This is a post-peer-review, pre-copyedit version of an article published in Higher Education Policy. The final authenticated version is available online at: http://doi.org/10.1057/s41307-021-00232-2.

Full reference: Fryer, T (2021) Conceptualising graduate outcomes with critical realism, *Higher Education Policy*, http://doi.org/10.1057/s41307-021-00232-2.

Conceptualising graduate outcomes with critical realism

Tom Fryer, University of Manchester

Abstract

Graduate outcomes are becoming increasingly prominent within higher education (HE) policy, driven by national governments keen to demonstrate 'value for money'. The majority of HE policy in this area uses narrow economic metrics, such as employment status and salary, often derived from national surveys of graduates. This paper uses critical realist philosophy to develop a set of foundational concepts (graduate functionings, graduate capabilities and graduate outcomes) that illuminate the key characteristics and mistakes of this HE policy. It is shown that the narrow economic metrics used in policy are graduate functionings not graduate outcomes—they describe how graduates function in the world, rather than how HE influences these functionings. Using graduate functionings to assess the quality and value of HE is an ontological mistake. This judges HE institutions by what graduates do, which may or may not be influenced by HE, rather than considering what HE institutions actually contribute and change. This means that HE policy risks producing inaccurate and misleading conclusions. The paper concludes by recommending how policy could adopt these foundational concepts to better assess the quality and value of HE, offering more appropriate accounts of how HE impacts graduates.

Keywords

Graduate Outcomes, Higher Education, Higher Education Policy, Critical Realism, Philosophy, Learning Outcomes

Introduction

Over the past 20 years, higher education (HE) policy has become increasingly concerned with graduate outcomes. There has been a growing focus on instrumental outcomes, for example the World Bank's HE strategy prioritises the aim of 'creating programmes that connect with the labour market' (World Bank, 2020). This concern with graduate outcomes has been particularly pronounced in Anglophone contexts, driven primarily by calls for 'value for money' by national governments in the face of increasing funding demands by the sector. For example, in order for Australian universities to be eligible for public funds, they are required to publish a statement concerning the graduate attributes they develop through their programmes (Barrie, 2006). Similarly, in the UK metrics associated with graduate employment form a key part of the Teaching Excellence and Student Outcomes Framework, a nationwide accountability exercise that attempts to measure teaching quality at higher education institutions (Tomlinson et al., 2018). This focus on graduate outcomes is also increasingly influential in other contexts, including China (Guo et al., 2019) and Kenya (McCowan et al., 2018).

The importance of graduate outcomes is reflected in the widespread use of large-scale national graduate outcome surveys. These surveys are no small undertaking for policymakers, for example the Graduate Outcome Survey is the largest annual survey in the UK (HESA, 2020). Jackson and Bridgstock (2018) analyse the ways these surveys have been

implemented in several contexts, including their various timings, outcomes and logistics. In a number of contexts, such as Australia, Ireland and the UK, these surveys are now shifting to second iterations. This reflects the high degree of policy attention in this area, leading to demands for refined metrics and data collection on a broader range of outcomes (Jackson and Bridgstock, 2018).

HE policy on graduate outcomes tends to have a number of common features. For purposes of clarity, the remainder of this section focuses on the case study of the UK, describing how these features play out in this specific context. The first feature of HE policy is that graduate outcome metrics have been incorporated into accountability regimes that aim to measure the value or quality of HE (Austin, 2019). Recent policy and rhetoric from government ministers has positioned graduate outcome metrics as 'robust' measures of programme quality and value (Department for Business, Energy & Industrial Strategy, 2020). This has been contrasted with other measures, such as student perception data from the National Student Survey. The UK government has gone as far as calling for a 'radical, root and branch review of the National Student Survey' on the basis that 'its results do not correlate well with other, more robust, measures of quality…[such as] progression to highly skilled employment' (Department for Business, Energy & Industrial Strategy, 2020). Similar dynamics can be seen in the policy detail surrounding the restructuring regime, a set of principles and practices that the UK government would adopt if a university falls into financial difficulty (Department for Education, 2020). This restructuring regime asks universities to explain how they are 'refocussing provision on high-quality courses, defined as courses with strong learner outcomes (eg low dropout rates and large proportions of

graduates finding highly-skilled employment)' (p.9), in order to have access to emergency funding (Department for Education, 2020).

The second feature of HE policy on graduate outcomes is that the operationalised metrics are often narrow and economic. Jackson and Bridgstock (2018) note that the most common measures from the surveys include: employment status, 'high-skilled' employment and salary. In the UK, the Graduate Outcome Survey measures all three of these narrow economic measures 15 months after students graduate (HESA, 2020). The salary data is also captured in an additional dataset, the Longitudinal Education Outcomes dataset, that uses tax and benefit records to track graduate salaries over the entirety of their career (Morris, 2017). The Graduate Outcome Survey does capture a broader range of data, including graduates' subjective wellbeing and their reasons for taking a particular job (HESA, 2020)—this expanded scope is a key difference with the survey's predecessor, the Destinations of Leavers from Higher Education. However, despite this collection of data on a broader range of outcomes, HE policy in the UK, such as the recent restructuring regime, continues to prioritise narrow economic metrics.

The third feature is that despite the operationalised metrics tending to be narrow and economic, the policy text and surrounding discourse often recognises the broad value of a university degree. For example, the Office for Students, the primary HE regulatory body in the UK, has a strategic objective related to graduate outcomes that states: 'All students, from all backgrounds, are able to progress into employment, further study, and fulfilling lives, and their qualifications hold their value over time' (Office for Students, 2020). This recognition that HE has a role to play in helping students lead 'fulfilling lives' goes far

beyond a narrow focus on employment and salaries. Hence, there is a tension between the operationalised metrics and the stated intensions of policymakers. From 2015, a number of projects on 'learning gain' were funded in an attempt to develop an alternative way to think about the value of HE (Evans et al., 2018). However, it is notable that this effort did not lead to a metric that has been adopted within policy—the foundational concepts introduced later in this paper help to explain why this is the case.

As graduate outcomes have gained prominence in HE policy, a diverse academic literature has developed in response. Some of this literature largely accepts the framing of graduate outcomes in narrow economic terms (Clarke, 2018). This type of literature considers how economic graduate outcomes vary between different institution types (Pigden and Moore, 2019) or subject areas (Robst, 2007). There is also a well-developed literature that takes a more critical approach. One strand of this critical literature makes a theoretical critique of HE policy's use of narrow metrics, arguing that employment is not simply caused by the quality of teaching received (Tomlinson, 2012). Other studies adopt a form of immanent critique, accepting the narrow economic framings of graduate outcomes, but highlighting the ways that disadvantaged students face additional barriers transferring their degree into the labour market (Pitman et al., 2017). There have also been some attempts to develop broader conceptualisations of graduate outcomes, whether Melanie Walker and Samuel Fongwa's (2017) work with the capability approach, or Paul Ashwin's (2020) study of the transformative effects of HE and other related approaches that consider learning outcomes (Coates and Zlatkin-Troitschanskaia, 2019; Kinzie, 2019).

Although this critical literature offers a number of powerful critiques of HE policy, much of this literature has been *in response* to policy and has so far had a limited impact (Coates and Zlatkin-Troitschanskaia, 2019). The paper adopts a different approach. Instead of centring a critique of a specific HE policy or graduate outcome metric, this paper takes a step back and asks: how should graduate outcomes be conceptualised? This task is achieved using critical realist philosophy, which results in the development of three foundational concepts: graduate functionings, graduate capabilities and graduate outcomes. Only after this theoretical work has been completed does the paper return to contemporary HE policy, applying these concepts to illuminate its key characteristics and mistakes, as well as making recommendations for future policy.

The following structure is adopted:

- The first section introduces critical realism and explains two of its key ontological conclusions: structured reality and the nature of causation.
- The second section uses these critical realist conclusions to develop three foundational concepts and illustrates their utility for conceptualising the impact of HE on graduates.
- The final section identifies the characteristics and mistakes of current HE policy, and offers recommendations for the future.

Introducing critical realism

Critical realism is a philosophical theory that provides an ontological account of the basic characteristics of reality and an epistemological theory of how we produce knowledge (Gorski, 2013). The core theory was developed by Roy Bhaskar in the 1970s, within his first

two books: *Realist Theory of Science* (2008) and *Possibility of Naturalism* (2015). In these, Bhaskar both exposes the weaknesses of two rival theories, positivism and philosophical constructivism,[1] and develops critical realism as a robust account of the basic characteristics of the world.

Questions of ontology can seem far removed from the day-to-day concerns of researchers and policymakers. However, all research and policy has an implicit ontology. Our ontological assumptions influence what we look for and how we intervene in the world. Positivists go looking for universal quantitative laws, philosophical constructivists seek collections of stories about people's experiences, and critical realists search for causal mechanisms that act as tendencies. Bhaskar (2008) argues that these philosophical theories are not all equally true, some offer more accurate accounts of the way the world actually is. If the world does not consist of universal laws, then positivism pushes us to look for things that do not exist. Similarly, if the world consists of more than just experiences, then philosophical constructivism ignores important parts of the world. If critical realism offers the theory of ontology that most closely corresponds to the basic characteristics of the world, an assumption that this paper holds, then social science research and policy should adopt its conclusions in order to avoid these ontological mistakes.[2] This is not the place to fully justify the strength of critical realism compared to its rivals, and several detailed accounts already exist (see Collier, 1994; Gorski, 2013), however the end of this section will briefly address this point.

Critical realism and structured reality

The first of critical realism's conclusions that is important for our goal of conceptualising graduate outcomes is that the world is 'structured' (Bhaskar, 2008). It is structured in the sense that there are different domains of reality. For our purposes, it is essential to distinguish between two of these domains:[3]

1. Domain of the actual – this domain contains *events* that happen in the world, eg a tree falls in a forest.
2. Domain of the real – this domain contains *causes*. Causes are the powers that make events occur, eg the disease that damages a tree's roots and causes it to fall.

Amber Fletcher (2017) uses the metaphor of an iceberg to describe the different domains of our structured world. At the top of the iceberg is the domain of the actual, which contains events. These events are the things that we can perceive around us, such as a tree falling, the number of people in full-time employment or the voting behaviour of women in an election. In contrast, at the bottom of the iceberg, underneath the water, is the domain of the real. This domain contains the causes, and it is these causes that lead to the events we see. Some of the causes of the events mentioned above could be: the disease in the tree roots, economic structures influencing employment rates, and class structures impacting voting behaviour. The iceberg metaphor helps to illustrate two aspects of the relation between the domains: causes in the domain of the real underlie events, and the causes are more challenging to study than the events. It is easier to observe events at the top of the iceberg, such as the employment rates of different groups, whereas the causes of these employment rates cannot be perceived in the same way.

Critical realism argues that not only is the world structured into the domain of the actual and the domain of the real, it also has *depth* (Bhaskar, 2008). When we develop a causal explanation of an event, this cause can be explained by another deeper cause. If different employment rates are caused by differences in peoples' skillsets, then we could ask what causes these skillset differences. One possible cause of these skillsets differences is different levels of engagement with education; engagement with education may tend to increase peoples' skillsets which then causes the observed employment rates.

Critical realism and the nature of causation

The second critical realist conclusion relevant to our purposes concerns the nature of causation. This theory explains that causes act as *tendencies* outside of scientific experiments (Bhaskar, 2008). This means that a cause does not completely determine an event, it only *tends* to lead to an outcome. This applies in natural science as much as social science. For example, if a drug was found to be 100 percent effective during a controlled medical trial, this does not guarantee that it will be effective outside the scientific experiment. We have good reason for thinking that it will tend to be effective—the chemical structure of the drug will not spontaneously change as it moves outside the trial. However, outside of experimental conditions, several causes can act at once to change, enhance or inhibit the effect of the drug. It is possible that drinking grapefruit juice inhibits the drug, but this was not identified in the trial as participants only drank bottled water. In this way, the drug will *tend* to be effective, but this is not perfectly predictable as there are many other causes in operation.

When we study the impact of a drug, we are required to offer a causal explanation of *how* the drug causes certain changes in the body and how this is impacted and influenced by other causal factors. While the statement: 'Drug A cures 97.6% of patients' might be an accurate description of how the drug has functioned in the past, this attempt to describe a universal law is an inappropriate account of causation (Collier, 1994). Firstly, it is shallow, giving us no sense of how the drug functions to cure a particular illness. Secondly, it mistakenly assumes causation involves universal laws rather than tendencies. The drug does not have a constant and universal impact, it is a tendency that could operate quite differently in different contexts with different people—just as the drug worked differently in the presence of grapefruit juice.

Table 1: Summary of critical realist terms and conclusions

Concept	**Explanation**
Domain of the actual	The domain of reality that contains events.
Domain of the real	The domain of reality that contains causes.
Event	An occurrence that happens in the world.
Cause	The underlying influence(s) on events.
Causal tendencies	Causes act as tendencies; a cause *tends* to create a certain outcome. However, as many causes can act and interact at the same time, this is not perfectly predictable.

Having outlined critical realism's account of some of the basic features of the world, let us briefly return to the consider the weaknesses of its rival theories. Positivism goes looking for universal laws of events (Bhaskar, 2008). This makes two mistakes, looking for causes in the domain of the actual not the domain of the real, and assuming that causation involves simple universal laws rather than nuanced explanations of how causes operate. Similarly, philosophical constructivism also ignores the domain of the real, avoiding discussion of causes and focussing only on experiences in the domain of the actual—this is problematic because it ignores a key part of the world (Bhaskar, 2015). Critical realism offers a theory

that better corresponds to the ontology of the world. In order for research and policy to both avoid looking for the wrong things, and to stop ignoring large parts of the world, it is necessary to adopt critical realist conclusions.

Using critical realism to develop foundational concepts

This section uses critical realist conclusions about the ontology of the world to develop three foundational concepts related to the study of HE's impact on graduates (see Figure 1). These foundational concepts are:

- Graduate functionings – what graduates do.
- Graduate capabilities – the causes of what graduates do.
- Graduate outcomes – the influence of HE, on the causes of what graduates do.

Figure 1: The three foundational concepts and their components

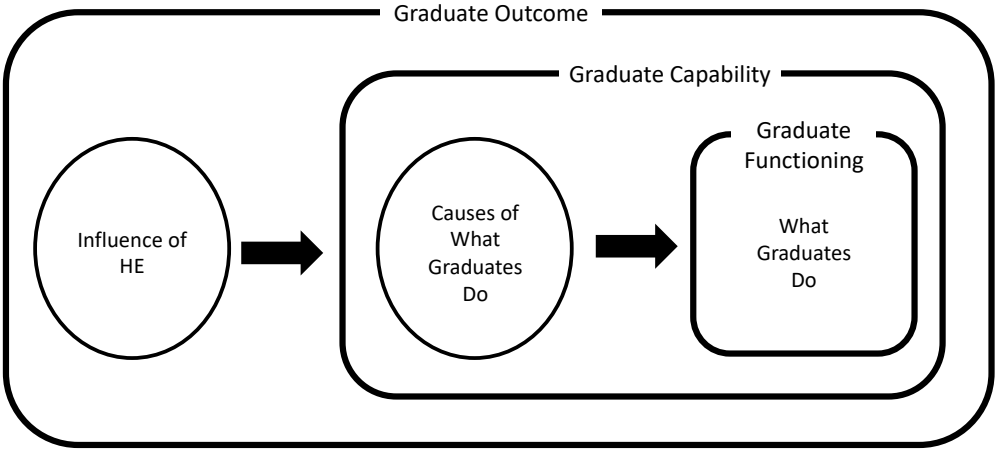

The concept of graduate functionings utilises critical realist conclusions about events. Critical realism explains that events are the things that happen in the domain of the actual—it is perfectly possible to talk about an event without discussing its underlying causes. A graduate functioning is simply an event that involves graduates. This could be the number of

graduates who are employed at a given point in time, the attitude of graduates towards lifelong learning, or the extent to which graduates adopt environmentally sustainable behaviours. There is no limit to what could be considered a graduate functioning. As long as an event involves graduates, or a sub-set of graduates, it is a graduate functioning. The question of whether some functionings are more valuable than others falls beyond the scope of this paper.

Critical realist theory demonstrates that the world consists of more than just events. To focus only on events ignores causes in the domain of the real. In order to avoid this partial focus, we need to develop other concepts that go beyond graduate functionings to consider both causes and events. This paper proposes that the cause of a graduate functioning, or the cause of what graduates do, is a *graduate capability*. For example, the higher rates of graduate employment compared with non-graduates could be caused by their skillsets. In this case, the impact of a graduate's skillset on their employment rate is a graduate capability. However, a particular functioning could be caused by a number of different graduate capabilities. An alternative capability might result from the social connections of graduates, which provide networks that can be used to gain employment. A graduate capability refers to the cause of a particular functioning, but there can be several diverse capabilities that act to cause the event.

The third concept, graduate outcomes, builds on Bhaskar's (2008) theorising about depth—causes are always caused by other causes. If a graduate capability is the cause of what graduates do, then a graduate outcome is the influence of HE on the graduate capability. In other words, a graduate outcome looks at the influence of HE on the cause of a particular

functioning—it looks at the cause of the cause. Returning to the example of graduate employment rates, a graduate outcome would seek to explain how HE influences the relevant graduate capabilities. If it is graduate skillsets that cause the patterns of employment, then a graduate outcome would involve the influence of an aspect of HE, say a problem-based pedagogy, on the skillset of graduates. HE could influence several different graduate capabilities, and these could be influenced in several different ways. This makes graduate outcomes complex, a feature that will be explained in more detail later.

It is useful to distinguish a sub-set of graduate outcomes, which this paper will term *educational graduate outcomes*.[4] The concept of graduate outcomes imposes no limit on the way in which HE influences graduate capabilities; the change could result from the teaching function of HE (eg pedagogy) or from a non-teaching function (eg extra-curricular activities). However, an educational graduate outcome refers to the influence of *HE teaching* on a graduate capability, and excludes the influence of non-teaching aspects. In this way, assessments of the influence of non-teaching aspects of HE, such as the impact of extra-curricular activities or the impact of living independently, would not constitute educational graduate outcomes. This explanation that educational graduate outcomes are a subset of graduate outcomes becomes important in the next section when we consider HE policy that attempts to measure HE teaching quality—these policies should seek educational graduate outcomes.

Now that these foundational concepts have been defined, it is possible to use critical realism to understand their nature and how we can produce knowledge about them. Through reflecting on the ontology of each concept we can answer questions like: can any

of the concepts be quantified, can we go and observe these concepts, and what form must our knowledge take?

Graduate functionings have a fundamentally different nature to graduate capabilities and graduate outcomes. Graduate functionings are events, whereas the other two concepts refer to both events and their underlying causes. As graduate functionings are events, they can often be expressed in relatively simple quantitative terms, eg '80% of graduates are employed 18 months after graduation'. It is true that there may be debates and conceptual issues when producing this description, say in deciding whether an artist working on a portfolio should be included in the 'unemployed' category. However, graduate functionings remain relatively easily identifiable through observation, albeit theoretically-informed observation, because they are events.

In contrast, graduate capabilities and graduate outcomes are of a fundamentally different nature, which has important implications for how we produce knowledge about them. These concepts refer to both events and their underlying causes. Critical realism explains that causes act as tendencies, and that events are often determined by several causes interacting in complex ways (Bhaskar, 2008; Collier, 1994). This means that the study of graduate capabilities and graduate outcomes does not involve simple observation. Instead, they require the development of causal explanations of how the events came about—this often involves the creation and application of theories about how these causes operate. For example, if we wish to discuss a graduate outcome related to employment rates, this would involve explaining how a certain aspect of HE (say, problem-based pedagogy) influences a certain graduate capability (say, the skills of graduates to work independently), and how in

turn this capability influences employment rates (say, though improving performance in employer assessment days). This cannot be simply observed, and it cannot be described in a simple way, because it is not an event. To try and express these nuanced causal explanations as simple quantitative statements in the form: 'University X led to Y% impact on Outcome Z' makes an ontological mistake. No matter how much we may wish it, only events are amenable to simple quantitative description. When we are dealing with causes, our knowledge must take the form of nuanced causal explanations of how these causes operate, and how this varies for different people in different contexts (Pawson and Tilley, 1997). This reflects that causes act as tendencies, and many causes can act at the same time. This is explained through a concrete example below.

Applying the foundational concepts to a concrete example

This paper has developed three foundational concepts using critical realist philosophy. This helps us to distinguish events (graduate functionings), from the causes of these events (graduate capabilities), and the influence of HE on these causes (graduate outcomes). These concepts enable researchers and policymakers adopt an appropriate ontological position, helping them avoid searching for things that do not exist or ignoring large parts of the world. The concepts are useful as well as necessary. Below, their utility is demonstrated by applying them to the topic: the influence of HE on graduates' attitudes to environmental sustainability. This example is hypothetical, dealing with fictionalised capabilities and outcomes, rather than assessing the actual literature and evidence base on this topic. This hypothetical approach enables the focus to be on illustrating the three concepts, rather than being diverted by questions related to the current evidence base.

Figure 2: Applying the concepts to understand the influence of HE on graduates' environmental attitudes

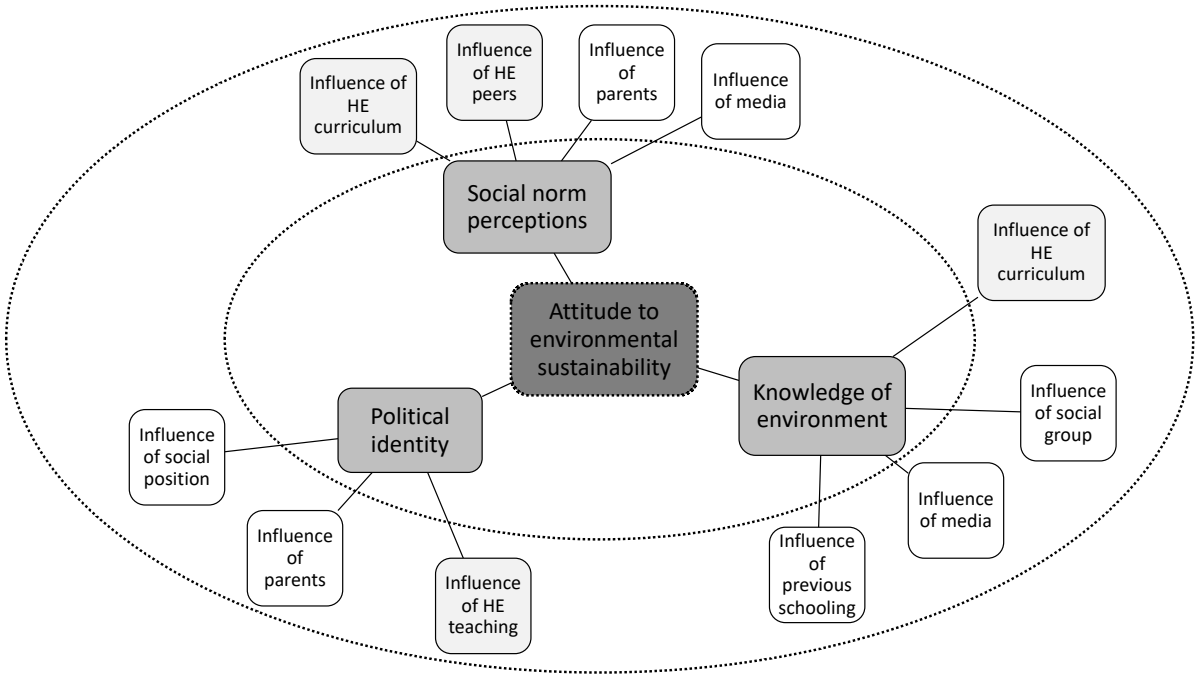

When approaching this research topic, the foundational concepts help to identify different aspects of the problem. These aspects correspond to the three layers in Figure 2, with the graduate functioning in the centre, graduate capabilities in the middle layer, and graduate outcomes in the outer layer. The concept of graduate functionings pushes us to consider which events are relevant to this topic area. This example focuses on one particular functioning: graduate attitudes to environmental sustainability, and how these attitudes vary for different sub-populations.

The concept of graduate functionings also helps us to avoid misinterpreting these patterns in graduate attitudes as directly revealing the influence of HE. Graduate functionings are events, and make no attempt to explain underlying causes. Graduate attitudes to the environment could be caused by HE, but equally they could be caused by something unrelated. It is plausible that graduates may have read more about the environment in the

media, and it is this, rather than any cause related to HE, that influences their attitudes to the environment. In this way, it would be a mistake to assume that a graduate functioning tells us about the influence of HE. If we wish to understand the influence of HE on these attitudes we must go beyond descriptions of these graduate functionings to consider causes.

The concept of graduate capabilities pushes us to consider the causes of the observed patterns of graduate attitudes. The middle layer of Figure 2 depicts several possible capabilities that could cause the graduate attitudes. For example, one possibility is that graduate attitudes are caused by their level of environmental knowledge—graduates, or some sub-populations of graduates, might have greater knowledge of the climate crisis, which causes them to hold more environmentally sustainable attitudes. If this is true, the graduate capability refers to the knowledge graduates have about the environment, and the causal impact this has on attitudes. This is by no means the only possible causal explanation, a number of different graduate capabilities could underlie a particular graduate functioning. Figure 2 details how it is possible that the environmental attitudes of graduates do not stem from greater knowledge, but are instead caused by perceptions of the social norms. Graduates might tend to adopt sustainable attitudes because they perceive this as the socially acceptable thing to do, rather than because of any substantial environmental knowledge.

This highlights that for a given graduate functioning there are a number of possible graduate capabilities that could cause the event to occur. In fact, a graduate functioning is likely to be caused by several of these capabilities at once. Graduate attitudes to the environment are

likely caused by graduates' environmental knowledge, their perceptions of social norms, and their political identity, to name only three possible underlying capabilities. These capabilities may not only act at the same time, but they could also interact with each other in complex ways. Although perceptions of social norms might tend to influence graduates' attitudes to environmental sustainability, this may not apply to graduates with a high level of knowledge about the environment. It is possible that the attitudes of these 'high-knowledge' graduates is less influenced by social norms, and instead their attitudes are determined largely by their knowledge base. In this way, while perceptions of social norms may tend to influence graduates' attitudes to the environment, this might not occur for this specific subpopulation. Therefore, to understand the various graduate capabilities that underlie a particular functioning, we need to produce nuanced explanations of how these capabilities operate, when they have an impact, and how they interact with each other. This cannot be captured by a shallow positivistic statement in the form 'Perceptions of social norms predict 75% of the variation in environmental attitudes'. As was explained in above, this fails to explain how environmental attitudes are caused, and misrepresents causes as universal laws rather than tendencies.

To understand the influence of universities on promoting graduates with sustainable environmental attitudes, we need to go even further and consider how HE influences these capabilities. This is the outer layer of Figure 2. For example, a particular HE curriculum could influence graduates' perceptions of social norms, which then causes their attitudes to environmental sustainability. However, this outer layer that involves graduate outcomes is complex. Firstly, it is complex because different aspects of HE can influence the same capability, eg both HE curricula and HE peer groups could influence graduates' perceptions

of social norms. Secondly, not only can different aspects of HE influence the same capability, HE can influence multiple different capabilities. It is possible that HE curricula influences graduate capabilities associated with environmental knowledge *and* perceptions of social norms. We are left with a complex picture in which several aspects of HE can influence several graduate capabilities, which can then interact in complex ways to produce the resulting graduate functioning.

However, Figure 2 highlights that assessing graduate outcomes has yet another element of complexity. Many of the 'causes of causes' in the outer layer of Figure 2 have nothing to do with HE. Graduates' greater knowledge of the environment, may be caused by curricula content from secondary school or from information received from the media. It is possible that the observed graduate functionings are not at all influenced by the fact that graduates attended HE, they could be fully determined by these non-HE related causes. To observe a graduate functioning, and even to explain how a graduate capability causes this functioning, tells us nothing about the role of HE.

Even when an aspect of HE does influence a graduate functioning, non-HE causes may interact with this in complex ways. For example, even if we have explored how a particular HE curriculum influences graduates' knowledge about the environment, it is possible that non-HE causes act to complicate this picture. Imagine if one social group tends to have less trust in scientific institutions. For this social group, it is possible that even if the HE curriculum does lead to greater scientific knowledge about the environment, this may not then translate to more concerned attitudes because these graduates tend to distrust the institutions that produce this knowledge. If this is the case, we would need to explain why

this social group tends to demonstrate greater levels of scientific distrust. In this way, to make a statement about a graduate outcome is to offer a nuanced causal explanation of how HE influences a graduate capability, or a number of different graduate capabilities, as well as to explain how this influence of HE interacts with other non-HE causes of graduate capabilities. Given the ontology of the world, the study of graduate outcomes is inevitably complex.

The key takeaway is that the influence of HE on any graduate functioning is complex. The three concepts help us understand why this is the case. When we talk about graduate outcomes, we are talking about the causes of causes of a functioning. This is compounded by the fact that many causes can operate at the same time, and they can interact with each other in complex ways. However, the concepts do help us navigate our way through this complexity. They help us know what to look for: we must identify and distinguish events (graduate functionings), the causes of these events (graduate capabilities), the influence of HE (graduate outcomes), and the influence of non-HE related factors. The concepts promote a view of knowledge that corresponds to Pawson and Tilley's (1997) explanation that our knowledge should explain how things work, who they work for, and in what circumstances. Applied to our topic area, this involves understanding how HE influences capabilities that in turn cause particular functionings, and how this varies for different graduates, in different circumstances. Necessarily, this involves rich and nuanced causal explanations.

Implications for HE policy

This section considers the implications of the three foundational concepts for HE policy in this area, focussing specifically on two aspects of this policy: national graduate surveys and

assessments of HE quality and value. The concepts help clarify the key characteristics and mistakes of contemporary policy. They also suggest ways that future policy could overcome these issues and develop more appropriate assessments of HE quality and value.

Characteristics and mistakes of current HE policy

The beginning of this paper described the recent trend to conduct national surveys of graduates (Jackson and Bridgstock, 2018). The outputs from these surveys then go on to play a key role in HE policy. In the UK, metrics on graduate employment are used to both assess quality in the Teaching Excellence and Student Outcomes Framework (Tomlinson et al., 2018) and to identify low-quality courses in the recent restructuring regime (Department for Education, 2020). But, what *are* the outputs of these surveys? Using the foundational concepts, it is possible to see that the surveys produce graduate functionings. They produce descriptions of events involving graduates, such as their employment rates, salaries or sense of personal wellbeing.

Recognising that these national graduate surveys produce graduate functionings helps us understand both the importance and the limitations of these surveys. It is essential to know how graduates function in the world, and how this may vary for different sub-populations. Without this, we cannot begin to assess the impact of HE on graduates or the ways that we could develop HE to better promote valuable functionings for all graduates. However, we must remember that graduate functionings are descriptions of events. By themselves, they tell us nothing about causes—we cannot assess the influence of HE by looking at a graduate functioning. Instead, graduate functionings provide hypotheses and guidance for future causal research (Danermark et al., 2001). For example, if we observe different employment

rates for graduates between two universities, this *could* suggest that one university is having a greater impact. However, this can only be understood by explaining if and how one university has this impact on graduate capabilities, and how this influences the employment rates. To reiterate, graduate functionings do not refer to causes, but they can highlight patterns and differences between subpopulations that then guide future research.

A problem arises when HE policy uses graduate functionings from these national surveys *but misinterprets them as graduate outcomes*. This is seen in policies that seek to assess the quality or value of HE. It is a mistake to use graduate functionings in this way. We have seen how a graduate functioning could be caused by something completely unrelated to HE, such as the social background of students influencing their employment rates (Clarke, 2018). If employment rates are used to assess HE quality, then HE institutions with disproportionally disadvantaged student intakes may be labelled 'low-quality', even if they offer a very high quality education (Tomlinson et al., 2018). This use of graduate functionings to conclude about the quality and value of HE is problematic because it gives no indication of why one university has a greater employment rate, hindering our ability to identify and learn from best practice about how HE can influence graduate capabilities.

A related issue is that the policies assessing HE quality and value often seek simple comparisons between different HE institutions or degree programmes, sometimes in order to provide rankings. However, this is an ontological impossibility. As much as we may desire simple and quantitative ways to compare HE quality and value, the ontology of the world does not allow this. Finding a graduate functioning that 'Employment rates from University A and B are 90% and 80% respectively' does not demonstrate that University A is better

than University B. This simple move from a graduate functioning to an assessment of quality or value does not address whether HE caused this change, the way that HE brought about this change and how non-HE causes influenced this functioning. If these policies seek to drive quality and enhance value, it is ironic that they show such little interest in what exactly it is that HE can do to bring about this quality and value. Proposals for how this could be better achieved, are given below.

A further characteristic revealed by the foundational concepts is the overly narrow economic focus of much HE policy. This is not a novel insight, complimenting the work of other scholars (Tomlinson, 2012; McCowan et al., 2018), but this paper arrives at this from a more philosophical perspective. There is nothing in the nature of the foundational concepts that justifies a narrow economic focus. The concept of graduate functionings is open to the ways graduates function in all spheres of life, not just the economy. Similarly, graduate capabilities and outcomes are not only interested in the causes of economic functionings, but any functioning of graduates that we deem valuable. If policymakers wish to continue to prioritise narrow economic functionings, the foundational concepts place the burden on policymakers to justify this approach. To be reasonable, they would have to evidence why they assume HE only influences economic functionings or why only these economic functionings are valuable. This shifts the burden of proof from those advocating for broader graduate functionings, to those pushing narrow economic ones.

Recommendations for future policy

The critical realist-informed concepts not only help to identify the characteristic features and weaknesses of these HE policies, they also suggest how future policy could be reformed

to avoid these mistakes. Firstly, they demonstrate the importance of national surveys of graduates. These national surveys should continue to be run, in order to gather important information about graduate functionings. If there are valuable functionings beyond employment status and salary, then efforts should be made to broaden these surveys and other forms of data collection. This broadening has already been seen, at least to some extent, within the Graduate Outcome Survey in the UK (HESA, 2020). Although it is beyond the scope of this paper to justify which broader graduate functionings are valuable, it is clear that there is nothing ontological within the foundational concepts that justifies the disproportionate focus on economic graduate functionings.

Secondly, HE policy that uses graduate functionings within accountability regimes to assess the quality and value of HE, such as the Teaching Excellence and Student Outcomes Framework (Office for Students, 2021) and the recent restructuring regime that targets 'low-value' courses (Department for Education, 2020), should also be reformed. Assessing the quality and value of HE is an important task, allowing the evaluation of how HE influences graduates, whether some subpopulations of graduates gain more than others, and how it could be refined to better promote valuable functionings. However, in order to achieve this, policies must stop misrecognising graduate functionings as graduate outcomes, and abandon the related desire to find simple comparisons between different universities or programmes (Tomlinson et al, 2018).

Instead of graduate functionings, policies should seek graduate outcomes. If a policy seeks to assess teaching quality, then it is not just graduate outcomes but educational graduate outcomes that should be sought. This involves the creation of nuanced causal explanations

about how certain aspects of HE, say a particular pedagogy, influences graduate capabilities, which then contributes to the observed patterns of graduate functionings. Put another way, the goal of accountability regimes should shift from using graduate functionings to make ontologically problematic arguments about value or quality, to detailed rigorous research into how certain forms of HE influence graduates, and how in turn this influences the way they function later in life. This is no simple matter. The influence of HE will be impacted by many non-HE causes, such as social background and previous education, and HE may influence several different graduate capabilities. However, as a collective effort, research could come to understand how various aspects of HE influence these graduate capabilities, when this influence tends to occur, and which graduates this tends to impact (Pawson and Tilley, 1997). A sector armed with these causal explanations would be in a position to intervene and guide HE to have a greater impact on valuable graduate functionings for everyone. Although this task is complex and time consuming, this approach is necessary, given the ontology of the world.

Compliance with Ethical Standards

Conflict of interest. The author declared no potential conflicts of interest with respect to the authorship and publication of this article.

Notes

1. This paper uses the term ‘philosophical constructivism’ to refer to a philosophical theory that holds an irrealist ontology. Specifically, one that denies causes exist. Although, ‘constructivism’ is often used in social science to refer to any research that

recognises knowledge is theory-dependent, this idea is compatible with both critical realism and philosophical constructivism.

2. Critical realist philosophy is compatible with many different social science theories. Although some social science theories may be excluded as ontologically problematic, for example those that do not allow for individual agency, critical realism does not help to select from the range of compatible theories. This would be the role of social science, not philosophy.
3. Critical realism actually argues there are three domains of reality: empirical, actual and real (Bhaskar, 2008). The empirical contains events *experienced* by agents, whereas the actual contains all events whether experienced or unexperienced. This distinction is less relevant to the purposes of this paper.
4. There is a similarity between an 'educational graduate outcome' and a 'learning outcome', which is used in some parts of the literature (Coates and Zlatkin-Troitschanskaia, 2019). However, there is some ontological ambiguity in the latter term. A learning outcome can be used to describe a graduate functioning, eg what graduates know. Equally, it could refer to an educational graduate outcome, eg assessing the impact of HE teaching on a particular graduate capability.

References

Archer, M., (2008) *Realist Social Theory: The Morphogenetic Approach*, Cambridge: Cambridge University Press.

Ashwin, P., (2020) *Transforming University Education: A Manifesto*, London: Bloomsbury Academic.

Austin, I., (2019) 'Governance Architectures in Higher Education Arising from Extra Information on Learning Outcomes', *Higher Education Policy* 32(4): 617–637.

Barrie, S.C., (2006) 'Understanding What We Mean by the Generic Attributes of Graduates', *Higher Education* 51(2): 215–241.

Bhaskar, R., (2008) *A Realist Theory of Science*, 3rd edition, London: Verso.

Bhaskar, R., (2015) *The Possibility of Naturalism: A Philosophical Critique of the Contemporary Human Sciences*, 4th edition, Oxford: Routledge.

Clarke, M., (2018) 'Rethinking graduate employability: the role of capital, individual attributes and context', *Studies in Higher Education* 43(11): 1923–1937.

Coates, H., & Zlatkin-Troitschanskaia, O., (2019) 'The Governance, Policy and Strategy of Learning Outcomes Assessment in Higher Education', *Higher Education Policy* 32(4): 507–512.

Collier, A., (1994) *Critical Realism: An Introduction to Roy Bhaskar's Philosophy*, London: Verso.

Danermark, B., Karlsson, J.C., Jakobsen, L., & Ekstrom, M., (2001) *Explaining Society: An Introduction to Critical Realism in the Social Sciences*, London: Routledge.

Department for Business, Energy & Industrial Strategy, (2020) *Reducing bureaucratic burden in research, innovation and higher education*, London: UK Government.

Department for Education, (2020) *Establishment of a Higher Education Restructuring Regime in Response to COVID-19*, London: UK Government.

Evans, C., Howson, C.K., & Forsythe, A., (2018) 'Making sense of learning gain in higher education' *Higher Education Pedagogies* 3(1): 1–45.

Fletcher, A.J., (2017) 'Applying critical realism in qualitative research: methodology meets method', *International Journal of Social Research Methodology* 20(2): 181–194.

Gorski, P.S., (2013) 'What is Critical Realism? And Why Should You Care?', *Contemporary Sociology* 42(5): 658–670.

Guo, F., Luo, Y., Liu, L., Shi, J., & Coates, H., (2019) 'Analysing Mechanisms for Evaluating Higher Education Outcomes in China', *Higher Education Policy* 32(4): 557–575.

HESA, (2020) *About the survey*. Graduate Outcomes, https://www.graduateoutcomes.ac.uk/about-survey, accessed 11 October 2020.

Jackson, D., & Bridgstock, R., (2018) 'Evidencing student success in the contemporary world-of-work: renewing our thinking', *Higher Education Research & Development* 37(5): 984–998.

Kinzie, J., (2019) 'Taking Stock of Initiatives to Improve Learning Quality in American Higher Education Through Assessment', *Higher Education Policy* 32(4): 577–595.

McCowan, T., Oanda, I., & Oketch, M., (2018) 'Towards a National Graduate Destinations Survey in Kenya: An Exploratory Study of Three Universities', *Higher Education Policy* 31(1): 97–119.

Morris, D., (2017) 'A beginner's guide to Longitudinal Education Outcomes (LEO) data', *Wonkhe*. https://wonkhe.com/blogs/a-beginners-guide-to-longitudinal-education-outcomes-leo-data/, accessed 27 November 2020.

Office for Students, (2020) *Outcomes performance measures*, Office for Students. https://www.officeforstudents.org.uk/about/measures-of-our-success/outcomes-performance-measures/, accessed 1 June 2020.

Office for Students, (2021) *About the TEF*, Office for Students. https://www.officeforstudents.org.uk/advice-and-guidance/teaching/about-the-tef/, accessed 5 February 2021.

Pawson, R., & Tilley, N., (1997) *Realistic Evaluation*, London: Sage.

Pigden, L., & Moore, A.G., (2019) 'Educational advantage and employability of UK university graduates', *Higher Education, Skills and Work-Based Learning* 9(4): 603–619.

Pitman, T., Roberts, L., Bennett, D., & Richardson, S., (2017) 'An Australian study of graduate outcomes for disadvantaged students', *Journal of Further and Higher Education* 43(1): 1–13.

Robst, J., (2007) 'Education and job match: The relatedness of college major and work', *Economics of Education Review* 26(4): 397–407.

Tomlinson, M., (2012) 'Graduate Employability: A Review of Conceptual and Empirical Themes', *Higher Education Policy* 25(4): 407–431.

Tomlinson, M., Enders, J., & Naidoo, R., (2018) 'The Teaching Excellence Framework: symbolic violence and the measured market in higher education', *Critical Studies in Education* 61(5): 627-642.

Walker, M., & Fongwa, S., (2017) *Universities, Employability and Human Development*, London: Palgrave Macmillan UK.

World Bank, (2020) *Tertiary Education*. World Bank. https://www.worldbank.org/en/topic/tertiaryeducation, accessed 19 November 2020.