objects of science doesn't change the fact that, at a given point in time, researchers should aim to produce theories that better correspond to the world.

The idea of the transitive and intransitive objects of science also links to a wider point about what truth is in our research.

From a critical realist perspective, truth is essentially the correspondence¹² of the transitive to the intransitive objects of science. Now, although this might be the definition of truth, it's debatable how useful this is in practice, like when we're deciding between one theory and another.

This is an area we think critical realism could do more. In RTS Bhaskar argues that we assess truth and validity through processes that are specific to our disciplines and doesn't really go into much more detail than that. Other people talk about 'judgemental rationality' as one of the three key features of critical realism (alongside a realist ontology and a relativist epistemology), but again, this hasn't been super well fleshed out.¹³

We're not going to solve this issue, but a couple of quick points.

One, critical realism highlights the importance of reflexivity. It's essential to take the time to think about whether our explanations are a good fit for our data, or whether our research likely increases or hinders human flourishing. In other words, we need to think about the theories we're using, where they came from, who they benefit, and who they exclude. This reflexivity is essential to help our work be more robust and have a positive impact on the world.

Two, we think critical realism could learn from pragmatist scholars. Correspondence theories of truth (i.e. the idea that truth is the correspondence of the transitive to the intransitive domains) might be an okay formal definition, but it's pretty useless at helping us select between two competing theories or explanations.¹⁴ To be fair, that's not what the definition is trying to do—the correspondence theory of truth only wants to give us a definition of *what* it means for something to be true. However, if critical realism wants to inform our scientific practice, then we think we need something more from our theory of truth. In this, pragmatist scholars such as Hasok Chang (2022) make the case for judging truth by the coherence of scientific practices.

We think that critical realism could do more to engage with this pragmatic work, as it might help provide more practical and relevant insights for researchers in the social and natural sciences.

Okay, that's the end of Section A.

Before we move onto the implications of critical realism for our methodology, we thought it might be useful to summarise some of the implications we've been talking about in a big table.

¹² Although Bhaskar himself claimed not to adhere to a correspondence theory of truth, most critical realists do. For a discussion see Ruth Groff (2000).

¹³ There was a recent special issue in the Journal of Critical Realism on this issue, see Robert K. Isaksen (2022) for more.

¹⁴ Hasok Chang (2022) explains that the correspondence theory of truth can be placed within a wider group of *truth by comparison*. This is a bit like judging the truth of a drawing of friend, by comparing this with the person's actual appearance in the world. Chang argues that this type of truth by comparison plays little role in scientific practice, instead we tend to judge truth by coherence.

Here you go.

Summary of implications

Ontological conclusions	Implications for research
- Experiments are <i>closed</i>, and the world is <i>open</i>. - The world is structured. It contains experiences, events and causes. - We should explore the causes that underlie events.	- You can distinguish experiences, events and causes to help conceptualise your research area. - You can produce causal explanations, rather than finding natural laws. - You can ask causal research questions.
- Causes only <i>tend</i> to bring about particular events outside of experimental conditions. - Theory of natural necessity: propertied things have powers and act to bring about events.	- You can seek nuanced explanations of how causes bring about events. - You shouldn't seek universal natural laws.
The world is stratified. It has emergence and depth, albeit complex.	- You can use interdisciplinary research methods. - You're justified in producing social science explanations, rather than reducing everything to individual behaviour or neuroscience.
Agency is possible, as humans are able to act on the basis of reasons.	- You should use social theories and descriptions that acknowledge agency. - You should value people's perspectives and worldviews. - You can use people's proto-scientific theories as a starting point. - Your social science research can critique false ideas and be emancipatory.
- Social structures are real. - Social structures constrain and enable people's action. - Social structures and agency interact: - Structures set the context. - People then act in these contexts on the basis of their concerns. - This action reproduces or transforms the structures, conditioning future action.	- You should use social theories and descriptions that acknowledge social structures are real. - You could use Archer's 3-step process to unpack agency, structure and their relation.
Transitive and intransitive objects of science.	You should incorporate reflexivity into your research design.

Section B: Implications for methodology

We're going to zoom out a little in this section. Instead of looking at the specific implications of Bhaskar's conclusions, we're going to think more broadly about the methodological lessons.

Perhaps most importantly, critical realism's conclusions should prompt us to re-evaluate the strengths and limitations of quantitative and qualitative methods, rather than accepting the conventional wisdom that *quants is causal; qual is descriptive*.

We're aiming to make this a really tangible discussion. Hopefully, if you're just about to embark on some research design, whether in a grant application or your PhD, this is exactly up your street.

We'll start with quantitative research, then switch to qualitative work, and end by briefly discussing mixed methods designs.

Before starting, one quick comment that might help you to understand where we're coming from: we think that methodologies and methods are often compatible with several philosophical positions. Interviews can be used by positivists and constructionists alike. The difference is that the former seek universal laws and the latter the co-production of stories. It's true that some methods are tied to particular philosophical traditions (e.g. it's hard to imagine positivistic discourse analysis or constructionist multi-level modelling) but we think there's often some flexibility.

It's the fact that methods often are flexible with respect to their philosophical underpinnings that explains why we need to think through what critical realist-inspired quantitative, qualitative and mixed methods research could look like.

Implications for quantitative methods

We've seen that critical realism justifies the aim to develop causal explanations through our research. This isn't a Humean approach, but one that stresses the need to look at the underlying causes of event regularities.

As we mentioned, it's not uncommon to see claims that quantitative methods are best placed to produce causal knowledge. There's plenty of methods classes that teach people quant is for causal and qual methods are descriptive.

This is a load of tosh.

First of all, this doesn't really describe the way many quantitative researchers think about their work. Sure, you'll definitely be able to find some quantitative work that adopts a naïve positivistic position, but if you look at the debates around quantitative methods, you won't find a bunch of people out with placards shouting: 'Quants is causal, Quants is causal'.

Instead, literature on quantitative methodology has lots of important debates, such as on how to interpret quantitative outputs or why we should avoid data dredging (where huge datasets are analysed for any significant relationships, whether or not there's any proposed

connection between the variables). These debates centre on the complications and pitfalls of moving from quantitative analysis to a claim about the world, rejecting the idea that this is a simple process of describing patterns of events. For this reason, we don't think high-quality quantitative work is automatically positivistic, despite what might be implied in some introductory methodology textbooks.

The second reason it's rubbish to say quantitative research is the only route to causal knowledge relates more closely to critical realism. If causal knowledge is about understanding the reason why two events are connected as they are, rather than just detecting the event itself, then quantitative research isn't particularly well placed to do this.

The output of a regression model doesn't offer you a nuanced causal explanation, in and of itself. It only tells you that a given set of variables (i.e. the ones you chose to whack in your model) have a certain association with your dependent variable. In other words, the output of a regression model isn't an explanation—it's an empirical pattern that, itself, requires an explanation.¹⁵

Again, there's not many quantitative researchers who would disagree with this.

Okay, so if we're throwing the conventional teaching on quants in the bin, and rejecting the simple equation of quants and causal knowledge, what's the alternative?

First and foremost, we believe quantitative methods are valuable research tools. Not all research needs to have a quantitative component, but these approaches do offer much to social science. For example, there's some aspects of the world that are much easier to access via quantitative methods, and we'd be missing out if we limited our work to qualitative research designs.

Secondly, we think it's helpful not to present a straw man of quantitative research. While positivistic language does creep into many quantitative papers, especially around prediction, effect sizes and 'the model for', it's worth resisting the temptation to dismiss this work as useless from the outset. As we said, we think high-quality quantitative work already has considerable compatibility with a critical realist worldview, i.e. a critical realist approach to quantitative research has strong parallels to best practice advocated within quantitative research traditions.

So, what is high-quality quantitative research?

High-quality research offers a clear explanation of why it investigates a relationship between two or more variables. It doesn't matter what type of stats are used, if you don't have a reason for looking at a particular relationship, then it's going to be super hard to generate any useful findings. Similarly, all control variables need to be justified and you need to think about their relationship with other variables (Kohler et al, 2023).

¹⁵ We're grateful to Dave Elder-Vass for this succinct way of characterising the output.

To put this another way, good quantitative research is *theory-driven*. It's theory driven in the sense that there's a clear theory about how the two or more variables might be related to each other.

Now, this theory-driven quantitative work, can do at least two different things:

1. It can provide some evidence for or against a particular theory.
2. It can estimate the *effect* of a particular cause in a given context.

On the first of these, quantitative research can test a theory's claims about how variables are related. This then provides some evidence for or against the theory.

For example, if we expect to find that students from lower socio-economic backgrounds are disadvantaged within university admissions procedures, we could look to see if this is the case. If we do find that students from lower socio-economic backgrounds have lower rates of access to prestigious universities, especially when we control for academic achievement, this is some compelling evidence in favour of the theory that university admissions place some applicants at a disadvantage.

Note, this quantitative research won't provide us with a nuanced explanation of *how* and *why* some students are placed at a disadvantage. Remember, quantitative research doesn't give us the causal explanation itself, but it can provide evidence for or against the existence of social inequalities in university access.

However, this evidence-giving isn't the only thing quantitative research can do. It can also begin to estimate the *effect* of a particular cause.

Revisiting our example of socio-economic background and university access, quantitative research can indicate whether socio-economic background seems to have a small or large impact. While it won't give us a single number that captures the universal impact of socio-economic background, it will give us some insight into the relative importance of this feature in a given context. For example, what is the likelihood of a working-class student from a state-educated background gaining an offer from a selective university in the UK, compared with their middle-class and privately educated peers?

Now, at first glance, this might seem a bit controversial from a critical realist perspective. If quantitative methods can assess causal effects, are we not jumping back to the idea that quant research is for causes?

No.

There's an important difference between *causal effects* and *causal explanations/processes*.¹⁶

Critical realism has shown us that it's important to develop causal knowledge in the form of causal explanations. These causal explanations involve some description of the *process* by

¹⁶ For more on this distinction see: Dretske, 1991; Hirshman & Reed, 2014; Goertz & Mahoney, 2013.

which two things are related or the *process* by which a particular event regularity comes to be.

Quantitative research doesn't provide us with a detailed understanding of these causal processes. A regression model alone won't tell you the process that links two variables.

At most (if it's a theory-driven analysis) it will provide us with some evidence for or against a particular hypothesised causal process and begin to estimate its importance.

If it's causal processes you seek, and this is at the heart of critical realist approaches to causal knowledge, then it's theory-driven qualitative research that's best placed to do this (more on this in the next section).

Quantitative research can provide evidence for these processes and estimate their effects, but it doesn't give you the *process* itself.

It's probably worth clarifying that theory-driven quantitative research tends to consider a causal process in a *particular part of the world*. In other words, we're not using quantitative methods to search for universal relationships between variables but looking at the variation in impact on different actors in different contexts.

Again, this reflects much of what good quality quantitative research already does. Most well-designed quantitative research doesn't end up arguing that it's found some universal relationship between variables, at least not in social science.

To take a quick example, there's no universal relationship between: 1) manager/employee relationships, and 2) wellbeing at work. There's huge differences between different countries or different employment sectors, on what appropriate manager/employee relationships look like. A relationship in one context may promote wellbeing but do the exact opposite in another. Think about academics and professional footballers. While a footballer can expect substantial aspects of their diet and leisure time to be managed, we can only imagine how this idea would be received in a Sociology department faculty meeting.

Again, good quality quantitative research recognises this complexity. It's not seeking a universal relationship between A and B—and instead gives a lot of detail about sampling processes and representativeness, rather than claiming a simple universal relation. This is also why high-quality quant research explicitly conceptualises the relations between variables and the ways they might interact with one another. These techniques allow us to answer how a third variable changes the original relation between A and B in a given context.

To revisit a point we made above, there's not really specifically critical realist methods of research. Instead, critical realism has the capacity to unify and bring together different methods. To us, the question isn't whether quantitative methods are compatible with critical realism (it's clear they can be), but more practical questions such as:

- Does this particular research question make sense? Is it theoretically grounded?
- What's the best method to answer this particular research question?
- What's the best practice for using this method?

- What assumptions am I making about the world in this method?

Some of you may take these questions for granted when thinking about a quantitative design. If that's the case, great! As we've said, we think high-quality quantitative work would do this. But there's also plenty of work that appears to forget this, so hopefully this is still useful.

Before moving on to qualitative work, it's probably worth highlighting that there are some critical realists who don't believe many quantitative methods are compatible with critical realism. Or at least they're sceptical about some of these techniques. If you're looking for some references, just as a counter point to the vision we've put forward, you could look at some of Tony Lawson's (2012) work that argues mathematical modelling is unsuitable for explaining the social world, particularly as it relates to economics. There's also Andrew Sayer's (2010) work that has explored the limitations of quantitative methods in social science research.

Don't panic though, we're not alone in seeing a substantial space for quantitative methods in critical realist inspired research. There's lots of people like Wendy Olsen and Jamie Morgan (2005); Hubert Buch-Hansen (2014); Doug Porpora (2015) and Peter Næss (2004) making this case.

We thought we'd end this sub-section with a few practical recommendations on quantitative methods that have been useful for us:

- For those interested in econometrics: Angrist, J.D., & Pischke, J.S. (2014). *Mastering 'Metrics: The path from cause to effect*. Princeton University Press.
- For a critical paper on causal analysis in quantitative sociology: Kohler, U., Class, F., & Sawert, T. (2023). Control variable selection in applied quantitative sociology: A critical review. *European Sociological Review*, 40(1), 173-186.
- A personal favourite of Cristián's is a practical guide for statistical analysis on RStudio: Urdinez, F., & Cruz, A. (2020). *R for political data science: A practical guide*. CRC Press. There's also a Spanish-language open access version on [Github](#)
- For a critical appraisal of regressions in critical realism: Porpora, D.V. (2015) *Do realists run regressions?*. This is the second chapter of Porpora, D.V. (2015). *Reconstructing Sociology: The critical realist approach*. Cambridge University Press.
- For network analysis in critical realism: Buch-Hansen, H. (2014). Social Network Analysis and Critical Realism. *Journal for the Theory of Social Behaviour*, 44(3), 306–325.
- Quite a big book for almost anything on data analysis using regression and multilevel modelling: Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.
- Another personal favourite of Cristián's is a specific example of theoretically-based explanations using quantitative methods: Torche, F. (2005). Unequal But Fluid: Social Mobility in Chile in Comparative Perspective. *American Sociological Review*, 70(3), 422–450.

Implications for qualitative research

Okay, we've spent a fair bit of time talking about quantitative research and the ways that critical realism challenges the conventional wisdom (or at least adds more nuance) that quantitative research is the route to causal knowledge.

However, we've said a bunch of times that critical realism also challenges the idea that qualitative research can only be descriptive. Instead, we can use qualitative work to reach explanations.

While conventional critical realist thinking would argue that qualitative work produces causal explanations, we're going to extend this to argue that qualitative methods are best placed to produce explanations, whether or not these explanations are causal. In other words, good qualitative work should *explain* its object of interest.

As we said, following Bhaskar's approach to causation, you'll hear many critical realists say things like: "Qualitative work is best placed to produce nuanced causal explanations". This is because it's qualitative studies that allow us to investigate how causes work, how they might interact with other causes, and how they produce the experiences and events that we're interested in.

This is a powerful argument, and one that we agree with. It completely flips the normal focus on quantitative methods for causation, with qualitative methods relegated to providing context or a bit of extra detail. When we focus on *causal processes* rather than *causal effects*, it's qualitative methods that are best suited to give you knowledge about these processes and mechanisms.

Take the example of the relationship between the experience of poverty during childhood and future earnings. This may seem like a textbook example of a relationship that begs for a quantitative design, but there's a bunch of things that we can unpack between these two variables. For example, if you want to understand how experience of poverty during childhood impacts future life projects then the best way to answer this is to conduct in-depth qualitative work that focuses on people's lived experiences (this is a bit like some of Archer's (2003) work).

But, if you want to quantify the *effects* of childhood poverty on future earnings, then you need a quantitative design. This will probably be a quantitative approach that follows a 'life course' approach (Spini et al., 2017) – this would likely involve a survey and following individuals over time, asking them about their childhood and then running a statistical analysis to explore how their current salary relates to their childhood.¹⁷

Notice that both the qualitative and the quantitative study are *explaining* something. However, the thing they're trying to explain is different. While the qualitative study seeks to explain how people's lived experiences of poverty impact their future life projects (this is explaining a causal process), the quantitative study explains the relation between childhood poverty and future earnings (this is explaining a causal effect).

¹⁷ See examples of: Berg, et al., 2017; Breen and Goldthorpe, 2001; Sirniö, Martikainen, & Kauppinen, 2013.

This argument that qualitative work is particularly useful for producing explanations of causal processes is similar to that made by other scholars, such as Tavory and Timmermans (2013), Lichterman and Reed (2015), Katz (2015), and Small and Cook (2023).

In this way, an explanation aims to show why something happened the way it did. This involves investigating the causes that produced the phenomenon. To do this, you need to identify some plausible causes and then put them to the test with your qualitative data. There's then a back and forth between possible mechanisms and testing with your data, in order to develop your explanations (see for e.g. Tavory and Timmermans (2022)).

Now, to do this type of qualitative work, you need to be well-versed in the relevant theory of your topic. We don't mean you need to be well-versed in more abstract social theories, whether Bourdieu, Marx, or Archer, but instead here we're talking about more 'mid-range' theories that give you alternative explanations and mechanisms to understand your data. For this reason, high-quality qualitative research also needs to be theoretically founded, although this doesn't deny the fact that previously untheorised processes may be needed to explain your data.

Okay, so far, we've talked about causal explanations and how qualitative methods can produce this kind of knowledge. But causal knowledge isn't the only kind of explanation. Other explanations don't talk about causal mechanisms but *how* a given phenomenon is *constituted* (Pacewicz, 2022).

This second type of explanation looks at three interrelated things: 1) the existence of some phenomenon (i.e. X exists); 2) the relation of the phenomenon to other things (i.e. X is a case of...); and 3) the categorisation of the phenomenon (i.e. X consists of $X_1, X_2, \dots, X_i$ elements).

To put this more concisely, a *constitutive argument* explains that a phenomenon has certain properties that make it what it is.

These *constitutive arguments* are different from causal arguments. Rather than showing that a phenomenon happened because of a given causal mechanism, constitutive arguments look at the properties of a thing and how those properties make it what it is.

Let's try to clarify this with the example of Cristian's empirical work on *scientific uncertainty*. We could offer two types of explanation about uncertainty: 1) we could identify things that trigger uncertainty and why this happens; and 2) we could provide an explanation of *what* uncertainty is. The first is causal and the second is constitutive. Although they're related it's worth differentiating between a question about 'what causes uncertainty' versus 'how can we better define uncertainty'. Let's expand on this further.

In my fieldwork and thanks to the specialized literature I realized that there are things that cause uncertainty in a laboratory. For example, a microbe that is difficult to recognize, a contaminated assay, or the question of funding. All these events generate uncertainty—they make things unclear and make us question what will happen next.

However, these causes of uncertainty don't really tell us how to define uncertainty. While the causes of uncertainty might help us to start thinking of the definition, they don't get us to a consistent and comprehensive definition of scientific uncertainty that applies in different places.

Instead, it was through ethnographic work in three laboratories that I saw uncertainty is experienced by scientists as a questioning of an expected course of action due to the possibility of committing errors in lab-work.

This conclusion tells us two things: 1) what uncertainty is (the questioning of expected courses of action); and 2) how it is caused in the lab (by the apprehension of the possibility of making mistakes in lab work).

Causal arguments and constitutive arguments are both explanatory as they show why something happened the way it did. This is fairly obvious for causal arguments—they explain the causes that gave rise to something. But constitutive arguments are also explanatory because they allow us to deepen and systematise our conceptualisations of phenomena, seeing why it is what it is.

In the case of uncertainty, the constitutive argument allows us to see why uncertainty is different to ignorance—the former can't be reduced to not-knowing, since uncertainty happens even when knowledge wasn't an issue (instead it was the desire to avoid the possibility of mistakes).

In this way, both causal and constitutive arguments are explanatory—it's just that these respond to different questions.

Now, of course, we're not saying that qualitative research always needs to give causal and constitutive explanations. That would be impossible. Or at least it would be too demanding.

What we are saying is that critical realism helps us to see qualitative research can provide explanations, and that these can be causal or constitutive.

Let's end this section by recommending a few pieces related to qualitative methods:

- A wonderful book for practical recommendations in qualitative research: Lareau, A. (2021). *Listening to people: A practical guide to interviewing, participant observation, data analysis, and writing it all up*. University of Chicago Press.
- On how to posit good questions: Mears, A. (2017) Puzzling in Sociology: On Doing and Undoing Theoretical Puzzles. *Sociological Theory*, 35(2), 138–146.
- For interview studies and their challenges: Small, M.L., & Cook, J.M. (2023). Using Interviews to Understand Why: Challenges and Strategies in the Study of Motivated Action. *Sociological Methods & Research*, 52(4), 1591–1631.
- Current approaches to social mechanisms: Ritz, B. (2023) Social mechanisms: bridging critical realist and pragmatist approaches. *Journal of Critical Realism*, 22(3), 404-410.

- On constitutive arguments: Pacewicz, J. (2022). What Can You Do with a Single Case? How to Think About Ethnographic Case Selection Like a Historical Sociologist. *Sociological Methods & Research*, 51(3), 931–962.
- Case selection: Small, M. L. (2009). ‘How many cases do I need?’: On science and the logic of case selection in field-based research. *Ethnography*, 10(1), 5–38.

Implications for mixed methods research

So far, we’ve explored the ways in which qualitative and quantitative methods can be used in a way that’s compatible with critical realism. In this process, we’ve seen the ways that these different methodologies have different strengths and weaknesses.

Given this, it’s a fairly obvious point to say that if we bring these two types of methods together in our work, this can be especially powerful. It’s by bringing them together, both quant and qual research, that we can begin to address the limitations of each, which potentially gives us a more comprehensive picture of the world.

Now, we’re absolutely not saying that all research needs to be mixed methods. Mixed methods work is time-intensive and difficult to do (e.g. you often need a team with a mix of skills) which just isn’t always possible. Similarly, the research that already exists in the field might influence whether or not a mixed methods approach is more or less desirable. If there’s already tonnes of qualitative work, you might not need to do this yourself.

Okay, so if mixed methods research can be powerful, because it addresses the relative weaknesses of both quant and qual research, can we say anything else?

Yes!

We think there’s a couple of mixed methods research designs that are particularly important (although this isn’t to say these are the only two):

- In many research areas, a sequential design with quantitative first and qualitative second makes sense. An initial quantitative study could look for patterns of events, which are then explored and explained through in-depth qualitative work. For example, we could initially use quantitative methods to see whether students’ political attitudes change during their time in higher education, and we could then perform qualitative analysis to assess the specific causes of some of these changes.
- In really under-explored research areas, it might make sense to do qualitative work first, just to gain some idea of what’s happening in terms of experiences and events. You’ll then be better placed to develop quantitative tools to assess a broader population. And of course, you could revisit qualitative methods to explore the potential causes for these patterns.

Now, there’s much more that could be said here, but a lot of this falls beyond the scope of RTS and its philosophical conclusions.

RTS and its conclusions allow us to evaluate the strengths and limitations of quantitative and qualitative methods, and give us some suggestions about what high-quality work looks like in these areas. However, these philosophical conclusions tell us less about how to

design our specific studies in response to a particular research area. This is necessarily a more specific and contextual question.

It's for this reason that the above discussion on mixed methods is brief. You absolutely should take the next step to think about how qualitative and quantitative methods could be used in your research area, but which one you choose, and whether/how you choose to combine these in a mixed methods approach has to respond to the needs of your particular research area. There's not cookie cutter approaches that are useful to discuss, or at least not ones that derive from the philosophical conclusions that Bhaskar draws in RTS.

That being said, we thought it might be useful to point you towards a study that we feel captures best practice in mixed-methods research. We'd really recommend checking out Robert Fishman and Omar Lizardo's (2013) *How Macro-historical Change Shapes Cultural Taste: Legacies of Democratization in Spain and Portugal*. In the words of Cristián: "It's just so well written – read it!".

Okay, let's wrap up by summarising the core contributions of critical realist work from RTS to our methodology, which can be broken into three main points:

Point 1: quant methods can be used in ways that are compatible with critical realism. Theory-driven work can both provide evidence related to this theory, as well as looking at causal effects. It's not great at describing and understanding causal processes though.

Point 2: qual research can provide in-depth work to develop explanations. This understanding may help us judge how transferable an explanation is to another context.

Point 3: mixed methods designs can be really useful in critical realist inspired research, as they can address the limitations of each type of research.

A quick goodbye

Well folks, we've reached the end of the road, and we hope this discussion has been useful.

We aimed to introduce you to research philosophy and give you a way to navigate through questions about your ontology and epistemology.

Also, we hope that this has given you a way to start a journey of getting to know critical realism. There's a lot more to explore, but we hope this account of RTS, which really is *the foundational text* of critical realism gives you a decent start.

We think of it as setting you off on a journey with a few words of encouragement and a metaphorical backpack stuffed with sugary snacks.

Good luck with this trip, if you decide to take it. And good luck with the rest of your research.

Appendix: An annotated bibliography on critical realism's main critiques and contending positions

Clearly, both of us think critical realism has some real strengths, but it's also important to recognise a range of critiques of this position. In some ways, we're both quite distant from core Bhaskarian critical realists, with Cristián in particular levitating more towards a sort of pragmatic realism.

Throughout the book we've tried to mention some of the debates around critical realism, but we've not really tackled some of the criticisms and critiques head-on. For those of you that are interested, we decided to add this appendix to point you towards some of these debates. We've written this as an annotated bibliography to help you to pick and choose any areas you want to explore.

1. Critical realism is not alone nor the best

Several publications note that critical realism can be a slightly closed tradition, one that has not always made extensive efforts to engage with other philosophical traditions, although of course there are exceptions for individual researchers. Some of these publications make the case that other traditions accomplish similar objectives in a more coherent and robust way. We also include arguments that propose other forms of realism or other post-positivist positions.

- Fuhse, J. A. (2022). How Can Theories Represent Social Phenomena? *Sociological Theory*, 40(2), 99–123.
 - This provides a critique of critical realism's concepts of truth and reality, and outlines a constructivist proposal.
- Lizardo, O. (2013). *Three thousand more words on critical realism*. Orgtheory.net.
 - This blog provides a digestible critique and positioning *viz a viz* critical realism, coming from one of the best North American sociologists.
- Navarrete, C. (2024). Rearticulating realism: chains of reference and epistemic success in the sciences. *Sociological Theory*, 42(4), 354-379.
 - This article puts forward a new set of arguments for both realism and naturalism. It is not necessarily inconsistent with critical realism, but it does start from a different argument and has different conclusions. It seeks to provide a non-ontological argument for realism because this is something that doesn't sit very well for scholars coming from non-realist positions.
- Reed, I. (2008). Justifying Sociological Knowledge: From Realism to Interpretation. *Sociological Theory*, 26(2), 101–129.
 - This is kind of an article version of Reed's (2011) book. It is a great analysis of how different positions understand sociological knowledge and a defense of interpretivism.
- Reed, I. (2011). *Interpretation and social knowledge: On the use of theory in the human sciences*. University of Chicago Press.
 - Chapter 2 focuses on realism, but Cristián really encourages you to read the whole book. It is such a well written book on social knowledge. As anything by Reed to be honest.

- Zhang, T. (2023). Critical realism: A critical evaluation. *Social Epistemology*, 37(1), 15-29.
 - This is a critical appraisal of critical realism in the context of philosophy of science. The main idea is that within the philosophy of science there are other options that do better than critical realism in reaching its realist objectives.

2. Naturalist critiques

There are also some scholars that try to link the foundations of critical realism more tightly to empirical research. We'd recommend:

- Kaidesoja, T. (2013). *Naturalizing critical realist social ontology*. Routledge.
 - A great book that provides a particular reading of Bhaskar's realist argument. Instead of framing it as a transcendental argument, Kaidesoja argues that we should frame the argument as an *inference to the best explanation* because these kind of arguments are open to empirical scrutiny.
- Sarkia, M., & Kaidesoja, T. (2023). Two approaches to naturalistic social ontology. *Synthese*, 201(3), 104.
 - Similar to Kaidesoja's book, this article provides two approaches to social ontology that are open to empirical scrutiny. The authors seek to maintain what is characteristic of social ontology, thinking about entities, but rendering the debate in such a manner that we can decide by looking at "empirical evidence".

3. Critical realism is not well grounded

Some publications argue that the ontology, epistemology and/or methodology associated with critical realism are inconsistent:

- Martin, J.L. (2015). *Thinking through theory*. W.W. Norton & Company.
 - Martin presents a critique of critical realism from pp.95-110. The book introduces people to theoretical thinking, how to think about social theory, and how to craft explanations. Martin's work on fields and culture are also one of the best pieces of social theory currently out there.
- Tsilipakos, L. (2014). Theoretical Procedures and Elder-Vass's Critical Realist Ontology. *Philosophy of the Social Sciences*, 44(6), 752–773.
 - This article analyses Elder-Vass's formulation of critical realism and argues that it is not well grounded in its usage of scientific practices or examples.

4. Actor Network theory

Some scholars working within actor network theory, or studies of science more generally, have also engaged with and critiqued critical realism:

- Baxter, J., & Chua, W. F. (2020). Using critical realism in critical accounting research—a commentary by two ANTs. *Accounting, Auditing & Accountability Journal*, 33(3), 655-665.

- This a critique of critical realism that insists there are no intransitive structural forces that can be identified by retroduction; there are only actor-networks.
- Kemp, S. (2022). Organizational de-structuring? Latour's potential contribution to the critical realist – pragmatist dispute. *Journal of Critical Realism*, 21(3), 1–22.
 - This article brings a pragmatist reading of Bruno Latour's work into conversation with Dave Elder-Vass. It debates the ontological status of organisations and in doing so posits challenges for critical realism.

5. Debates on cultural sociology

This group of literature includes publications that do not address critical realism directly but can be read as critiquing its interpretation of culture. The main point of the critique is Archer's view of culture as a system that has definite properties beyond human minds, as well as the more general idea that culture is a complex web of meanings. Both articles are part of a wider movement that re-thinks sociological concepts and methods in relation to cognitive neuroscience discoveries.

- Martin, J. L. (2010). Life's a beach but you're an ant, and other unwelcome news for the sociology of culture. *Poetics*, 38(2), 229–244.
 - In this article, Martin argues that common approaches to culture within sociology are problematic. He identifies a dilemma: either culture is a complex cognitive web that is largely shared by actors, or it is "stored" in people's head. Usually we commit to both, but this is unlikely due to our limited cognitive powers.
- Lizardo, O. (2023). An Analytical Approach to Culture. *Philosophy of the Social Sciences*, 53(4), 281–302.
 - Lizardo outlines a general approach to culture. The paper does this in an analytical manner, focusing on recent works in the cognitive sciences and its relation to sociology.

6. Emergence

There are also some publications that debate critical realism's concept of emergence, directly or indirectly.

- Hansson Wahlberg, T. (2014). Causally redundant social objects: Rejoinder to Elder-Vass. *Philosophy of the Social Sciences*, 44(6), 798–809.
 - This article argues that critical realism (specifically Elder-Vass's account) posits causally redundant social objects. The main idea is that emergent wholes are redundant because we can causally explain without them.
- Latour, B. (2020a). "Composing the New Body Politic from Bits and Pieces." and Latour, B. (2020b). "Afterword: Politics - A Glimpse at Bodybuilding." In *A Book of the Body Politic: Connecting Biology, Politics and Social Theory*, edited by B. Latour, S. Schaffer, and P. Gagliardi, 285–302. Venice: Fondazione Giorgio Cini
 - Both of Latour's chapters are in the same book, and while they're not specifically about emergence, they do offer some cautionary words against emergentist positions. This is a great read for anyone interested in the

relationship between “natural” phenomena and social sciences, as well as anyone that wants to understand why Latour and some people from science and technology studies are not persuaded by emergentism.

- Searle, J. R. (2016). The limits of emergence: Reply to Tony Lawson. *Journal for the Theory of Social Behaviour*, 46(4), 400–412.
 - This is a great discussion between Lawson and Searle. It outlines certain objectives that social ontology ought to secure, and how Lawson’s critical realist position may struggle to get at them.

We’re definitely missing some critiques, but we hope this is enough to get you started, for anyone who wants to take this further.

And that is definitely the end.

Unless you really want to read the references...

References

- Anderson, E. (2024). *Feminist Epistemology and Philosophy of Science*. The Stanford Encyclopedia of Philosophy. Available from: <https://plato.stanford.edu/archives/fall2024/entries/feminism-epistemology/>
- Angrist, J.D., & Pischke, J.S. (2014). *Mastering 'Metrics: The path from cause to effect*. Princeton University Press.
- Archer, M.S. (1996). *Culture and Agency*. Cambridge University Press.
- Archer, M.S. (2000). *Being human: The problem of agency*. Cambridge University Press.
- Archer, M.S. (2003). *Structure, agency and the internal conversation*. Cambridge University Press.
- Archer, M.S. & Elder-Vass, D. (2012). Cultural System or norm circles? An exchange. *European Journal of Social Theory*, 15(1), 93-115. <https://doi.org/10.1177/1368431011423592>
- Baxter, J., & Chua, W.F. (2020). Using critical realism in critical accounting research—a commentary by two ANTs. *Accounting, Auditing & Accountability Journal*, 33(3), 655-665.
- Berg, N., Kiviruusu, O., Karvonen, S., Rahkonen, O., & Huurre, T. (2017). Pathways from poor family relationships in adolescence to economic adversity in mid-adulthood. *Advances in Life Course Research*, 32, 65–78. <https://doi.org/10.1016/j.alcr.2016.07.001>
- Bhaskar, R. (2008). *A Realist Theory of Science*. Third Edition. Verso.
- Bhaskar, R. (2008). *Dialectic: The Pulse of Freedom*. Second Edition. Routledge.
- Bhaskar, R. (2015). *The Possibility of Naturalism: A Philosophical Critique of the Contemporary Human Sciences*. Fourth Edition. Routledge.
- Breen, R., & Goldthorpe, J.H. (2001). Class, mobility and merit the experience of two British birth cohorts. *European Sociological Review*, 17(2), 81–101.
- Buch-Hansen, H. (2014). Social Network Analysis and Critical Realism. *Journal for the Theory of Social Behaviour*, 44(3), 306–325. <https://doi.org/10.1111/jtsb.12044>
- Carmichael, M.A., Thomson, R.L., Moran, L.J., & Wycherley, T.P. (2021). The Impact of Menstrual Cycle Phase on Athletes' Performance: A Narrative Review. *International Journal of Environmental Research and Public Health*, 18(4), 1667. <https://doi.org/10.3390/ijerph18041667>
- Chang, H. (2022). *Realism for Realistic People*. Cambridge University Press.

Collier, A. (1994). *Critical realism: An introduction to the philosophy of Roy Bhaskar*. London: Verso.

Criado Perez, C. (2019) *Invisible Women: Exposing Data Bias in a World Designed for Men*. Chatto & Windus.

Danermark, B., Ekström, M., & Karlsson, J.C. (2019) *Explaining Society: Critical Realism in the Social Sciences*. Routledge.

Dretske, F. (1991). *Explaining behavior: Reasons in a world of causes*. MIT Press.

Fishman, R.M., & Lizardo, O. (2013). How Macro-Historical Change Shapes Cultural Taste: Legacies of Democratization in Spain and Portugal. *American Sociological Review*, 78(2), 213–239. <https://doi.org/10.1177/0003122413478816>

Elder-Vass, D. (2005). Emergence and the Realist Account of Cause. *Journal of Critical Realism*, 4(2), 315–338. <https://doi.org/10.1558/jocr.v4i2.315>

Elder-Vass, D. (2010). *The Causal Power of Social Structures*. Cambridge University Press.

Elder-Vass, D. (2012). *The Reality of Social Construction*. Cambridge University Press.

Elder-Vass, D., Fryer, T., Groff, R.P., Navarrete, C., & Nellhaus, T. (2023). Does critical realism need the concept of three domains of reality? A roundtable. *Journal of Critical Realism*, 22(2), 222–239. <https://doi.org/10.1080/14767430.2023.2180965>

Fryer, T. & Navarrete, C. (2024) *A short guide to ontology and epistemology: why everyone should be a realist*. Self-published, 3rd edition. Available from: <https://tfryer.com/ontology-guide/>

Fuhse, J.A. (2022). How Can Theories Represent Social Phenomena? *Sociological Theory*, 40(2), 99–123. <https://doi.org/10.1177/07352751221087719>

Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.

Goertz, G., & Mahoney, J. (2013). *A Tale of Two Cultures: Qualitative and Quantitative Research in the Social Sciences*. Princeton University Press. <https://doi.org/10.1515/9781400845446>

Graeber, D. (2015). Radical alterity is just another way of saying “reality”: A reply to Eduardo Viveiros de Castro. *HAU: Journal of Ethnographic Theory*, 5(2), 1–41. <https://doi.org/10.14318/hau5.2.003>

Groff, R.P. (2000). The Truth of the Matter: Roy Bhaskar’s Critical Realism and the Concept of Alethic Truth. *Philosophy of the Social Sciences*, 30(3), 407–435. <https://doi.org/10.1177/004839310003000304>

- Groff, R.P. (2017). Causal Mechanisms and the Philosophy of Causation. *Journal for the Theory of Social Behaviour*, 47(3), 286–305. <https://doi.org/10.1111/jtsb.12118>
- Groff, R.P. (2020). What Does the Doing? On Powers, Things and Powerful Things. *Synthese Library*, 417. Springer.
- Groff, R.P. (2021). Conceptualizing causal powers: Activity, capacity, essence, necessitation. *Synthese*, 199(3–4), 9881–9896. <https://doi.org/10.1007/s11229-021-03229-x>
- Hacking, I. (1999). *The social construction of what?* Harvard University Press.
- Hansson Wahlberg, T. (2014). Causally redundant social objects: Rejoinder to Elder-Vass. *Philosophy of the Social Sciences*, 44(6), 798–809. <https://doi.org/10.1177/0048393114521364>.
- Hempel, C.G., & Oppenheim, P. (1948). Studies in the Logic of Explanation. *Philosophy of Science*, 15(2), 135–175.
- Hirschman, D., & Reed, I.A. (2014). Formation Stories and Causality in Sociology. *Sociological Theory*, 32(4), 259–282. <https://doi.org/10.1177/0735275114558632>
- Isaksen, R. (2022). Special Issue: Judgemental Rationality. *Journal of Critical Realism*, 21(5), 589–91. <https://doi.org/10.1080/14767430.2022.2146930>.
- Kaidesoja, T. (2013). *Naturalizing critical realist social ontology*. Routledge.
- Katz, J. (2015). Situational Evidence: Strategies for Causal Reasoning from Observational Field Notes. *Sociological Methods & Research*, 44(1), 108–144. <https://doi.org/10.1177/0049124114554870>
- Kemp, S. (2022). Organizational de-structuring? Latour’s potential contribution to the critical realist – pragmatist dispute. *Journal of Critical Realism*, 21(3), 1–22. <https://doi.org/10.1080/14767430.2022.2049090>
- Kohler, U., Class, F., & Sawert, T. (2023). Control variable selection in applied quantitative sociology: A critical review. *European Sociological Review*, 40(1), 173–186. <https://doi.org/10.1093/esr/jcac078>.
- Lareau, A. (2021). *Listening to people: A practical guide to interviewing, participant observation, data analysis, and writing it all up*. University of Chicago Press.
- Latour, B. (2020a). Composing the New Body Politic from Bits and Pieces. In *A Book of the Body Politic: Connecting Biology, Politics and Social Theory*, edited by B. Latour, S. Schaffer, and P. Gagliardi, 19–38. Venice: Fondazione Giorgio Cini

Latour, B. (2020b). Afterword: Politics - A Glimpse at Bodybuilding. In *A Book of the Body Politic: Connecting Biology, Politics and Social Theory*, edited by B. Latour, S. Schaffer, and P. Gagliardi, 285–302. Venice: Fondazione Giorgio Cini

Lawson, T. (2012). Mathematical Modelling and Ideology in the Economics Academy: competing explanations of the failings of the modern discipline? *Economic Thought*, 1(1), 3–22.

Lichterman, P., & Reed, I.A. (2015). Theory and Contrastive Explanation in Ethnography. *Sociological Methods & Research*, 44(4), 585–635.
<https://doi.org/10.1177/0049124114554458>

Lizardo, O. (2013). *Three thousand more words on critical realism*. Orgtheory.net. Available from: <https://orgtheory.wordpress.com/2013/09/09/three-thousand-more-words-on-critical-realism/>

Lizardo, O. (2023). An Analytical Approach to Culture. *Philosophy of the Social Sciences*, 53(4), 281–302. <https://doi.org/10.1177/00483931231169313>.

Martin, J.L. (2010). Life's a beach but you're an ant, and other unwelcome news for the sociology of culture. *Poetics*, 38(2), 229–244. <https://doi.org/10.1016/j.poetic.2009.11.004>

Martin, J.L. (2015). *Thinking through theory*. W.W. Norton & Company.

Mears, A. (2017). Puzzling in Sociology: On Doing and Undoing Theoretical Puzzles. *Sociological Theory*, 35(2), 138–146. <https://doi.org/10.1177/0735275117709775>

Næss, P. (2004). Prediction, Regressions and Critical Realism. *Journal of Critical Realism*, 3(1), 133–164. <https://doi.org/10.1558/jcr.v3i1.133>

Navarrete, C. (2024). Rearticulating realism: chains of reference and epistemic success in the sciences. *Sociological Theory*, 42(4), 354–379. <https://doi.org/10.1177/07352751241291831>.

Navarrete, C. & Fryer, T. (2024). Redefining emergence: Making the case for contextual emergence in critical realism. *Journal for the Theory of Social Behaviour*, 1–18.
<https://doi.org/10.1111/jtsb.12414>

Nichol, A., Hastings, C. & Elder-Vass, D. (2023). Putting philosophy to work: developing the conceptual architecture of research projects. *Journal of Critical Realism*, 22(3), 364–383.
<https://doi.org/10.1080/14767430.2023.2217054>.

Olsen, W. & Morgan, J. (2005). Towards a critical epistemology of analytical statistics: Realism in mathematical method. *Journal for the Theory of Social Behaviour* 35 (3):255–284.

Pacewicz, J. (2022). What Can You Do with a Single Case? How to Think About Ethnographic Case Selection Like a Historical Sociologist. *Sociological Methods & Research*, 51(3), 931–962.
<https://doi.org/10.1177/0049124119901213>

- Porpora, D.V. (2015). *Reconstructing Sociology: The critical realist approach*. Cambridge University Press.
- Reed, I.A. (2008). Justifying Sociological Knowledge: From Realism to Interpretation. *Sociological Theory*, 26(2), 101–129. <https://doi.org/10.1111/j.1467-9558.2008.00321.x>
- Reed, I.A. (2011). *Interpretation and social knowledge: On the use of theory in the human sciences*. University of Chicago Press.
- Ritz, B. (2023). Social mechanisms: Bridging critical realist and pragmatist approaches. *Journal of Critical Realism*, 22(3), 404–410. <https://doi.org/10.1080/14767430.2023.2217040>
- Rutzou, T., & Elder-Vass, D. (2019). On Assemblages and Things: Fluidity, Stability, Causation Stories, and Formation Stories. *Sociological Theory*, 37(4), 401–424. <https://doi.org/10.1177/0735275119888250>
- Sarkia, M., & Kaidesoja, T. (2023). Two approaches to naturalistic social ontology. *Synthese*, 201(3), 104. <https://doi.org/10.1007/s11229-023-04105-6>.
- Sayer, A. (2010). *Method in social science: A realist approach*. Second Edition. Routledge.
- Searle, J.R. (2016). The limits of emergence: Reply to Tony Lawson. *Journal for the Theory of Social Behaviour*, 46(4), 400–412. <https://doi.org/10.1111/jtsb.12125>
- Sirniö, O., Martikainen, P., & Kauppinen, T.M. (2013). Intergenerational determinants of income level in Finland. *Social Forces*, 92(2), 463–490.
- Small, M.L. (2009). ‘How many cases do I need?’: On science and the logic of case selection in field-based research. *Ethnography*, 10(1), 5–38. <https://doi.org/10.1177/1466138108099586>
- Small, M.L., & Cook, J.M. (2023). Using Interviews to Understand Why: Challenges and Strategies in the Study of Motivated Action. *Sociological Methods & Research*, 52(4), 1591–1631. <https://doi.org/10.1177/0049124121995552>
- Spini, D., Bernardi, L., & Oris, M. (2017). Toward a Life Course Framework for Studying Vulnerability. *Research in Human Development*, 14(1), 5–25. <https://doi.org/10.1080/15427609.2016.1268892>
- Tavory, I., & Timmermans, S. (2013). A Pragmatist Approach to Causality in Ethnography. *American Journal of Sociology*, 119(3), 682–714. <https://doi.org/10.1086/675891>
- Timmermans, S., & Tavory, I. (2022). *Data analysis in qualitative research: Theorizing with abductive analysis*. University of Chicago Press.

Torche, F. (2005). Unequal But Fluid: Social Mobility in Chile in Comparative Perspective. *American Sociological Review*, 70(3), 422–450.
<https://doi.org/10.1177/000312240507000304>

Tsilipakos, L. (2014). Theoretical Procedures and Elder-Vass's Critical Realist Ontology. *Philosophy of the Social Sciences*, 44(6), 752–773.
<https://doi.org/10.1177/0048393112461055>

Uebel, T. (2013). "Logical Positivism" — "Logical Empiricism": What's in a Name? *Perspectives on Science*, 21(1), 58–99.

Urdinez, F., & Cruz, A. (2020). *R for political data science: A practical guide*. CRC Press.

Zhang, T. (2023). Critical realism: A critical evaluation. *Social Epistemology*, 37(1), 15-29.
<https://doi.org/10.1080/02691728.2022.2080127>.