The London School of Economics and Political Science

Laws in the Social Sciences

Catherine Greene

A thesis submitted to the Department of Philosophy, Logic and Scientific Method of the London School of Economics for the degree of Doctor of Philosophy, London, December 2017

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 67,455 words.

Abstract

The social sciences are often thought to be inferior to the natural sciences because they do not have laws. Bohman writes that "the social sciences have never achieved much in the way of predictive general laws—the hallmark of naturalistic knowledge—and so have often been denied the honorific status of 'sciences'" (1994, pg. vii). Philosophers have suggested a number of reasons for the dearth of laws in the social sciences, including the frequent use of *ceteris paribus* conditions in the social sciences, reflexivity, and the use of 'odd' concepts. This thesis argues that the scarcity of laws in the social sciences is primarily due to the concepts that social scientists often work with. These concepts are described as Nomadic and are characterised by disagreement about what can reasonably be included within the scope of a concept. The second half of the thesis explores the implications of this analysis. It argues firstly, that counterfactual analysis is problematic when using Nomadic concepts. Secondly, it argues that taking an intentional perspective on behaviour often involves the use of Nomadic concepts so, if social scientists do hope to formulate laws, then they are more likely to succeed if they focus on behaviour that is not intentional.

Index

Summary

Chapter 1 discusses three methodological problems that are taken to explain why the social sciences have had considerable difficulty in finding laws. It argues that these methodological problems, which have traditionally been seen as separate, are in fact a consequence of the sorts of 'odd' concepts that social scientists often use. The concepts used by social scientists are thought to be 'odd' because it is often difficult to give necessary and sufficient criteria for inclusion within a concept. This is extensively discussed by Bradburn (2011), Little (1993), and Crasnow (2015) among others. The concepts used in the social sciences are, it is argued, not as precise as paradigmatic scientific concepts, such as 'length', or 'mass'. This is not a controversial point. It is also uncontroversial to note that the social sciences often face extensive *ceteris paribus* clauses, and reflexivity. The contribution of this chapter is to argue that these problems, the use of 'odd' concepts, extensive *ceteris paribus* clauses, and reflexivity are inter-related. Specifically, reflexivity is one reason why concepts used in the social sciences are often 'odd', and the existence of extensive *ceteris paribus* clauses is a consequence of using these 'odd' concepts.

Given the importance of this 'oddness' of the concepts social scientists often use, a number of explanations of it are discussed. These are whether the concepts used in the social sciences are essentially contested, whether they are vague, and whether they are multiply realisable. It will be argued that, while some concepts fit all these descriptions, these descriptions do not fully account for the 'oddness' of many concepts social scientists use. Concepts used in the social sciences have been described in a variety of ways, and the philosophers doing so invariably use different examples when doing so. Examples of essentially contested concepts include 'democracy' and 'work of art', examples of multiply realisable concepts include 'money' and 'markets'. The question posed by this is how we know which of these categories a specific concept belongs in. Woodward (2016) argues that we need a way of deciding whether concepts are good ones for using in analysis. Chapter 1 concludes by arguing that, while broadly correct, Woodward's proposal is not well suited to the social sciences.

Chapter 2 proposes a new way of thinking about concepts used in the social sciences. It will be argued that the various ways in which social science concepts have been described; that they are essentially contested, that they are multiply realisable, and

that they are vague can all be subsumed within this framework. Concepts used in the social sciences are described as Nomadic. The implications of accepting this Nomadic framework are firstly, that this provides an *a priori* way of deciding whether concepts are useful for social science research; secondly, it explains why the social sciences have had so much difficulty discovering regularities; and thirdly, that current proposals for making concepts more precise are unsatisfactory.

Chapter 3 demonstrates the implications of accepting the framework for analysing social concepts outlined in Chapter 2. The structure of the concepts used in any particular case affects the viability of counterfactual analysis. Specifically, it is argued that Woodward's (2003) idea of a 'hypothetical experiment' is difficult to apply to many social science cases. This is because, often, social science concepts allow different social scientists to construct multiple, plausible, but inconsistent counterfactual scenarios, which there is little principled way of choosing between. This makes it difficult to agree about the truth value of a counterfactual.

Chapter 4 begins with the observation that taking an intentional perspective on behaviour usually means explaining behaviour in terms of beliefs and desires. These beliefs and desires are often expressed using concepts that are Nomadic, which means that there will often be a number of plausible explanations of behaviour which it is difficult to decide between. One possible response to this is to seek to redescribe beliefs and desires in other terms. McIntyre's (1996) argument in support of this is reviewed, but found unconvincing. The second part of the chapter analyses an example of regular behaviour- Middle bias, with the aim of understanding whether it is intentional. It will be argued that middle bias behaviour does not fit within the usual understanding of what intentional action is. It is described as 'derivatively intentional' because this behaviour is only intentional because of wider behaviour, of which it is a part. This chapter concludes that behaviour that is 'derivatively intentional' is describable using concepts that are less Nomadic.

Chapter 5 deals with a number of clarifications and possible concerns. These are firstly, the possible implications the Nomadic framework has for social policy. Secondly, showing that less Nomadic concepts from earlier chapters may be derivatively intentional. Finally, suggesting other possible examples of derivatively intentional behaviour.

Chapter 1: Introduction

Outline

This chapter begins with the observation that the social sciences are often thought to be inferior to the natural sciences. One of the main justifications for this view is the dearth of laws in the social sciences. Three possible reasons for the lack of laws in the social sciences are then discussed. These are the existence of extensive *ceteris paribus* clauses, reflexivity, and the use of concepts that are 'odd' when compared to the concepts natural scientists often use. I argue that understanding the concepts used by social scientists is critical, firstly, because reflexivity partly explains why many concepts are 'odd' and secondly, because extensive *ceteris paribus* clauses are to be expected when social scientists are using 'odd' concepts. The second part of the chapter argues that none of the ways in which the concepts social scientists use have been characterised is satisfactory, and suggests that a new way of understanding these concepts is needed. Finally, this chapter shows that the methodological problems faced by the social sciences; the existence of extensive *ceteris paribus* clauses, reflexivity, and multiple realisability, are a consequence of the 'oddness' of the concepts that social scientists often use.

1. **Why do people think that the social sciences are inferior to the natural sciences?**

It is often suggested that the social sciences are inferior to the natural sciences. Fritz Machlup (1994, pg. 5) correctly notes that deciding why this is the case depends on how we interpret 'inferior', and what we consider the 'social sciences' to be. He discusses nine points of comparison between the natural and social sciences, but I will focus on one, which is the purported absence of laws in the social sciences, because this is often seen as the most important point of comparison between the social and natural sciences. Illustrating this, Bohman writes that "the social sciences have never achieved much in the way of predictive general laws—the hallmark of naturalistic knowledge—and so have often been denied the honorific status of 'sciences'" (1994, pg. vii).

The absence of laws is judged to be a critical mark of inferiority because of the central role that laws play in explanation, at least according to Hempel's extraordinarily influential Covering Law model of explanation (Hempel, 1962/ 1998) [1]. Salmon *et al* write that this model is "central to all discussions of the subject" of explanation (Salmon *et al*, 1999, pg. 14), while Staley writes "The discussion of explanation within contemporary philosophy of science […] really begins with Carl Hempel and Paul Oppenheim's *Deductive-Nomological* (D-N) model of explanation" (2014, pg. 199. Italics in original). Hempel argued that a successful explanation of an event, in the natural or social sciences, shows that this event was to be expected, given particular facts of the world prior to the occurrence of the event, and general laws that show how the prior facts led to the event. This view has been influential in the natural and social sciences. For example, McIntyre says that it is the 'intuition' behind Hempel's covering law model that has inspired explanation, and the way in which laws are seen, in the social sciences (McIntyre, 1996, pg. 4). If the social sciences do not have laws, therefore, they are inferior to the natural sciences because they cannot explain events.

Nevertheless, despite the influence of this approach, the applicability of the Covering Law model has been questioned in the natural sciences (see, for example, Salmon *et al*, 1999, Ch 1). Part of this criticism hinges on the difficulty with specifying exactly what a law is. Illustrating this, Roberts writes that "The concept of a law has been terribly difficult to explicate, and the literature on laws is rife with disagreement" (2004, pg. 153). Some philosophers, such as Hume and Ayer (Ayer, 1956/1998) argue that laws are true universal generalisations. 'All cats are mammals' is an example of a true universal generalisation. This is perhaps the most intuitive description of what a law is. However, this view is problematic because it does not exclude vacuous laws—laws that refer to things that don't exist. An example is 'All hobbits are under three feet tall'. This account also makes it difficult to distinguish between laws and accidental generalisations. The following two statements are true universal generalisations: 'there are no gold spheres a mile wide' and 'there are no uranium spheres a mile wide'. However, there is a reluctance

[1] Hempel also identified a form of probabilistic explanation, less frequently mentioned in discussions of Hempel's philosophy of science, known as the inductive-statistical explanation. In these cases, the prior facts about the world, combined with a statistical law, or generalisation, make the event we are seeking to explain very likely. He defines a statistical probability as, "roughly speaking, the long-run relative frequency with which an occurrence of a given kind... is accompanied by an 'outcome' of a specified kind" (1962/1998, 689).

to accept the statement about gold spheres as a law because there is no reason why this has to be the case. It seems entirely accidental. Uranium, on the other hand, cannot be combined in such quantities. However, defining laws more precisely has proven difficult. Some philosophers have sought to explicate a necessary relation that they take to exist between the predicates of laws (for example, see Dretske 1998, Armstrong 1983). Some philosophers even deny that laws play a role in successful explanation in the natural sciences (for example, see Cartwright 1998, Van Fraassen 1980, Giere, 1999).

Roberts provides a pragmatic way through this impasse by finding a point of common agreement within this dispute. This agreement is about regularities. Roberts distinguishes between three kinds of regularities:

> **Strict regularities:** Universally quantified conditionals, holding throughout the universe. Such as 'All copper conducts electricity'.
>
> **Statistical regularities:** Characterise an unrestricted domain.
>
> Such as 'Any atom of uranium-238 has a probability of 0.5 of decaying within any time-interval of 4.5 billion years'
>
> **Hedged, or ceteris paribus, regularities:** Regularities that are qualified by admitting that they have exceptions in various circumstances (Roberts, 2004, pg. 153-4).

Roberts writes that "Most philosophers who have written on the topic of laws of nature allow that laws may be (or may entail) either strict or statistical regularities" (2004, pg. 154). Some also allow that they might be hedged regularities (2004, pg. 154). So, this provides a useful starting point for the social sciences because, if there are no regularities in the social sciences, it is highly unlikely that there will be laws. However, he says, a regularity is not necessarily a law. The fact that there are no gold spheres with a radius of a mile is a strict regularity, and it is most probably true. Nevertheless, this is not a law because, Roberts says "No law that we have any inkling of rules out gargantuan spheres of gold; it just happens that to be the case that there are no such spheres" (Roberts, 2004, pg. 154). But what additional features a regularity must have in order to count as a law is unclear (See Roberts, 2004, 154-157 for an overview). Nevertheless, if the social sciences have regularities, then there are, potentially, laws. This thesis therefore discusses 'regularities' rather than 'laws'.

Nevertheless, even if we talk of 'regularities', the social sciences have, so far, provided very few examples. Why is this? Three reasons have often been given for the inferiority, or peculiarity of the social sciences. These three reasons are:

i. The reflexivity of social sciences (see Cartwight 2007, Hacking 1995, Day 2012)

ii. The necessity of *ceteris paribus* condition in the social sciences (see Scriven 1994, McCauley, 2004)

iii. The use of 'odd' concepts (see Bradburn 2011)

The following section discusses each of these in turn, and concludes that the use of 'odd' concepts is of primary importance. I argue that this is because reflexivity is one reason why the social sciences often deal with 'odd' concepts, and that the necessity of *ceteris paribus* conditions is a consequence of the 'oddness' of many concepts that social scientists use. In other words, if we understand the 'oddness' then the other two reasons are subsumed within this. The following section therefore concludes that understanding why concepts that social scientists use are 'odd' is worthy of further investigation.

i. Reflexivity

Cartwright provides a nice summary of the 'reflexivity' of the social sciences. She says, "people change in response to the way we study them, the way they conceive themselves, or in reaction to what they think will happen" (2007, pg. 40). As this quote highlights, there are two aspects to reflexivity. The first is people reacting to the ways in which they are classified. The second is people acting in anticipation of a law. Both of these pose a problem for finding regularities in the social sciences because behaviour changes, and may change in a way that makes it no longer regular.

The first aspect of reflexivity is described by Hacking. People change the way they see themselves when they become aware of what they believe is an accurate generalisation. Hacking describes human kinds as classifications of types of people that we use in an attempt to change these types of people or their living conditions. He says that we can only 'improve the person' if we know what kind of person we are dealing with (Hacking, 1995, pg. 351). Human kinds are ways of classifying people that are used in the social sciences, and Hacking does not provide an exact definition. He does illustrate what he means when he writes that these are kinds about which social scientists would like to have systematic, reliable knowledge that can be used to predict behaviour. Pregnant teenagers, child abusers, and those suffering from multiple personality disorder, are examples of human kinds. The important thing about human kinds, according to Hacking, is that when we apply a kind to a person it not only changes the way we view that person, and how we act towards them. Additionally, it may also change the way the person views themselves, which may lead them to behave in new ways, such as behaving in the way that they believe a person of that kind behaves, or conversely in such a way that the 'kind' no longer applies to them. As people of a kind change their behaviour this, in turn, leads social scientists to change the way they think about the kind. This process generates the so-called "looping effect" of human kinds (Hacking, 1995, Ch. 12). The greater the moral connotations of a kind, the greater the potential for this looping effect because people are more likely to change their behaviour in response to a classification that labels them as particularly praiseworthy or blameworthy.

Day (2012) discusses the second aspect of reflexivity. He says that generalisations about human behaviour can be 'used'. For example, we can improve our ability to appear trustworthy by acting in accordance with the 'law' that looking at people directly when talking to them is an indicator of trustworthiness. Adopting the behaviour to disguise untrustworthiness means that the law no longer holds. Day's worry does not just apply to individual behaviour. He says that becoming aware of a law stating that "money increases the chances of political power" (Day, 2012, pg. 67) may motivate changing the system upon which this law depends. This is different to the natural sciences where expectations about laws do not affect how a law behaves in a particular situation.

There are two things going on here. Firstly, when people believe that the consequences of some policy or action are regular, or predictable, they can change their behaviour in the light of these expectations, which can make the effects of a policy unpredictable. Secondly, people often respond to the terms that are applied to them, which in turn changes the way they are analysed. If these terms form part of a law, this suggests that the regularity expressed may not apply as behaviour changes in the future. Further evidential support for these two phenomena is found in Bergenholtz and Busch, who write that "Various studies seem to show that exposure to economics changes how economics students behave in social science labs" (Bergenholtz & Busch, 2016, pg. 29).

They also describe a 'self-falsifying' theory, as follows, "a recent study on racial bias among basketball referees shows that the racial bias disappeared after an academic study and media attention highlighted the problem" (Bergenholtz & Busch, 2016, pg. 27). The reflexivity phenomenon appears relatively well confirmed empirically.

This is not a feature of the natural sciences because many natural phenomena, such as electrons, or planets, or chemical elements, do not change as a result of generalisations that are made about them. On the face of it, this suggests a significant difference between the social and natural sciences. McIntyre says this appears to make the possibility of finding laws remote, because relationships between variables are unlikely to remain constant for long. "Human systems […] are not stable enough over time for us to understand the relationship between underlying variables, due to the constant realignment in the network of underlying influences." (McIntyre, 1996, pg. 21). Steuer says the "single greatest difference" between the social and natural sciences is that "society in all its aspects is, or may be, changing" (2003, pg. 36). This is in contrast with the natural sciences, he says, because the ageing of the universe, for example, "is slow for all practical purposes". Even biology, where change occurs much more quickly, remains subject to continuing "underlying principles" whereas "the social universe is a place of deep seated and of superficial change". (Steuer, 2003, pg. 35).

McIntyre responds to the problem of reciprocal complexity, or reflexivity, by arguing that the complexity of phenomena depends on the level at which we describe it. In other words, complexity is derived from our description of the phenomena, rather than inherent in the phenomena itself. Even if we grant that at some level social phenomena may be too complex for us to discover laws, this does not mean that the phenomena will be too complex at any level. He cites Scriven's claim that if physicists sought to explain how a particular leaf falls from a particular tree then physics would be in a similar state to social science in relation to laws. He says that "the possibility of social science laws is crucially dependent upon how the phenomena are framed by our level of engagement with them" (McIntyre, 1996, pg. 23). In a later paper, he elaborates on what 'levels of engagement' are, suggesting that levels are synonymous with "modes of explanation" (McIntyre, 2000, pg. 107). He gives the example of a suicide, which is amenable to a number of different explanations "depending on the level at which the questioner was enquiring about the behaviour" (McIntyre, 2000, pg, 108). These explanations are a psychological one (which considers how the man saw himself), explanation by his neighbours (which considers traits he may have inherited from his father), a sociological explanation (in terms of cultural factors), a medical explanation (in terms of the specific effects of drugs he took). McIntyre writes that the interests in this case are diverse and that each of these explanations are "legitimate". Understood in this way, levels of enquiry are really just different ways of explaining phenomenon. Picking the right level, or explanation, is key to formulating generalisations for McIntyre.

In the 1996 book, he considers the problem that social phenomena are perhaps too complex to find generalisations at the level of inquiry in which we are interested. We may be able to discover stable relationships at another level, but we just aren't interested in these. McIntyre does not give an example, but we might find generalisations about human brains at the level of brain chemistry, but find these unilluminating from a social science perspective. McIntyre's response is that the problem of complexity is due, not to the levels of enquiry in which we are interested, but "is due, perhaps, to the very nature of the natural kinds and descriptive terms used to characterise the phenomena *at* this level of inquiry" (McIntyre, 1996, pg. 27). The problem, as he sees it, is with the way we describe and think about phenomena, not, necessarily, the phenomena itself. Rosenberg (1995) gives a nice example that shows how the use of certain categories can hamper the discovery of laws. He begins with the supposition that we define 'fish' as 'aquatic animals' and try to discover a generalisation about how 'fish' breathe. We might begin by looking at fish, and decide that they breathe through gills. However, then, we might find dolphins and whales, which lead us to specify these animals as exceptions to our generalisation. Unfortunately, we then find some crabs, starfish, lobsters, jellyfish and barnacles, and think there is no longer any point in our generalisation about fish because there are just too many exceptions to it. However, if we define 'fish' as 'scaly aquatic vertebrates' a generalisation about the breathing mechanisms of fish is possible (Rosenberg, 1995, pg. 16).[2] This example illustrates that things may indeed become less complicated if we describe them differently.

McIntyre provides a possible solution to the second type of reflexivity, which is that explanation at different levels, or different types of explanation, may not suffer from this reflexivity. However, as discussed above, it remains somewhat unclear what levels are, and how we are to know which ones are the right ones for formulating generalisations.

[2] However, there is disagreement about whether the concept of a 'fish' is well-formed because some fish, like salmon, are more closely related to mammals than to other fish.

More problematically, McIntyre does not suggest a solution to the first type of reflexivity that Hacking discusses; which is the self-conscious behaviour of human beings responding to how they are categorised, or theorised about. Even if a law were formulated at some other level of description, whatever this may be, McIntyre gives no reason for thinking that such a law would be isolated from the reflection of human beings categorised by this law.

Reflexivity in the social sciences does therefore present a significant problem for social scientists. This is because it means that any regularities or explanations we do find may only be relevant for a short period of time. While McIntyre provides a partial response to this worry, by suggesting that social scientists analyse phenomena at different 'levels' he fails to show that self-reflection will cease to be a problem if we do so.

ii. The necessity of *ceteris paribus* condition in the social sciences

The second reason why the social sciences can be seen as inferior, and which makes regularities difficult to formulate, is the diverse and complicated *ceteris paribus* conditions attached to any purported generalisation, or law. McCauley discusses this in terms of 'invariance principles' (McCauley, 2004, pg. 2). For example, Galileo suggested dropping a ball from the mast of a uniformly moving ship on a smooth sea to show that the ball would fall parallel to the mast. The experiment is intended to illustrate that relative motion does not matter. McCauley says that the experiment "would have made no sense were the earth not a local inertial frame for times on the order of seconds or minutes. Nor would it have made sense if initial conditions like absolute position and absolute time mattered" (McCauley, 2004, pg. 2). Laws of physics are grounded in local invariance principles, with respect to frames moving at constant velocity, local translational invariance, and local time-translational invariance. These invariances, says McCauley, "form the theoretical basis for repeatable identical experiments whose results can be reproduced by different observers independently of where and at what time the observations are made, and independently of the state of relative motion of the observational machinery" (2004, pg. 2). McCauley contrasts this with economics, where there is no invariance. He says, "Because the laws of physics […] are based on local invariance principles, they are independent of initial conditions like time, absolute position in the universe, and absolute orientation. We cannot say the same about markets" (2004, pg. 4) Any economic regularity can be changed by people, at least in principle. Furthermore, economic regularities may vary by country and depend upon initial conditions. If McCauley is right, then it is more difficult to set up experiments in the social sciences because the *ceteris paribus* conditions may be different each time we set up the experiment and therefore difficult, if not impossible, to discover. Unlike with physics there is no 'invariance' that grounds such experiments.

The question for the social sciences is whether McCauley is wrong, and whether it is possible to set up the world in such a way as to see law-like behaviour. One reason that this might be difficult is because often there is little consensus about why certain relationships, such as the infrequency with which democracies go to war one another (the democratic peace thesis), hold in the social sciences. There are a variety of different explanations of democratic peace which focus on different factors, including the decentralised decision-making process of a democracy, and the economic structure of democracies (see Maoz *et al* (1993)). Without consensus about the underlying factors that give rise to the regularity it isn't clear how we would know how to set up an experiment, to say nothing about the difficulties (pragmatic or moral) faced in even setting up an experiment. In Galileo's demonstration of relative motion we can, to a great extent, fill in the *ceteris paribus* conditions by conceptual analysis or by reflecting on limiting factors. With democratic peace, even if the decentralised decision-making process and economic structure were the only limiting factors how could we test the democratic peace thesis by trying to identify the exact properties of a democracy which contribute to peaceful relations with other democracies?

Cartwright (2005) disagrees that there is a fundamental difference between the natural and social sciences in this regard, arguing that we can set up the social world in the same way as the physical world. Her terminology for this is a 'nomological machine'.[3] She says we need this 'nomological machine' to see the relationship described by a law in physics. A nomological machine is a fixed arrangement of factors with stable capacities that, in the right sort of environment, gives rise to the kind of regular behaviour that we can describe as law-like. "Where there is a nomological machine, there is law-like behaviour but we need parts described by special concepts before we can build a

[3] Nomological machines will be discussed here as they relate to the social and natural sciences. However, it is important to note that, for Cartwright, the concept of a nomological machine is quite broad. According to her view, a vending machine would also qualify as a nomological machine.

nomological machine." (Cartwright, 2005, pg. 57). For example, it is easy to set up a simple nomological machine to observe the law-like behaviour suggested by Newton's second law of motion. This law states that when an object is acted on by an outside force, the strength of the force equals the mass of the object multiplied by the resulting acceleration (F=ma). We can observe this relationship by simultaneously dropping various objects of different weights from the same height. Their acceleration is approximately constant due to the force of gravity acting on them, so both hit the ground at the same time. The object with the greater mass will make a bigger impact on the ground, which can be seen if they are dropped into loose material, such as flour. In cases like this we have set up a 'machine' which enables us to observe the relationships described by the law. Cartwight believes that nomological machines are not only found in the natural sciences. She argues that we can have nomological machines in economics, and the following paragraph describes such a nomological machine.

She considers the example of an "economic lever that multiplies money" (Cartwright, 2005, pg. 142). This is analogous to a mechanical lever that multiplies force. The money multiplying force is the following: A central bank has a monopoly on money creation. This money is transferred to the commercial banks, who multiply it based on their reserve ratio. The reserve ratio is the amount of money that banks keep to meet their obligations; the remainder is lent out. The total money in the economy as a whole is therefore larger than the money issued by the Central Bank. The factors relevant to determining the multiplication effect are the banks' reserve ratios, and the amount of money deposited with banks. Her numerical illustration of this is as follows:

Money issued by the Central Bank: £100

Commercial bank receives £100.

Commercial bank retains £20 and lends out £80 (reserve ratio of 1:5)

The borrower deposits this £80 in a retail bank (100% deposit rate).

The retail bank retains £16 and lends out £64 (reserve ratio of 1:5)

The total money in the economic system now is £100 + £80 + £64: £244.

This process continues

This is described by the following equation:

$M= ((1+cu)/(re+cu))H$

Where M is the total money in the economy

H= the money issued by the Central Bank

cu= the rate of currency to deposits

re= the reserve ratio

This equation, according to Cartwright, describes a "socio economic machine that would give rise to a regular association if it were set running repeatedly and nothing else relevant […] happens" (2005, pg. 143). The 'nothing else relevant' corresponds to *ceteris paribus* clauses. Importantly, she notes that this equation rarely describes accurately the amount of money in an economy, because other things interfere. However, she argues that the observation of a regularity is not the only thing that matters. What we want from theories is knowledge that enables us to adjust relevant factors to give us a predictable outcome. In other words, we want to be able to set up a nomological machine (Cartwright, 2005, pg. 144). Cartwright suggests that there is less separating the social and natural sciences in this regard. Is it more difficult to set up such nomological machines in the social sciences though? I will argue that here are two reasons for thinking so.

Firstly, although experiments are difficult to set up in the natural sciences, in the social sciences we are not often in the happy situation of even knowing that our nomological machine has been set up successfully. For example, Cartwright gives the example of a model which, if run repeatedly gives rise to the following regularity:

> "All profitable projects which have no initial sunk costs and whose initial returns are at least as large as the returns from the alternative use of the assets will be undertaken" (Cartwright, 2005, pg. 146).

This model tells us the parts that make up the machine, how the parts are assembled and the rules for calculating the effects. The components are two game players, who have the same discount rates, who are 'motivated only by greed', 'operate under perfect certainty', and 'are perfect and costless calculators' (Cartwright, 2005, pg 146). Again, skirting what it means to have 'perfect certainty', the worry is that even if we know what perfect certainty is, we do not know when we have it. Even if a player were to tell us that they are operating under perfect certainty or are motivated just by greed, we have no way of knowing whether this is true. This is because it relates to the mental states of people, which are essentially inaccessible. For example, they may say that they are motivated by greed, and behave as if they are, but in reality they may also be applying some other rule—such as 'when the other player has lost everything you can stop being greedy'. The other player may not explicitly be aware of this additional rule, and it will only become apparent through their behaviour under a particular set of circumstances. Furthermore, behaviour is rarely as transparent as this—observers often disagree about someone's motivations. We can agree that if conditions of perfect certainty are met then behaviour will conform to a particular model, but how do we know when the conditions of the model are met?

Now, clearly this model is theoretical and needs no commitment to the realism of its assumptions, such as perfect certainty. My purpose is not to criticise this example on the basis that it is unrealistic, nor on the basis that that the assumptions need to be checked empirically. The criticism is more general than this. The worry is that when *ceteris paribus* conditions include references to motivations and mental states it is difficult to know whether they hold. If our aim is purely theoretical this need not worry us; but if we are seeking to use theories practically to generate predictable outcomes then it is a concern.

Davidson is helpful for expanding on this point. He says that when we attribute a belief or desire to a person, we assume that a diverse range of background attitudes are also true. These are, most notably, rationality and consistency. This is why we are hard pressed to say what would convince us that a person preferred *a* to *b*, *b* to *c* and *c* to *a*, because, says Davidson, we only make sense of such preferences against background assumptions of rationality and consistency (Davidson, 1974/2002, pg. 237). Furthermore, Davidson argues, we cannot understand what a person means when they say something, without knowing a lot about their beliefs. Davidson's argument for this is that to interpret verbal behaviour we have to be able to tell when a speaker believes that one of his sentences is true. However, we can't judge the truth of sentences in isolation; sentences are taken to be true on the basis of both what a person believes, and what they meant by their words. Therefore, "just as we cannot infer beliefs from choices without also inferring desires, so we cannot decide what a man means by what he says without at the same time constructing a theory about what he believes." (Davidson, 1974/2002, pg. 238). The significance of this for the present discussion is that firstly, in order to make sense of the behaviour of others we must assume a substantive set of background beliefs, conforming to certain criteria such as rationality or consistency, which may not be accurate. Secondly, judging whether players have certain beliefs or motivations is not only difficult, but not discernible in isolation, without consideration of their wider beliefs. Both of these make it extraordinarily difficult to know whether *ceteris paribus* conditions referring to motivations and mental states hold.

The second reason for thinking that nomological machines are more difficult to set up in the social sciences is that we often don't of know what the *ceteris paribus* conditions are. I noted above that there is no consensus about the reasons for the purported relationship between democracy and peace. When we do not know why there is a relationship between two variables we are unable to specify the *ceteris paribus* conditions that must hold in order for us to build a 'nomological machine'. We may say that democracies do not go to war with each other *ceteris paribus*, but without specifying the content of our *ceteris paribus* clause this is an empty statement. In effect, we are saying that as long as nothing occurs to make this relationship not hold, then it holds.

Roberts (2004) suggests a more sophisticated understanding of what this might mean. Roberts calls laws requiring *ceteris paribus* clauses 'hedged laws', the general form of which is "Whenever A happens, B happens, unless there is an interference" (Roberts, 2004, pg. 163). What matters, in judging the status of such laws is how we understand 'interference'. The proposal that he concedes may be viable in the natural sciences is that interference is understood as a class of circumstances and, although we can't precisely identify all the constituents of this class, we understand what 'interference' means in our particular context. Roberts argues that social scientists may not know what interference means in a particular case because of the enormous physical complexity, and multiple realisability of the underlying physical realisations of social kinds. The general point can be made more simply, while sidestepping a discussion of multiple realisability and social kinds. Roberts does this when criticising Kincaid's (2004) discussion of supply and demand. Roberts says that while the law of supply and demand may be a useful rule of thumb, there are many conditions in which it might be false. These include the imposition of price controls, irrationality on the part of customers or suppliers, humanitarian or other concerns which over-ride economic concerns, and many others. Roberts says that any attempt to specify all the possible conditions is hopeless because these conditions are diverse, and also have diverse causes. He does not mention this, but we could also add that these conditions might interact with each other. In summary, according to Roberts even if we concede that in the natural sciences there may be some ambiguity in what 'interference' is, in a particular context, this ambiguity is vastly more significant in the social sciences.

To summarise, although specifying *ceteris paribus* clauses may be problematic for some cases in the natural sciences it is nevertheless often more difficult in the social sciences. This is because, firstly, it is often very difficult to know what the relevant factors are that need to be controlled for, and secondly, even if we do know what some of them are, where they relate to people's mental states we are unable to know whether they are true.

iii. The social sciences use 'odd' concepts

A number of social scientists have also sought to highlight, if not the inferiority, then at least a peculiarity of the social sciences because of the concepts social scientists use. To cite just a few; Bradburn (2011) describes many concepts that social scientists use as *Ballung* concepts; which he contrasts with "scientific concepts that refer to specific features, such as age, minimum wage etc." (Bradburn, 2011, pg. 53). *Ballung* concepts are not like scientific concepts because these concepts "sort things into categories based on a loose set of criteria in which members of the same category do not share a specific set of features." (Bradburn, 2011, pg. 54). Examples of *Ballung* concepts include 'poverty' and 'disability'. In other words, scientific concepts are associated with precision, whereas there are no necessary and sufficient criteria for membership of a *Ballung* concept, they are associated with less precise associations between members of the same category. This is clearly a worry for social scientists seeking to find regularities because if we are seeking to find a regularity between 'poverty' and something else, and 'poverty' is a *Ballung* concept, then we will be seeking to find a regularity referring a concept for whose members there isn't a feature that they all share.

Little (1993) provides a slightly different characterisation of the concepts social scientists use. He describes them as 'cluster concepts', which he describes as encompassing "a variety of phenomena that share some among a cluster of properties" and that concepts like 'riot', 'revolution', 'class', and 'religion' are a "class of social entities that share a common causal structure", however, this causal structure is not "homogenous". (1993, 190). The example he discusses in detail is 'market economies', which have many things in common, but which also have causally significant differences. The USA and Japan are both market economies and have many properties in common, but the apparent prevalence of longer time-horizons in Japan leads to differences between these market economies. Little believes the value in "making use of cluster concepts […] in the social sciences is that it permits us to group social entities together in ways that emphasise their common features" (1993, 191). There are other ways of defining cluster concepts but this chapter will only discuss Little's because his definition makes for a clear distinction with *Ballung* concepts. For Little, members of a cluster concept have some feature in common. In other words, we can give a necessary criterion for inclusion within a group. This is in contrast with *Ballung* concepts, which, following Bradburn's characterisation, do not have any single feature in common, although a subset of members of a group may have a feature in common. In other words, it is not possible to give necessary or sufficient criteria for inclusion within a Ballung concept. There are different ways of demarcating the difference between *Ballung* and cluster concepts, and it is even possible to argue that there is no difference between the two. This will not be discussed here because it has no bearing on the subsequent discussion.

Despite these different ways of thinking about the concepts social scientists use, these different ways of describing the concepts that social scientists use are not usually seen to be in opposition to one another. This is perhaps because the discussion the literature is not about whether, say, 'poverty' is a *Ballung* concept or a cluster concept, but on how such concepts are to be used in practical research. For example, Crasnow (2015) describes "ordinary use" concepts as "vague and ambiguous" (2015, 7). She notes that "Such concepts have been referred to in a variety of ways in the philosophical literature. Sometimes they are called 'cluster concepts', Nancy Cartwright has called them *Ballungen* concepts […] The idea has similarities as well to Wittgenstein's notion of family resemblance…" (2015, 7). Thereby suggesting that these terms can be used interchangeably. The key thing, for her, is that such concepts are 'vague and ambiguous', and the remainder of her article is devoted to discussion of how to make such concepts more precise. Similarly, Adcock and Collier (2001) begin their process for making social science concepts more precise by beginning with the category "Background Concept", which is "The broad constellation of meanings and understandings associated with a given concept" (2001, pg. 531). Bradburn also follows his discussion of *Ballung* concepts with a discussion of how these concepts are to be defined more precisely, and measured (Bradburn, 2011, pg. 54-57). Although these authors recognise that social science concepts are problematic as they stand, rather than spending time characterising them, they concentrate the ways in which they can be made more useful for research. I believe this is the wrong approach. The following section takes the first step towards demonstrating this by showing that, of the three reasons discussed in this section, the 'oddness' of social science concepts is the most important because reflexivity and *ceteris paribus* conditions are a reason for, and a consequence of, the concepts social scientists use.

The 'oddness' of social science concepts is primary.

The preceding section introduced the observation that the concepts social scientists use are often 'odd', or peculiar. This section demonstrates that the 'oddness' of social science concepts is of primary importance for understanding the alleged inferiority of the social sciences. It demonstrates this by showing that the other two main points of difference between the natural and social sciences, which are reflexivity and extensive *ceteris paribus* conditions, are grounded in this 'oddness'. Reflexivity because it partly explains the 'oddness' of the concepts social scientists often use, and *ceteris paribus* conditions because these are partly a consequence of the 'odd' concepts that social scientists use. Specifically, reflexivity is partly responsible for changes that occur in social science concepts. Furthermore, if social science concepts are often *Ballung* concepts then the existence of extensive *ceteris paribus* conditions should be no surprise.

Reflexivity:

This section demonstrates that reflexivity influences the concepts that social sciences use and that this, in turn, partly explains why they are often described as *Ballung*, or cluster concepts. Hacking's work on autism illustrates how the concepts social scientists use change due to reflexivity. He demonstrates that the concepts social scientists use evolve because of changing behaviour amongst those who are classified as 'autistic'. This, I will argue, is one reason why such concepts may be described as *Ballung* concepts.

Hacking proposes that classifying people as autistic may change their behaviour, and the behaviour of people who interact with them. He argues that people with autism are developing ways of describing their condition, for which there are few generally accepted descriptions of autism recognised as valid by high-functioning people with the condition. This is happening through a combination of diaries written by people with autism, and fictional portrayals of autistic characters. This is in turn is leading to greater understanding of the condition (Hacking, 2009). Tsou (2007) believes that autism is a difficult example for Hacking, because, he says "individuals with such conditions […] are *not fully aware* of their condition or how they are classified" (Tsou, 2007, pg. 332. Italics in original). This makes it problematic to see how people with autism can respond to their classification. Hacking's response to these sorts of worries is that it is not only the reaction of an individual with autism that is important, but the reactions of the people with whom they do interact (Hacking, 1999, pg. 115).

McGreer (2009) explicitly sets out to assess whether ideas about autism are evolving in the ways Hacking suggests, and concludes that, broadly speaking, they are. She argues that the way we think about the condition is changing as people with autism describe their condition, and that this, in turn, affects the behaviour of people with autism. Her methodology is to assess both the extent to which knowledge about autism is being transformed by what people with autism are saying about themselves, and the effects public discussion of autism and the proliferation of fictional characters with autism are having on people diagnosed as autistic.

This example illustrates how the concepts social scientists use change over time. Firstly, the meaning of the concept 'autism' changes as experts and the wider community understand the condition differently, partly as a result of the behaviour of people with autism. The underlying behaviours of people with autism are also changing as they alter their behaviour in response to environmental changes. These environmental changes are, in turn, partly influenced by changes in understanding the concept 'autism'. When we refer to 'autism' at different points in time it is possible both that we mean something slightly different by the concept, and that the phenomenon will have undergone changes too. With the benefit of hindsight some of these changes may be easy to identify, but they are likely to be more difficult to track while they are happening. This is in contrast to the natural sciences where concepts do not change dynamically over time. Furthermore, this liability to change gives a reason why many concepts may be described as *Ballung* concepts.

The concept, such as autism, is tracking a changing use of the concept, and changing behaviour that is described by this concept. These changes are likely to be difficult to track as they are happening, so it is natural that the concept refers to social phenomena for which it is difficult to specify necessary and sufficient criteria for inclusion within the concept.

Ceteris paribus conditions:

Where the concepts that social scientists use are *Ballung* concepts, or even cluster concepts, the existence of complicated and diverse *ceteris paribus* conditions should be no surprise. This is because these concepts, such as poverty, include a diverse range of phenomena. Therefore, when formulating a regularity there are many potential factors and situations that we need to exclude from the regularity. Revisiting the democratic peace thesis from the *ceteris paribus* section above illustrates this. The concept of democracy is, arguably, a cluster concept rather than a *Ballung* concept because democracies do have features that are common to them all (particularly of we restrict ourselves to modern democracies) which are the criteria that the population of a political region have the power to change the government. Political scientists may dispute the extent to which this happens, or whether this is the only criteria for 'democracy', but this is not the central concern here. It is arguable that 'democracy' is a cluster concept rather than a *Ballung* concept. If we accept that there is something all (modern) democracies have in common, this does not mean that there won't also be many differences between them too. Therefore, if we are seeking to formulate a regularity about democracies, then we are likely to need a variety of *ceteris paribus* conditions to exclude other features of democracies, or other variables relating to democracies that complicate, interfere, or otherwise render obsolete, the generalisation we have highlighted. Extensive *ceteris paribus* conditions should therefore be expected when we are dealing with 'odd' social science concepts.

Summary

The social sciences have often been thought to be inferior to the natural sciences, in part because they have failed to find many satisfactory examples of laws. A review of the philosophy of science literature suggests that we should discuss 'regularities' rather than 'laws'. Nevertheless, regularities are no easier to find in the social sciences than laws. I discussed three reasons that have been given for this inferiority of the social sciences. Reflexivity, the preponderance of *ceteris paribus* clauses, and the use of 'odd' concepts, and concluded that these do point to significant differences between the natural and social sciences. I suggested that the use of 'odd' concepts is a particularly important difference because reflexivity, as well as illustrating how the social world changes, also shows how concepts used in the social sciences can refer to changing aspects of the social world over time. This, in turn, leads to different things being discussed when we use these concepts over time and makes them more easily describable as *Ballung*, or cluster concepts. Furthermore, when we are dealing with *Ballung*, or cluster, concepts then the extensive *ceteris paribus* clauses should be no surprise.

So far, this discussion has not explored the precise sense in which social science concepts are 'odd'. It seems clear that reflexivity is partly responsible, but other explanations are discussed in the literature. The following section discusses three possible explanations why the concepts social scientists use are 'odd'. These are:

i. The social sciences deal with essentially contested concepts

ii. The concepts social scientists use are vague

iii. The concepts social scientists use are multiply realisable

I conclude that these explanations are unsatisfactory, before suggesting a new approach.

i. Do the social sciences deal with essentially contested concepts?

Gallie proposed that some concepts are 'essentially contested'. The examples he gives include 'work of art', 'democracy' and 'social justice'. These concepts are essentially contested because they are characterised by disputes about what, really, counts as a 'work of art', or 'democracy' (Gallie, 1957, 169). He writes that "there are disputes […]which are perfectly genuine: which, although not resolvable by argument of any kind, are nevertheless sustained by perfectly respectable arguments and evidence" (Gallie, 1957, pg. 169). In other words, despite dispute, argument, and the presentation of evidence, no agreement is reached. Gallie gives seven criteria for a concept being judged 'essentially contested'. These are;

1. The concept must represent some kind of valued achievement

2. The achievement must be internally complex

3. The achievement is describable in a variety of different ways

4. The achievement is liable to change as circumstances change

5. Scholars recognise that their use of the concept is likely to be contested by other scholars

6. The concept is derived from an original example, which is recognised as such by all the scholars using the concept.

7. Continuous competition between scholars "enables the original example's achievement to be sustained and/or developed in optimum fashion"

   (Gallie, 1957, pg. 171-172, 180).

The difference between essentially contested concepts and more scientific concepts is that "Competition between scientific hypotheses works successfully largely because there are acknowledged general methods or principles for deciding between rival hypotheses." (Gallie, 1957, pg. 179). In other words, Gallie believes that, when using essentially contested concepts, debate continues indefinitely without resolution. If social science concepts are 'essentially contested' this might explain some of the 'oddness' of these concepts. Continuing contestation might explain why some concepts are '*Ballung* concepts' because continuing debate might lead to changes in what we consider 'poverty' to be. McKnight (2003) points out that it is impossible to define essentially contested concepts in terms of necessary and sufficient conditions, which is also the defining characteristic of *Ballung* concepts. Concepts not defined by necessary and sufficient conditions will be disputed by scholars who dispute either the particular use of the concept, or the particular descriptions of the concept. McKnight says, "The open-endedness of the concepts concerned means that we cannot lay down in advance laws for their future application and any attempted statement of such conditions will itself be disputable." (McKnight, 2003, pg. 262). It seems, therefore, that *Ballung* concepts and essentially contested concepts have much in common.

Essential contestation and reflexivity may also be related because one of the criteria that essentially contested concepts must meet is that they represent a valued achievement, which is describable in a variety of ways. Collier *et al* in their discussion of Gallie suggest that "it is highly plausible that the positive normative valence attached to these concepts is important in spurring debates over their meaning." (Collier, 2006, pg. 216). Their view is that 'normative valence' is, at least in some cases, responsible for debates about the meaning of concepts. Freeden argues that essentially contested concepts may also be "disapproved and denigrated phenomena" (1996, pg. 56). This could lead to changes of the type Hacking describes when people respond to how they are described, which in turn affects how they are analysed by social scientists. Similarly, the debate over the meaning of 'poverty' may well influence the behaviour of people judged to be 'in poverty' in the ways that Hacking suggests.

Despite the immediate appeal of the notion of 'essential contestation', this account is unsatisfactory. Firstly, because not all concepts that social scientists use have obvious normative valence. Ehrenberg (2011) says that despite the importance for social scientists of the concept of 'government', "it is not illuminating of the concept of government to call it essentially contested. We do not usually see the important debate over the concept of government as consisting in what a government is, or what counts as a government. We do see it with respect to what counts as a just government, that is clearly a debate about the meaning of justice (a good candidate for essential contestation if ever there was one)." (Ehrenberg, 2011, pg. 215). Secondly, if Ehrenberg is correct, and I think that it is plausible that he is, then the critical question is whether the concept of 'government' is like a 'scientific concept', characterised by precision. This is because, if essential contestation explains why many social science concepts are *Ballung* concepts, or subject to reflexive change, then concepts that are not essentially contested should not have these characteristics. It seems clear that this isn't the case. Even though 'government' isn't essentially contested, it may still be a *Ballung* concept or a cluster concept. The Oxford English Dictionary defines government as 'the system by which a state or community is governed'. Given the almost infinite variety of ways in which this is done, it is reasonable to think that 'government' is a *Ballung* concept, or if we do manage to specify something that all governments share, a cluster concept. Either way, citing Bradburn again, the concept 'democracy' does not "refer to specific features, such as age, minimum wage etc" (Bradburn, 2011, pg.53). Gallie's essential contestation is therefore an interesting proposal for what is going on with social science concepts, but is incomplete because a concept with an absence of essential contestation may still be 'odd'.

ii. **Are Social Science concepts vague?**

The second proposal for why the concepts social scientists use are 'odd' is that they are vague. We saw above that Crasnow describes many concepts that social scientsts use as "vague and ambiguous" ((2015, 7). Vagueness seems to work as a general, catch-all description of what is going on. However, it is unfortunately the wrong word to use because this use of the term 'vague' is at odds with the way it is used in the philosophy of language.

The usual understanding of vagueness is that it is primarily due to the presence of boundary cases. To cite one recent example "Many if not most predicates in natural language are vague, meaning that they admit of borderline cases, that is, cases to which they neither fully apply nor fully fail to apply" (Douven, 2016, pg. 1). Similarly, the Cambridge dictionary of Philosophy begins the entry on 'vagueness' with the following "a property of an expression in virtue of which it can give rise to a 'borderline case"' (Audi *et al*, 1999, pg. 945). The standard example of a borderline case is a 'heap'. One grain of sand is not a heap. Two grains of sand are not a heap either. So, if we say that *n* grain of sand is not a heap, then *n*+1 grains of sand is not a heap. But, this appears to allow us to conclude that 50,000,000,000 grains of sand is not a heap (1 billion grains of sand takes up approximately one 30cm x 30cm cube). This is called a Sorites paradox.

The presence of borderline cases is sometimes a problem in the social sciences, as the following examples illustrate. The concept of poverty lacks precise boundaries. It is possible, and highly likely, that people will disagree about whether a group of people is 'in poverty'. This is regardless of whether just one meaning of 'poverty' has been specified. For example, we may specify that 'poverty' means 'lacking adequate provisioning of the necessities of life', and specify these even more precisely, such as '2000 calories per day', 'accommodation', and 'clothes'. A particular person may get 2000 calories most days, but not every day. Do we count them as satisfying this criterion for being in poverty, or not? Determining where to draw the boundary between poverty, and non-poverty based on the amount of calories consumed per day is similar to the 'heap' and 'non-heap' example in the sense that small, relatively homogenous units (grains of sand or single calories) can be added together. Judging where to draw the boundary is a question of how many of such units we think are necessary for being a heap, or not in poverty. However, not all examples in the social sciences are like this. Sometimes, deciding where a boundary lies is not a matter of judging how many, relatively homogenous, units are necessary, but judging how many separate factors are present. Not all cases of 'oddness' in the social science concepts are instances of vagueness like the Sorites paradox. The following example illustrates this point.

Cartwright and Runhardt (2014) write that one of the indicators used to decide whether a country is, or is not, experiencing a civil war is the 'military death count'. The 'military death count' sounds like a precise measure, but in many cases may not be so. This is firstly because both sides in such a conflict have incentives to overstate the number of deaths on the other side, while understating the level of deaths on their own side. Secondly, even if a dead person is found by an observer it is often unclear how to classify this. If we find a man fatally shot, wearing civilian clothes, carrying a gun, near a combat zone, how do we decide whether he is a military or civilian death? This is not a question about what a 'military death' means; we are clear that is means the number of military personnel who have died. Our problem is determining whether the person is, or is not, a member of the military.

There are a number of factors that determine whether a death counts as a military death, which might include carrying a gun, being dressed in combat fatigues, being found close to the location of a recent conflict and any number of others. Having none of these features clearly means that it is likely not a military death, while having all of them means that it is a military death. Even though we can tell these cases apart, this still does not make it a simple matter to say whether having, say, four out of five characteristics counts as a military death because some characteristics may be more salient than others, and salience may vary by context. For example, the importance of gun carrying might vary depending on the prevalence of gun carrying among the local population. The military death example therefore differs from the sand example because there are a number of features that are relevant for judging the boundary between military and non-military deaths, while in the sand example it is just a matter of how many grains of sand count as a heap. In principle, it might be possible to weight these features to indicate the extent to which a death is military, however, this is problematic because these weightings may vary by context. This is not merely a case of not knowing whether someone is a member of the military because we lack sufficient information to decide. Even if we have all the relevant information, it is a case where it may be indeterminate whether someone is part of the 'civilians' or the 'military'. In a civil war, much of the population maybe neither one nor the other, so deciding which features are relevant, and which should be judged most important in a particular case is difficult to do. It is plausible that that there is "no definite answer to the question, 'How many of which conditions are necessary for the term to apply?'" (Audi *et al*, 1999, pg. 947).

The idea of 'vagueness' as it is discussed in the philosophy of language literature does not match what is going on the social sciences very closely. This is because, although there are often boundary issues with concepts that social scientists use, these are also often more complicated than the examples discussed in the philosophy of language. However, the etymology of the word 'vague' is interesting. Carney and Miller (2016) write that it derives from the Latin verb *vagari*, which means to wander or roam. The noun vague in English means "an indefinite expanse so that a vague, perhaps close to its 'wandering' origin, relates to terrain and is thus immediately connected to considerations of space and geography […] Thus a vague, as an indefinite terrain, is the 'place' where wandering or roving occurs, a place of the nomad" (Carney & Miller, 2016, Ch. 2). This idea of undefined space will be discussed in greater detail in Chapter 2. Nevertheless, it is unhelpful to use a word that has been given such a clear meaning in the philosophy of language.

The quote from Crasnow above also mentions ambiguity. Is it ambiguity that is the source of the oddness of many concepts used in the social sciences? I do not believe that it is. Ambiguity is when people are unclear about which, of a number of meanings, they have in mind when they use a concept. The oft used example is 'bank'. If I just say 'bank' I may have either a river bank or a financial institution in mind. The discussion of civilian deaths illustrates that, however many times social scientists might be unclear about whether they should run towards the nearest river or their local high street when told to 'head for the bank', there is more plaguing the social sciences than ambiguity. Social scientists are not unclear about what they mean by a military death; even though this concept is unambiguous, they are still unsure whether what they are looking at counts as a member of this group.

iii. Are social science concepts multiply realisable?

The final proposal for understanding why social science concepts are odd is that they are multiply realisable. Roberts, in a debate with Kincaid, argues that the reason that the social sciences do not have laws is because social phenomena are multiply realised (2004). He takes issue with Kincaid's assertion that the law of supply and demand is a *bona fide* law. Kincaid formulates this law as follows:

> "Changes in price cause changes in the quantity demanded and quantity supplied.
>
> Changes in supply and demand curves cause changes in price" (Kincaid, 2004, pg. 177)

Roberts writes that all kinds of things could count as a 'market' for the purposes of assessing the law of supply and demand. When this is the case, it means if we have a law stating that "all F-systems in circumstances C exhibit behaviour G" (Roberts, 2004, pg. 162) there will be many physical instantiations of F and G and C, some of which will not be covered by the law. He argues further that even if we do not know of any such instantiations at the moment, the 'law' will be fragile because we have no reason to believe it is impossible that such instantiations could appear.

He uses a slightly far-fetched example to demonstrate this point. Suppose that the 'F-system' is currency markets and C is the level of inflation. He considers the case where humans undergo psychological changes, such that they become humanitarians and reject their previous acquisitive ways ("the sudden widespread imitation of Mahatma Ghandi" (Roberts, 2004, pg. 162). These humans, argues Roberts, are a physical instantiation of an F system, just as much as more acquisitive humans are. However, the behaviour (G) will differ greatly in this case. He summarises:

> "The basic point here is that no matter which social-scientific kinds F, C and G are, there are likely to be kinds of physical system that constitute social systems that instantiate F and C but, under normal physical

evolution, lead to bizarre outcomes, which will not be covered by G" (Roberts, 2004, pg. 162)

He suggests that the sorts of things social scientists often talk about can be multiply realised, and the well-known example of 'money' demonstrates this. Roberts notes that 'money' is realizable in very many ways, including letters of credit, precious metals, shells, and promissory notes. The multiple realisability of money appears uncontroversial and, consequently, his argument works better in abstract form, when it refers to F systems in circumstances C exhibiting behaviour G. If 'F', 'C', and 'G' refer to things like money, then it is clear that there will indeed be many ways in which each of these variables can be realised. When laws include terms that are multiply realisable then some of these realisations may not be covered by the law.

If social phenomena are multiply realisable, this may partly explain why the concepts that social scientists use are 'odd'. When discussing a 'market' or 'money' there are a wide variety of things we can include within this group. This suggests one reason why social science concepts can be described as *Ballung* or cluster concepts (*Ballung* concepts if there is nothing that all members of a group share, or cluster concepts if there are some common features).

However, this is an unsatisfactory explanation for the simple reason that not all social phenomena are multiply realizable. Tellingly, Roberts does not show that 'supply' and 'demand' are multiply realisable. Kincaid gives a clear specification of the laws of supply and demand, and the key concepts he uses are 'supply', 'demand', and 'price'. If Roberts seeks to show throw doubt on the nomological status of these statements, why does he not show that 'supply', 'demand', or 'price' are multiply realised? Indeed, in his more abstract formulation of his worry about multiple realisability he says that "all F-systems in circumstances C exhibit behaviour G" (Roberts, 2004, pg. 162) there will be many instantiations of F and C and G. This does not exactly match Kincaid's formulation of the law of supply and demand. Kincaid makes no reference to a 'system'. Roberts takes the relevant system to be a 'market' of some sort. Let us grant that this is reasonable. This would change the first half of Kincaid's formulation of supply and demand to the following:

All markets in circumstances of changes in price exhibit changes in supply and demand.

Let us grant that this is reasonable. Let us also accept that a 'market' is multiply realisable. Even if we do so, multiple realisability does not pose a significant threat to Kincaid's law of supply and demand. This is because we can accept disagreement over what counts as a market, but nevertheless find some examples of markets where we are happy to assess the law of supply and demand. In doing so, the other concepts in this formulation 'supply', 'demand', and 'price' are not, obviously, multiply realisable. In other words, Roberts only shows that multiple realisability is a problem for deciding where to *assess* the law of supply and demand; he does not show that it is a problem for the concepts used by Kincaid in his formulation of the law itself. Rather than targeting the terms Kincaid uses in his formulation of supply and demand Roberts imports the idea of a market to demonstrate multiple realisability. As such, his argument is unconvincing.

The law of supply and demand is one of the better confirmed examples of a regularity from the social sciences. However, as Roberts discusses above, one of the main criticisms of this law is that it requires wide ranging *ceteris paribus* conditions. Roberts notes that some concepts that social scientists use are multiply realisable, but not that 'supply' or 'demand' are multiply realisable. Multiple realisability is therefore a poor explanation of the why this law requires so many *ceteris paribus* conditions.

An alternative proposal

The discussion above illustrates that understanding why the concepts that social scientists use are 'odd' is critical to understanding the reasons for the purported inferiority of the social sciences. It also showed that the currently available explanations of this 'oddness' are unsatisfactory. However, much of the contemporary discussion of concepts in the social sciences sidesteps explanation of this 'oddness' in favour of proposing how concepts are to be used in research. This section demonstrates why this is the wrong strategy.

As discussed above, philosophers characterise the concepts that social scientists use in a variety of ways, and, importantly, they do so using different examples. The examples of *Ballung* concepts include 'poverty' and 'disability'. Hacking's examples include 'teen aged pregnancies' and 'autism'. Little's examples include 'market economies'. Roberts' uses the examples of 'a market' (by which he means a financial market), and ‘money’. The use of such varied examples hides one important question, and a number of slightly less important questions. The main question is what is it about a concept such as ‘poverty’ that means that it can be described as a *Ballung* concept, and what is it about the concept like ‘market economy’ that makes its characterisation as a cluster concept appealing, and what is it about the concept ‘money’ that means that it is multiply realisable? The related questions are: Is ‘money’ multiply realisable, *and* a cluster concept? Do *‘Ballungness’,* or clustering, or reflexivity, come in degrees? Do concepts that are multiply realisable all have the same number of realisations? If not, why not?

Answering these questions is important because currently, philosophers of social science write that some concepts are ‘not the right ones’ for analysis, while some are, without providing any principled way of distinguishing between them. In other words, they proceed in a piecemeal fashion, concept by concept, with no structured way of thinking about the work a particular concept can do for us in a particular case. For example, Cartwright says that nomological machines need “parts described by special concepts”. She says that everyday concepts like “irritability and inaccuracy will not do” (2005, pg 57). It is intuitively clear that ‘irritability’ won’t feature in most nomological machines. But Cartwright provides nothing over and above the immediate and obvious appeal of this statement in its support. What, exactly, is it about ‘irritability’ that means it won’t do, while ‘perfect information’ will do?

Woodward (2016) largely agrees with the need to take a structured approach to choosing which concepts to use. He begins by acknowledging that philosophers are familiar with the idea that some variables are bad, or unhelpful, for philosophical analysis. He illustrates this with the obvious example of ‘grue’. Woodward rejects the idea that ‘bad’ variables need no further analysis, that it is enough that we can recognise that they are bad, because there is no obvious way of telling ‘good’ and ‘bad’ variables apart. This is consistent with the argument above. Despite this, I will show that Woodward’s approach is not the correct one.

Woodward’s analysis in this 2016 paper relies on his manipulationsist account of causation (outlined in detail in Woodward, 2003). This will be discussed in detail in Chapter 3, so will be taken for granted here, particularly as it has little bearing on his comments about variable choice, except in his description of what scientists seek to achieve with variables. Woodward’s criteria for good variable choice are:

1- "Choose variables that are well-defined targets for (single) interventions in the sense that they describe quantities or properties for which there is a clear answer to the question of what would happen if they were to be manipulated or intervened on. […] Possible candidates for variables failing to meet this criteria […] include 'age', 'gender' and 'obesity'."

2- "Choose variables that have unambiguous effects on other variables of interest under manipulation, rather than variables that have ambiguous or heterogeneous effects"

3- "Choose variables that are defined in such a way that we can in principle manipulate them to any of their possible values independently of the values taken by other variables. […] This excludes, for example, variables that are logically or conceptually related."

4- Choose variables that "are relatively *sparse* in the sense that they postulate relatively few causal relationships among variables, rather than many."

5- "Choose variables that allow for the formulation of cause-effect relations that are as close to deterministic as possible or at least relations that exhibit strong correlations between cause and effect."

6- "Look for variables that allow for the formulation of causal relationships that are *stable* in the sense that they continue to hold under changes in background conditions".

7- "In general, look for variables such that the resulting graph accurately represents dependency relations, avoids unexplained correlations in exogenous variables, structure in residuals, and causal cycles with no obvious empirical rationale or interventionist interpretation."

(Woodward, 2016, pg. 1054-1055, italics in original)

These criteria are practical in the sense that they highlight the role that variables should perform in scientific analysis. They should be good ones for manipulation, for example. Woodward is explicit about this, saying that what we consider to be a 'good' variable depends on the work we want that variable to do (2016, pg. 1057). Woodward rejects an *a priori* approach because, he says, such an approach commits us to using variables that are either the fundamental categories of physics (or an as yet specified perfect physics), or variables that are reducible to these fundamental categories. I will firstly provide a reason for thinking that his criteria are unsatisfactory before showing that these worries suggest that we do, in fact, need an *a priori* analysis of variables.

The problem with these criteria, at least as far as they relate to the social sciences, is that we need to know the work that they do, before knowing that they do in fact do this work. For example, we need to know that 'obesity' is not a well-defined target for intervention, and that it doesn't have unambiguous effects on other variables. Woodward does not say how we know this. In advance of attempted, and failed, attempts to use 'obesity' in analysis it is difficult to see how we are to know whether 'obesity' has these characteristics. Given this, the criteria seem redundant in cases like this; we will simply settle on the variables, or concepts, that have yielded successful analysis. My central concern is that Woodward does not explain why 'obesity' is a bad variable, and specifically, what it is about 'obesity' that makes it a bad variable. If we know this, then we are much more likely to be able to avoid other 'bad' variables without first trying to use them and finding them wanting.

Woodward does comment on the social sciences specifically, he says that causal analysis in the social sciences of often plagued by heterogeneous effects of causal variables. He also writes that this is exacerbated by the fact that, often, the variables that social scientists use are proxies for the actual variables they want to analyse. The example he discusses is 'education', for which we might use years spent in school as a proxy. However, as Woodward correctly notes, the quality of schools varies widely, which results in unstable effects of 'education' as it is measured. The obvious suggestion is that we should work on measuring 'education' better, but, Woodward writes that "there are obvious practical limits on our ability to do this and so we may be stuck with heterogeneous or 'ambiguous' variables" (2016, pg. 1070). This, I believe, is precisely the question that needs answering. What is it about 'education' that makes it heterogeneous?

While Woodward is correct that we need a way to differentiate bad from good variables, or concepts, his focus on how these variables perform in a particular analysis is unenlightening because it relies on knowing which variables are, in fact, successful for explanation. As he correctly notes, the heterogeneity of many social science concepts, such as education, makes it difficult to judge whether this is a good variable in advance of attempts to use it. Furthermore, Woodward takes it for granted that it is clear whether a variable is a 'well-defined target', or has 'unambiguous effects'. The heterogeneity he points to may make it a matter of heated debate whether smaller class sizes improve education, for example- some social scientists may fervently believe that this is true, while others may equally fervently believe the opposite.

In summary, Woodward is entirely correct to suggest that social scientists need a way to distinguish between good and bad variables. His emphasises how well those variables perform in causal analysis. As he realises, heterogeneity is a particular problem in the social sciences, and this is a problem that needs explanation. What is it about 'education' that means that it is heterogenous? Once we have answered this question, social scientists will be in a better position to assess whether 'education' is a 'good' variable, and how best to work with it.

Chapter 2 sets out a framework that attempts to unify the various strands of discussion in this chapter, with the aim of providing an *a priori* way of thinking about whether concepts are more, or less, useful for social science research. Drawing on the historic meaning of the word 'vague' as a place where wandering or roaming occurs, social science concepts will be described as Nomadic. Chapter 2 therefore sets out the different ways in which concepts social scientists use roam, or wander, through the social landscape. One way in which they do this is through reflexivity, which leads them to occupy shifting areas over time. This nomadic behaviour will be shown to explain why the social sciences are often plagued by extensive *ceteris paribus* clauses, and why they have historically failed to find many convincing examples of regularities.

Chapter 2: Nomadic Concepts

Chapter 1 argued that understanding why many concepts used in the social sciences are 'odd' is of primary importance for understanding why the social sciences have struggled to find regularities. This chapter sets out a framework for thinking about this 'oddness' in a systematic way. Concepts that are often used in the social sciences will be described as Nomadic. The word 'Nomadic' is intended to convey the shifting and changing path of a concept across the social landscape. As with a nomadic tribe, or group, the boundaries of this tribe change over time. The landscape thorough which these nomads roam should be thought of as representing the social world, in other words, all phenomena that might be of interest to a social scientist. This chapter discusses the ways in which concepts used in the social sciences are Nomadic, before outlining two implications of adopting this approach. Firstly, it explains why the social sciences have found so few examples of regularities and secondly, it suggests that proposals made by Crasnow (2015), Cartwright & Bradburn (2011) and Goertz (2006, 2008) and Adcock & Collier (2001) for making concepts more precise are unsatisfactory.

The role of the Nomadic Concepts framework

Describing the concepts used in the social sciences as Nomadic allows social scientists to bring together the discussion in Chapter 1 into a coherent whole. Chapter 1 argued that the concepts used in the social sciences have been described in a variety of ways, but that there is no way of deciding, for example, whether a concept is a 'good' one for analysis, or not. The following framework allows this to be done.

The central aim of this approach is to demonstrate why social scientists can disagree about whether some specific aspect of the social world falls under the scope of a particular concept (for example, whether a particular political system is 'democratic'), and also why they can disagree about which parts of the social world fall under the scope of a particular concept (for example, which political systems fall under the scope of the concept 'democracy'). The first is a bottom up approach because it begins with a particular political system and asks whether it is 'democratic', while the second is a top down approach which begins with the concept of 'democracy' and searches for political systems to which this concept applies. Where concepts are Nomadic, both of these tasks are contentious because social scientists can potentially include a lot of the social world within the scope of such concepts. As such, these disagreements are not easily resolved. This contrasts with concepts that are more precise; when a concept is precise social scientists focus on the same aspects of the social world when using this concept. In other words, social scientists agree about which aspects of the social world fall under the scope of a particular concept. The following section describes the characteristics of Nomadic concepts, which are unclarity of boundaries, change over time, and many possible meanings. These characteristics demonstrate that there are a number of ways in which a concept can be Nomadic, and that concepts may be more or less Nomadic.

What is a Nomadic concept?

Nomadic concepts are defined by the following criteria:

1- A wide variety of social phenomena can be included within the scope of the concept. This results from these concepts having many possible meanings, unclear boundaries, and changing over time. These characteristics are not an all or nothing matter because these concepts can vary in the number of meanings they have, how unclear their boundaries are, and the extent to which they change over time.

2- The characteristics outlined in criteria 1 mean that disagreements about Nomadic concepts, and arguments making use of them, are difficult to resolve with academic analysis. Over time, analysis of a Nomadic concept leads to the incorporation of different social phenomena.

'Social exclusion' is an example of a Nomadic concept. Criteria 1 is satisfied because a wide variety of social phenomena can be included within the scope of the concept 'social exclusion'. There are many things 'social exclusion' can mean. Amongst other things, it can mean an inability to participate in economic life, a lack of interaction with the community (as in the case of the house bound elderly), or a refusal by other people to interact with the people in question in a constructive way (as in the case of children bullied at school), and a lack of access to normal means of communication. Following Sen (2004), we might also argue that people are excluded from society if they do not have access to certain capabilities. In this case technological capabilities are important. Each of these meanings does not have precise boundaries; for example, how little interaction do we need with the community in order for it to count as a 'lack of interaction'? Or, how little access to technology counts? We might agree that someone is socially excluded if they have no internet access and no phone. But what if they have a shared mobile and access to the internet at a local library, or their neighbour's house? We could argue this either way. Furthermore, these meanings can be expected to change during the time period of our analysis. We would not have considered children socially excluded if they lacked access to the internet in 1995, but we would arguably consider them so today. Some aspects of the concept of social exclusion have changed significantly over a relatively short period of time.

'Social exclusion' meets condition 2 because academic analysis has not resulted in agreement about what social exclusion is. For example, Axford writes that "the concept is used indiscriminately to describe myriad phenomena, from unemployment to being sexually abused, with some commentators even arguing that children as a class are excluded" (2010, 738). Over time, new ideas have been brought to bear on the concept, as illustrated by Sen's capabilities approach. It is not unreasonable to suppose that, in the future, social scientists will find new ways in which people are 'socially excluded'. Indeed, see Richardson & Le Grand (2002) for a paper outlining modifications to academic definitions of social exclusion based on discussions with people who are socially excluded. This is not to say that there is anything about 'social exclusion' that means that, in principle, it is impossible to reach agreement about the phenomena relevant to understanding this concept. In the future there might be such agreement. However, this would require agreement that 'social exclusion' has only one meaning, which has clear boundaries, and does not change over time. Were this to happen, then the concept would no longer be Nomadic. The following section discusses the characteristics of Nomadic concepts in more detail, with the aim of clarifying this definition.

Nomadic concepts have unclear boundaries

The meanings of 'social exclusion' have unclear boundaries because it is difficult to decide, for example, how much social interaction is needed for children to feel socially integrated. Chapter 1 discussed another way in which boundaries can be unclear, which was because a Sorities Paradox can be generated. The Nomadic concept framework is not being proposed just to add another descriptive term to the philosophical literature, and to be successful it must be able to incorporate the sorts of boundary issues discussed in Chapter 1. The following section illustrates how to do this. The advantage of this is that it allows for discussion of concepts that appear very different to one another within a single framework.

Characterising concepts as Nomadic is a different way of describing boundary problems. To see this, we need to think about the analogy with a Nomadic tribe. Firstly, consider a tightly bunched tribe that has put up a fence around its camp. In this case, the boundary issues will be almost non-existent. People living within the boundary are part of the tribe, and people living outside, are not. Once we remove the fence, things become more complicated. Members of the tribe may leave for periods of time, sometimes living within the tribe, and sometimes living elsewhere. In such a case, we might specify a number of days spent with the rest of the tribe as necessary for being a member of this group. In this case, we can describe the boundary as vague, in the philosophy of language sense. We are trying to draw a boundary around the tribe and adding, or subtracting, units that are identical for the purposes of this analysis. To make this precise a Sorites Paradox can be generated.

A person who spends 2 nights per year with the tribe is not a member of the tribe. A person who spends 3 nights per year with the tribe is also not a member of the tribe. So, if we say that *n* number of nights per year spent with the tribe does not denote membership of the tribe, then *n*+1 nights per year spent with the tribe does not denote membership of the tribe. But, this appears to allow us to conclude that 365 nights per year spent with the tribe does not denote membership of the tribe either.

Continuing with the current metaphor, there is clearly more to membership of a tribe than the number of nights per year spent with other members. Other criteria for membership may include biological descent, adoption of social practices, and proximity to the main tribe. For example, we may find people living at some distance from the main tribe, judging whether they are members of the same tribe will involve considering a variety of factors such as, for example, what both groups of people tell us, whether they appear to act in similar ways, whether they have similar beliefs, and the degree to which they interact. This is like the difficulty with judging whether a child is socially excluded, where we are no longer simply drawing a boundary between identical units, but attempting to come up with a list of criteria that help to decide whether or not someone is a member of a group. Where concepts that social scientists use have unclear boundaries in either of the senses outlined above, this contributes to them being Nomadic and increases the potential for disagreement about which aspects of the social world fall under the scope of the concept.

Nomadic concepts change over time

The concept of 'social exclusion' has changed over time; for example, through the inclusion of access to technological capabilities. This change is due to social scientists reflecting on the concept and finding new ways in which people might be 'socially excluded'. This is not the only way in which concepts can change. Chapter 1 discussed reflexivity, which is a different way in which concepts can change. As a brief reminder, the example of 'autism' illustrated that the meaning of concepts changes as people reflect on the concepts used to describe them, which leads them to change their behaviour, and this, in turn, leads to changes in the concepts used by social scientists to understand this behaviour. This can be illustrated by the description of Nomadic concepts. In the section on boundaries above, the landscape on which the tribe is placed was unimportant. This section illustrated how the landscape is critical to understanding how Nomadic concepts work.

The concept 'autism' tracks changing behaviour over time, which can be represented as the concept (or tribe) moving across the landscape, where the landscape represents the social world. This movement may be fast, or slow. Over time, the aspects of the social world to which the concept 'autism' applies change, and this should be understood as the concept moving across the landscape. Furthermore, the movement of the concept also changes the landscape itself as people with autism change their behaviour, and, through interaction with medical professionals and wider society, alter the social world. In other words, the concept moves over the landscape, but in doing so changes the landscape which, in turn, changes the concept.

There is one further way in which concepts that social scientists use change. We describe many things as games today, including diplomacy and tiddly winks. However, perhaps, in the future people will just use the word to refer to games played at home. There need be no particular reason for this, just a change in language and culture. Nevertheless, in the future, the link with our current understanding of the word 'game' remains, and we see this as a change in the use of the concept, rather than the invention of a new concept. This can be characterised as a situation where the concept moves across the landscape, but where there is little or no interaction between this movement and the landscape itself. The movement of the concept leaves the landscape unchanged.

To summarise, a concept can move across the landscape, and in doing so incorporates different aspects of the social world over time. Where this movement results just from the changing use of a concept, the concept and the landscape do not interact. Where there is reflexivity, the concept changes the social world, which, in turn leads to changes in the concept. This interaction between a concept and the landscape is not an all or nothing matter, and may happen to greater or lesser degrees. Where concepts change over time, this contributes to them being Nomadic because they move across the landscape (which is the traditional way of understanding a nomadic tribe). The changes in the aspects of the social world that are seen to fall under the scope of the concept over time increases the potential for disagreement between social scientists about what 'really' does fall under the scope of the concept. In a very basic sense, they are arguing about a moving target.

Nomadic concepts have many meanings

Many concepts that social scientists use have many different meanings. The concept of 'social exclusion' illustrates this. Other concepts used by social scientists also have many possible meanings; for example, Adcock and Collier note that 'background concepts' (which are the headline, pre-analysis, concepts, such as 'poverty') "routinely include a variety of meanings, the formation of systematized concepts often involves choosing among them" (Adcock & Collier, 2001, pg. 532). Similarly, Woodward writes that there are a number of different things that 'being female' can mean, for the purposes of analysing disparities in pay between males and females (2003, pg. 115). One further example is the concept of 'poverty', which may mean falling short of some minimum income, some shortfall versus society at large, or related to heath, education, or any number of other factors.

Where a concept has many meanings, this should be understood as different areas of the landscape falling under the scope of a concept. A concept with many meanings is like a tribe that exists in a number of locations. When using a concept with many meanings, social scientists need to specify which meaning they have in mind. However, in some cases this may not be possible because social scientists may not always be aware that they are using the same concept to refer to different parts of the landscape. The two following examples illustrate how this can happen.

Steel (2008, pg. 102) describes ambiguities in the meaning of '*ceteris paribus* law'. He says that "A striking feature of the philosophical literature is the variety of types of generalisation that are referred to by that label" (Steel, 2008, pg. 102). Citing Schurtz he divides *ceteris paribus* laws into two main types—exclusive and comparative. "An exclusive [*ceteris paribus*] law indicates an absence of factors that would produce exceptions to the law, whereas a comparative [*ceteris paribus*] clause asserts not that interfering factors are absent but that they are distributed identically between groups that differ with respect to the putative cause" (Steel, 2008, pg. 102). What he then describes is the debate between philosophers and economists using different meanings of '*ceteris paribus*' without awareness of this difference. The comparative meaning of *ceteris paribus* entails that some types of generalisations are relatively unproblematic. Philosophers subscribing to this meaning react incredulously to suggestions that *ceteris paribus* laws are meaningless or untestable. Steel cites a variety of literature to illustrate his point, including Earman and Roberts (1999) for arguments that *ceteris paribus* laws can serve no legitimate scientific purpose due to their open ended 'escape clauses' and Persky (1990) for an example of an economist dismissing such arguments as "foolish". Steel illustrates that different economists have a particular meaning of *ceteris paribus* in mind and their lack of awareness about other, potential, meanings of this term results in their incredulous reaction to economists who assume a different meaning. This is why they are so dismissive of the arguments of philosophers who have, without realising, a different meaning in mind. They think they are discussing the same thing, but in reality are not.

A further example is provided by McCauley, who argues that there are "at least six" different meanings of 'equilibrium' in finance and economics (2004, pg. 155). These are the idea of equilibrium fluctuations about a drift in price, the idea that market averages describe equilibrium quantities, the idea that the CAPM (Capital Asset Pricing Model) describes equilibrium prices, prices of options described by the Black and Scholes equation, the absence of arbitrage opportunities, and the idea that the market, and models of markets, define sequences of temporary price equilibria (McCauley, 2004, pg. 76). All of these differ from the way in which physicists understand 'equilibrium'. This is a concern because although many of the models used in finance and economics (in particular the Black and Scholes equation and some models of markets) apply models developed by physicists, McCauley argues that in various respects advocates of this approach to finance and economics are unaware of important differences between the use of 'equilibrium' in physics as compared to economics and finance. These differences are potentially important in understanding the extent to which the models apply to financial cases. The different meanings of concepts that social scientists use are not always transparent, even when we are dealing with concepts that have benefitted from intellectual refinement. '*Ceteris paribus*' is not a concept used in ordinary discourse, in the way that 'poverty' is.

To summarise, the concepts that social scientists use often have many meanings. In some cases, the different meanings a concept can have are transparent. In this case, the social scientists have a relatively clear picture of how a concept sits on the landscape and specify which meaning, or which aspects of the social world, they intend to discuss when using this concept. A more complicated case is when different social scientists are unaware that they are using a concept which inhabits different locations on the landscape.

Regardless of whether social scientists are aware of these different meanings, these meanings many be closer, or further apart on the landscape. For example, 'happiness' can mean a number of things, including economic security, satisfaction with family life, joy when thinking about a specific event, or general satisfaction with life. It is difficult (for me) to think about what all these meanings of 'happiness' share. They seem relatively unrelated to one another. Happiness in one area doesn't seem to relate to happiness in other areas. This can be represented as a case where the concept sits in different locations on the landscape. The aspects of the social world that we are directed towards when analysing 'economic security' are different from those we analyse when looking at 'satisfaction after a specific life event', such as marriage, or divorce. In the case of economic security this will include financial data, and attitudes towards this. In the case of event-specific satisfaction it will include self-reporting of mental states and expectations. This is not to suggest that there will be no overlap between these cases, in fact, we could be analysing the same people in each case, but that each meaning directs us towards different aspects of the social world.

This is in contrast with a concept like 'school leavers'. 'School leavers' can mean those leaving school having passed their compulsory exams, or those who, despite leaving school, are still receiving some form of education until they meet the required standards in English and maths. Both these meanings are closely related; in each case, we are talking about 16-year-old children leaving schools. Consequently, we can conclude that the concept 'school leavers' inhabits parts of the landscape that are closer together than does the concept 'happiness'.

Nomadic concepts can have many meanings, social scientists may all agree about what these meanings are, or they may not. Additionally, these meanings may be far apart, or close together on the landscape. There are two other complications relating to the different meanings that concepts used in the social sciences. The first is that there are many things we can mean, when we say that a concept has a particular meaning. To introduce a concept that will be discussed at length later in this chapter, we may argue that one of the things 'democracy' can mean is 'contestation', (in other words that there is competition and openness in the electoral process). However, there are also many things that 'contestation' can mean. For example, we might take it to mean that there are a number of candidates standing for political positions, or that there are a certain minimum number of political parties, or that the agendas of the political parties are sufficiently different, or that the media is independent of political parties, or any number of other things. A concept that is Nomadic because it has many meanings, like 'democracy', may have meanings that are themselves Nomadic for the same reason.

The second complication is that social scientists may be interested in the behaviour of individuals, as opposed to concepts like 'democracy' or 'poverty'. This is a problem of determining what the meaning of an action is. Taylor gives the example of 'the act of voting' which, depending on context, can mean "saving the honour of my party, or defending the value of free speech, or vindicating public morality, or saving civilisation from breakdown" (Taylor, 1971/1994, pg. 190). In other words, while we can observe a certain behaviour—somebody raising their hand during a vote— this action may mean any number of things. Without understanding what the action meant for the actor, or even what it meant for observers of the action, we are unlikely to get to grips with what is really going on. This situation can be represented in the Nomadic framework, as follows. When observing some behaviour social scientists may be unsure of which concepts this behaviour should fall under the scope of. Specifically, is the man raising his hand 'saving the honour of his party' or 'following orders'. Social scientists may disagree about whether these concepts, or a number of others, include this behaviour within their scope.

In conclusion, when concepts are Nomadic because they have many meanings this increases the potential for social scientists to disagree about whether specific aspects of the social world fall under the scope of this concept, and which aspects of the social world the concept covers. Specifically, when a concept has many meanings, which are agreed upon by social scientists they can usually just specify which meaning they have in mind. In other cases they might not agree about the different meanings the concept can have, or may even not be aware that other social scientists have a different meaning in mind. In the first case the concept is less Nomadic than in the second case, this is because there is greater potential for disagreement about which aspects of the social world fall under the scope of the second concept. These different meanings, regardless of whether they are agreed upon or not, may be further away, or closer together on the landscape. Additionally, each of these meanings may itself have a number of meanings. Finally, where it is possible to question what some specific behaviour means, this can be represented as a case where social scientists can disagree about which concepts this behaviour falls under the scope of.

Higher order/ methodological Nomadic concepts.

Social scientists do not just argue about concepts like 'social exclusion'. Social scientists also use other concepts that attempt to specify relationships between concepts, such as 'causation', and 'correlation'. These concepts can also be fitted into the Nomadic framework, and can be more, or less, Nomadic. Correlation is specified precisely in mathematical terms, whereas considerable debate surrounds the concept of causation. Correlation can therefore be represented as located at a single part of the landscape, with few boundary issues. The concept 'causation' is a different matter entirely, because there are many things that 'causation' can mean. Kim (2009, pg. 126) lists four approaches to analysing causation: the regularity analysis, the counterfactual analysis, the manipulation analysis and the probabilistic analysis. The regularity analysis began with Hume and Mill and sees causally connected events as instantiating general regularities between kinds of events. For example, an apple falling to the ground instantiates a more general regularity that objects fall to the earth. The counterfactual approach (Lewis 1973, 2005) says that *a* caused *b* if it is the case that had *a* not happened, *b* would not have happened. Counterfactual analysis will be discussed at greater length in Chapter 3. The manipulation analysis (Woodward 2003), which will also appear again in Chapter 3, says that causation should be understood as an ability to change, or manipulate something to produce a specific effect. In other words, *a* caused *b* if doing *a* brings about *b*. The probabilistic analysis says that *a* is a cause of *b* if the probability of *b* is greater when *a* has occurred than when *a* has not occurred.

Causation can mean a number of things, so has multiple locations on the landscape. It is also possible that new ways of understanding causation will be proposed, and that the meanings above will change as debate continues within the philosophy of science community. Concepts that can be described as Nomadic therefore include concepts that social scientists analyse, such as 'democracy', or 'poverty', and concepts that they use to understand connections between these concepts, such as 'causation'.

How is the Nomadic Concepts framework helpful?

The following section shows how social scientists can use this framework to think about concepts in a systematic way. The salient features of Nomadic concepts were discussed individually above, but in practice they will need to be considered together. The ways in which concepts used in the social sciences were described in Chapter 1 are that they are *Ballung* concepts, cluster concepts, essentially contested, multiply realisable, and vague. The following section demonstrates how these can be subsumed within the Nomadic concepts framework.

Ballung concepts

Ballung concepts "sort things into categories based on a loose set of criteria in which members of the same category do not share any specific set of features." (Bradburn, 2011, pg. 54). In other words, members of such categories have loose resemblances to one another and there are no necessary and sufficient conditions for membership of these categories. The examples cited in Chapter 1 were 'poverty' and 'disability'. *Ballung* concepts can be understood as being located in different places in the landscape (one for each of the things that these concepts can mean), these locations may overlap in some cases, but no location is shared by all meanings of the concept. Deciding on the exact boundary of each location is also likely to be difficult (because, for example, judging whether someone is, or is not, in 'economic poverty' will be contentious). Such concepts, and the different meanings associated with them may, or may not, change over time.

Cluster concepts

According to Little (1993), cluster concepts share some cluster of properties. As discussed in Chapter 1, one of his examples is 'market economies', which have many things in common, but which also have causally significant differences. Cluster concepts, according to this definition, differ from *Ballung* concepts because the different locations that the concept 'market economy' inhabit on the landscape overlap. The overlap represents that aspects of the social world that are shared by all meanings of 'market economy' (Little's common features include the existence of private companies, and well-developed corporate law). A cluster concept is likely to have boundary issues in many locations (for example, judging whether corporate law is 'well-developed' is likely to be contentious). Cluster concepts, and the meanings associated with them, may, or may not, change over time.

Essentially Contested Concepts

Essentially contested concepts can be represented as ones which have many meanings (as outlined in Chapter 1, Gallie describes this as being describable in a variety of different ways), and therefore have different locations on the landscape. However, it is unclear whether these locations overlap, or how far apart they are; is likely to differ depending on the specific concept. The concept must represent some kind of valued achievement. The concept changes over time, which indicates that it moves over the landscape and, additionally, there may be interaction between the concept and the landscape. Boundary issues are likely to be present. Essentially contested concepts may be *Ballung* concepts, or cluster concepts.

Essentially contested concepts are those which are subject to constant disagreement about what 'democracy', or 'social justice' mean. Seeing these concepts as Nomadic allows us to see why this is. When trying to specify what 'democracy' is (as described in Chapter 1, Gallie's examples of essentially contested concepts include 'democracy', 'social justice' and 'work of art'), social scientists can direct their attention towards a variety of locations in the landscape (which may, or may not, overlap), each of which is likely to have boundaries which are difficult to draw. As Gallie describes, social scientists, while recognising the validity of other points of view about what 'democracy' is, nevertheless argue persuasively for their own point of view. These concepts are also likely to change over time, meaning that the locations shift (for example, as the idea of democracy develops in the future). This means that it is possible to, reasonably, disagree about whether about the aspects of the social world that are subsumed within the concept of 'democracy'.

Multiple realisability

Concepts that are multiply realisable exist in different locations. For example, the concepts 'money', and 'market economy' are multiply realisable. In the case of 'money' this can include sea shells, fiat currency, currency backed by precious metals, letters of credit, bank deposits, and a number of other things. Some of these meanings, or locations may overlap (as in the case of letters of credit and bank deposits which both require a banking system), and others may not (as is the case between sea shells and bank deposits). These different meanings appear not to have significant boundary issues, at least as compared to other examples in this chapter, such as 'economic poverty'. Once social scientists have specified that, when they are talking about money, they really mean bank deposits, then it is relatively uncontroversial to determine what counts as a bank deposit. Some concepts which are multiply realisable may have boundary issues. Where a concept is multiply realised, but everyone agrees about what these realisations are, and each of these realisations is precise, and also do not change over time, then this concept is not Nomadic. This is best described as a case of ambiguity, where we just need to specify which of a number of things we intend to pick out when we use a concept.

Vagueness

In the philosophy of language, vagueness relates just to boundary issues of the kind where a Sorites Paradox can be generated. These are cases where it is difficult to draw a boundary between thing like heaps and non-heaps, or bald men and non-bald men. Nomadic concepts may have boundary issues of this sort as in the case where we are trying to determine whether a group of people are receiving an adequate number of calories per day.

Using the framework to think about concepts

The Nomadic concepts framework can be used to think in a more systematic way about the concepts that social scientists use. To demonstrate this 'social exclusion' will be contrasted with 'demand', and show that 'social exclusion' is more Nomadic than 'demand'. This is because there is more potential for disagreement about which aspects of the social world fall under the scope of the concept 'social exclusion' than 'demand'.

'Social exclusion': As decribed above, 'Social exclusion' has many different meanings, many of which do not have precise boundaries. The concept of 'social exclusion' has also changed over time, and can reasonably be expected to change further in the future. This concept can be described as Nomadic, and it is Nomadic in a variety of ways. These characteristics mean that analysis of social exclusion hasn't resolved what this concept 'really is'.

Now consider the rather different concept of demand, from economics: This is a more rarefied concept, which has benefitted from some degree of intellectual refinement. Recall from Chapter 1, Kincaid (2004) uses the law of supply and demand to justify his view that there are laws in the social sciences, so let us compare 'social exclusion' to 'demand' to see whether there is something different about these two concepts that justifies greater faith in generalisations about 'demand' than 'social exclusion'. With the analysis in the first portion of this chapter it is now possible to see how the concept of 'demand' differs from 'social exclusion'. The concept of demand in economics is the amount of a good or service that a customer is willing to buy at a particular price. Now, admittedly, some imprecision results from the notion of 'willing to buy'. I may be willing to buy 400 T-shirts at £5 each, but I may, due to time, transport, or other constraints only buy 200 T-shirts. Is my demand for T-shirts 400 or 200 at £5? In other words, do we mean actual demand, or potential demand? However, these two meanings are closely related because both apply where people are buying and selling something, and just measure, in different ways, the amount they want at a certain price. So, although the concept 'demand' may have different locations, these are close together, if not overlapping. Both meanings are also themselves well defined, although taking demand as the actual amount bought is the better defined of the two—it is synonymous with the number of T-shirts bought at a certain price. Determining a 'willingness' to buy at a certain price is harder to gauge, but, in principle is roughly determinable by asking people. The concept of 'demand' is unlikely to change over time. 'Demand' is not as Nomadic as 'social exclusion'. Consequently, there isn't extensive debate about what 'demand' is.

The implication of this is that, when social scientists discuss 'demand' they largely agree about the aspects of the social world that are relevant to analysing 'demand'. The concept occupies are relatively clear location on the social landscape. Social scientists are much less likely to agree about the aspects of the social world that are relevant when analysing 'social exclusion' because this concept is more Nomadic.

Before moving on, it is important to note that the framework outlined above is intended to help social scientists think about the concepts they are using in a systematic way. It is not intended as a tool to determine, once and for all, the structure of social science concepts. It is unlikely that social scientists will agree on this. For example, it is possible to disagree with the characterisation of 'social exclusion', or 'demand' sketched out above. However, with this framework in hand, it is possible to discuss this disagreement in a structured way by discovering how different social scientists specify a concept in terms of meanings, change over time, and boundary issues.

Implications

The following section discusses two implications of accepting this framework. The first is that it explains why the social sciences have had so much difficulty finding regularities; and the second is that it suggests that current proposals for making concepts more precise are unsatisfactory.

i. Making generalisations

Where there is the potential to include a wide variety of phenomena within the scope of a concept it is more difficult to formulate generalisations using such concepts. 'Generalisations' mean generalisations that are widely accepted by a heterogeneous group of social scientists. Tucker (2009) provides criteria that are helpful in explicating this idea; these are lack of coercion, heterogeneity (acceptance by social scientists from different backgrounds, who have few connections with one another, and who use different academic approaches), and size (Tucker, 2009, pg. 27-36). In other words, a generalisation can be seen as 'widely accepted' when this acceptance is uncoerced, comes from heterogenous social scientists and reflects the views of a sufficiently large number of social scientists. Tucker notes that judgements about whether a sufficiently large number of social scientists accept a generalisation will have to be made on a case by case basis.

Given this, how do we think about the statement 'High income inequality makes revolutions more likely'? Firstly, 'more likely' is a methodological concept that is Nomadic because there are many things this can mean. We can interpret this statement in a number of ways. We could interpret this statement as saying that 'High income inequality makes revolutions more likely than not'. In other words, that revolution is more likely than not having a revolution. A different interpretation can be taken from Northcott (2008), who argues that contrast classes are critical to making sense of causal claims. In other words, 'c causes e' should often be understood as meaning 'c rather than C* causes e rather than E*' (Northcott, 2008, pg. 112). The evaluation of counterfactual claims is central to his argument, which will not be addressed in detail here because Chapter 3 is devoted to the analysis of counterfactuals. Nevertheless, Northcott's wider point is relevant here. Northcott writes "often, at least in the special sciences, we have a *choice* of generalisations or models with which to analyse a case. This may impact in turn on choice of contrast class and thus, sometimes, make all the difference to causation" (Northcott, 2008, pg. 122). Later he says that linguistic stress and choice of vocabulary can "disambiguate which contrast classes are appropriate" (Northcott, 2008, pg. 122). Illustrating this with the current example suggests the following; when assessing the statement 'High income inequality makes revolutions more likely' we are implicitly contrasting 'revolution' with something else. This could just be 'not a revolution', or it could be something like 'seeking redress through political means', or even 'emigration', the potential contrast classes are numerous and will depend on context. In one case, we might be concerned with the reasons for revolution rather that peaceful protest, in another, with revolution rather than migration. The significance of this is that when social scientists are faced with a statement such as 'High income inequality makes revolutions more likely' there are a number of ways they can interpret 'more likely'. Furthermore, these different interpretations will affect whether they judge this statement to be true, or false.

Settling aside worries with 'more likely', as far as the generalisation as a whole is concerned we can, probably, off the top of our heads, think of a number of revolutions where income inequality appears to have been an important factor. However, are we willing to endorse this as a general statement? Probably not, and part of the reason why is that the concepts included in this generalisation are Nomadic; they have many meanings, which are not closely related to one another and lack precise boundaries.

The relevant concepts here are 'income inequality' and 'revolution'. Both of these terms can mean different things. Does 'income inequality' mean the actual difference between the highest and the lowest, the perceived income inequality, how rich the richest are, or one of a variety of other possible things, none of which are particularly well defined. 'Revolution' can mean political, violent, drastic democratic change, and a variety of other things, none of which are particularly well defined either. This means that when we think about this generalisation we can conceive of the relationship between the concepts 'income inequality' and 'revolution' being located on the social landscape in a wide variety of ways, all of which point us towards different conclusions regarding the plausibility, and strength, of this relationship. Anyone considering this generalisation can consider wide ranging and, possibly contradictory, interpretations. Even if analysis is limited to the consideration of past revolutions, different social scientists will disagree about whether specific historical events were, or were not 'revolutions' and whether 'income inequality' was, or was not, present. Of course, different social scientists may propose that this statement is true, or false, but it is unclear how we are to choose between them. Assuming that they both present well-chosen examples and argue them well we are unlikely to be able to come down conclusively on one side or the other. The opposing social scientists may be focusing on different aspects of the social world, all of which fall under the scope of the relevant broad concepts. The conclusion to the debate may therefore be 'sometimes yes and sometimes no'. Neither side in the debate is 'wrong', because what appears to be a debate is actually not about exactly the same things.

In contrast, we are happier to endorse the generalisation 'If the price of good X rises, then demand will fall'. The relevant concepts here are 'price' and 'demand'. The concept of 'price' is located more exactly on the social landscape. Admittedly, there is room to argue about differences in prices in different locations, or, if it is changing rapidly, which measure of price to use. 'Demand', as argued above, is also not very Nomadic. It is possible to focus on meanings and definitions of these concepts that make them synonymous with their numerical measure. Although we may dispute the strength of the relationship between price and demand, note the exceptions to it (as when price is taken to indicate the quality of a good), and argue about what social phenomena are relevant to the specific case in hand. For example, when assessing whether demand for a certain good will fall if the price rises, it may be debatable whether we also need to include price changes in substitute goods in our analysis. Neverthelesss, the generalisation is less controversial than the generalisation about income inequality and revolutions. The aim of this chapter has been to explain why this is true. While we may say that the concepts 'income inequality', 'revolution' and 'demand' all belong to the 'odd' concepts that social scientists use, the Nomadic framework allows us to see that they are not all 'odd' on the same way. These concepts are more or less Nomadic. The concept of 'demand' is less Nomadic than 'income inequality' and 'revolution'. These differences allow for greater disagreement about what social phenomena are to be reasonably included within the scope of the concepts 'income inequality' and 'revolution' than within 'demand'.

ii. Increasing precision

The second implication of this is that this framework allows us to understand when proposals for making concepts more precise are likely to succeed, or not. Much of the recent social science literature focuses on the methodological problem of how social science concepts should be made more precise (for example, see Crasnow (2015), Cartwright & Bradburn (2011) and Goertz (2006, 2008) and Adcock & Collier (2001)). This takes it as given that the concepts the social sciences deal with are 'odd' in some way, but sidesteps further analysis of this oddness by focussing on practical issues. It is often suggested that social scientists just stipulate which meaning of a concept is relevant. Consider 'disability', which may "mean different things depending on whether we are talking about particular individuals, about a policy goal, a variable in a psychological theory, or a characteristic of a group of individuals" (Bradburn & Cartwright, Discussion paper, pg. 3). Later, they say "whether a feature counts as in or outside the concept, and how far, is context and use dependent" (Bradburn & Cartwright, Discussion paper, pg. 4). And later, "If we want to turn the […] concept of poverty into a precise and unambiguous one, poverty concepts are bound to proliferate" (Bradburn & Cartwright, Discussion paper, pg. 7). In other words, depending on the purposes of our research, we focus only on certain aspects of a concept such as poverty, perhaps 'income per person below £15,000 per annum', and therefore consider features such as income and expenditure, rather than educational level or ethnic background. In any particular case, the features that interest us will differ. The suggestion is not that 'income per person below £15,000 per annum' is what 'poverty' *is*, but that this is the meaning of poverty that is relevant in this situation. In other situations we will use other aspects of poverty.

However, this is not always possible. As we saw above, there may be fundamental disagreement about which meaning is appropriate in a particular context. Bradburn and Cartwright do suggest an alternative approach that may work in these sorts of situations. They suggest using a table of features which lays out "the dimensions along which the family resemblances in question lie" (Bradburn & Cartwright, Discussion paper, pg. 10). An example of this the EU measure of 'social exclusion' which is a three-tiered measure. The first-tier lists seven or eight factors that pick out aspects of social exclusion in Europe, such as inequality. The second tier has a larger number of factors intended to make the characterisation more comprehensive. The third-tier lists society specific features (Bradburn & Cartwright, Discussion paper, pg. 7). In other words, we are interested in 'social exclusion', but there is no precise definition that satisfies our interest in this case. We are interested in 'social exclusion' as a whole. There are many things that this can mean but there appears to be nothing that all instances of 'social exclusion' share. We therefore must use a multi-faceted set of data to try to capture what we mean by the concept. In theory, this set of data should incorporate what most people think are the relevant similarities between instances of 'social exclusion', up to some sensible limit set by data gathering and computational constraints. This would allow for the inclusion of many different meanings but this does not quite solve the problem where people fundamentally disagree about whether something really is an aspect of poverty, because they would simply disagree about whether something should be included within the list of factors or not. I will now discuss an example from Crasnow (2015) in more detail to illustrate my worry with this overall approach for making social science concepts more precise. The worry is that it is not sufficient to just consider which meanings of a concept we might be interested in because these meanings are often Nomadic concepts themselves, so the aspects of the social world that can be included within the scope of these meanings may still be large enough to give rise to disagreement.

Crasnow advocates pluralism in the measurement of concepts, writing that "one standard may not be desirable and the sort of accuracy such a measure implies may not be appropriate" (2015, pg. 1). In other words, there may be no definitive answer to the question 'What is the best measure of democracy?' Instead, there may be a variety of measures that are appropriate in different contexts. While this is relatively uncontroversial, the rest of her paper illustrates that the process of determining how to measure democracy is more complicated than it appears. Crasnow notes that "ordinary use" concepts are described as "vague and ambiguous" (Crasnow, 2015, pg. 8). As Chapter 1 pointed out, the use of the word 'vague' is unfortunate because it has a very specific meaning within the philosophy of language, nevertheless, for the ease of expositon, it will be used here because this is the word that Crasnow uses. However, it should only be taken to mean a general uncertainty, or unclarity. The approach she discusses takes a concept, like 'democracy' or 'peace' and narrows it down to make it more precise. In other words, because these concepts are 'vague and ambiguous' we need to specify exactly what we mean, and the work we intend the concept to do. The suggestion is that social scientists need, somehow, to make the concepts less vague, taking into account the research they intend to do. There are a number of proposals for how to do this (see Goertz (2006, 2008), Crasnow (2015), Adcock & Collier (2001), Bradburn & Cartwright (2011)) but all aim to make the concept clearer so that it can be used for the purposes of research. In their differing ways, these philosophers suggest how to break the concept down into a number of levels. Most of the authors use three levels, but Adcock and Collier use four. This multi-level process is an attempt to refine the concept to make it clearer. The process is as follows.

At the first level is the ordinary use concept. This is the pre-analysis, everyday use of the word. At the second level, researchers think about which meanings of the concept are relevant, given what they want to research. For example, the concept of 'democracy' has many meanings. A democracy may mean the presence of 'political liberties', or 'popular sovereignty', or 'contested elections', or 'competition', or a variety of other things. The number of meanings of 'democracy' that are relevant in a particular context is likely to vary. For example, if we are concerned with ancient democracies we may focus on different meanings than if our interest is on post 1945 democracies.

At the third level, once we have decided which meanings are relevant for our purposes, we can decide on indices, or datasets, that allow researchers to establish the presence or absence of the attributes of democracy with which we are concerned. Adcock and Collier (2001) note that working out these levels is an iterative process, rather than a one-off exercise. Goertz (2006, pg. 8-9) provides a number of examples of how the concept of 'democracy' has been analysed. The key concept we begin with is 'democracy'; using Arat's framework the second level meanings in which we might be interested are 'participation', 'competitiveness' and 'coerciveness', while at the third level the data representing these second level meanings include 'competitiveness of nominations', 'party legitimacy', and 'party competition'. An alternative analysis of democracy, which Goertz takes from Coppedge and Reinicke, has one second level meaning, which is 'contestation', and the third level data are indices or datasets that show the extent of 'free and fair elections', 'freedom of expression', 'freedom of organisation' and 'pluralism of the media'. Examples might include the number of parties standing for election, or the number of independent media sources. (Goertz, 2006, pg. 8-9).

In other words, we begin with a pre-analytic, concept such as 'democracy'. Then we think about what we mean, in our particular context, when we use this concept. Once we have specified what we mean we can look for indicators, or datasets, which may be qualitative or quantitative, to gauge the presence or absence of the things in which we are interested. For example, we may specify that in a particular context 'democracy' means 'contestation'- that the political process is contested. Then, we need data that evidences the presence of absence of contestation. Following Coppedge and Reinicke we can use 'freedom of expression' and 'free and fair elections'. We then analyse any data we have access to, including the number of candidates standing in elections, and the number of independent media sources, and aggregate these, yielding a quantitative measure of the extent to which a political process is 'contested'. In this way, we have moved from the initial concept 'democracy' to a more specific set of data, such as the number of independent media sources and the number of candidates standing for election. This allows us to compare the extent to which countries or states are democratic and to look for explanations, or connections with other concepts.

This process appears sensible, but unfortunately, there are two problems. Firstly, the process of narrowing down the concept by specifying what is meant in a particular context appears to have settled what we mean by our concept. We say 'By 'democracy' I mean 'contested elections''. It appears therefore that we have settled what we mean by democracy by using other terms whose meaning is clear, or at least, more clear than 'democracy'. However, this is not the case. We could just as easily have 'contestation' or 'contested elections' at the first level and specify what we mean by 'contestation'. This proposal takes it as given that the concepts used at the second level are, in fact, less 'vague and ambiguous' than the first level concepts, when there is no reason why this should be so. 'Contestation' may be an aspect of 'democracy', but this does not mean that it is, *necessarily* a more precise concept. At the outset, it is important to reiterate how precision should be understood within the Nomadic framework. When a concept is precise there is little, or no, potential for disagreement about the aspects of the social world fall under the scope of a concept. The concept of 'demand' is more precise, or less Nomadic, than the concept 'poverty'. The following section illustrates the worries with Crasnow's approach.

Consider Crasnow's discussion of 'democracy' (2015). She describes 'democracy' as follows:

> "With any latent concept—which is an abstract concept—the appropriate concrete attributes to associate with that concept need to be specifically identified. These attributes will either be observable themselves or linked to subcomponents which will be observable and serve as indicators." (Crasnow, 2015, pg. 4).

This illustrates her three-level approach—at the first level is the abstract latent concept, at the second level are the relatively more concrete attributes, and at the third level are the data, in her words 'indicators'. Crasnow says, "it is generally agreed that there are two main sets of attributes through which democracy should be identified: competitiveness (contested elections) and participation (inclusiveness of the electorate)" (Crasnow, 2015, pg. 4).

For Crasnow, democracy implies competitiveness and participation. It appears that we have settled what we mean by 'democracy' by breaking it down into two constituent attributes—competitiveness and participation. This does not follow because, just as 'democracy' has multiple meanings, so do 'competitiveness' and 'participation'.

'Competitiveness' could mean the number of candidates standing, the ease with which candidates can put themselves forward for election, or the extent to which candidates' platforms differ. Crasnow offers two other potential meanings: the right to form political parties and freedom of the press (2015, pg. 5). Unfortunately, there are also a variety of potential meanings of 'freedom of the press'. This process of narrowing down a concept only works when we are not narrowing our concept down with other concepts that have multiple meanings. Crasnow is aware of this problem, for she says that:

> "Even if the concepts are seemingly well-specified questions of their consistency across context can arise. The term 'universal suffrage' does not mean the same thing in 1890 that it means in 1973. Questions may also arise about whether the concept is used in the same way for cross country comparative research" (2015, pg. 6).

I do not believe that democracy is well specified when we identify it with 'participation' and 'competitiveness'. It is not well specified because 'participation' and 'competitiveness' have multiple meanings. It is this poor specification which results in terms like 'participation' applying inconsistently over time.

When a concept has many meanings we can specify which meaning we have in mind. However, a concept cannot be made more precise with a concept that, although an aspect of the original concept, also has many meanings, especially when these meanings also have boundary issues, and change over time. This is because, as in the discussion of poverty, different social scientists may have different things in mind when they use such a concept. A natural reply is that Crasnow describes, if not *the* way, then at least *one* way that democracy can be measured. So, as long as we are clear about what she means, we can follow the logic of the proposal for measuring 'democracy'. Unfortunately, if the meaning Crasnow focuses on does not match our idea, or anyone else's idea about what 'participation' means, then the resulting analysis will have limited applicability.

The importance of this is illustrated by Crasnow's example distinguishing between 'universal suffrage' and 'universal male suffrage'. She says, citing Paxton, that a number of conclusions about democracy no longer hold if 'universal suffrage is treated as universal male suffrage' (2015, 6). Depending on whether we mean 'universal suffrage' or 'universal male suffrage', our conclusions about 'democracy' change. For example, our classification of a country as 'democratic' may differ depending on whether male suffrage is sufficient for a positive classification, or whether we need wider suffrage. Crasnow concludes that determining when we have the process of conceptualisation 'right' is an iterative process. However, recognising that the problem of what we mean reaches through all levels of conceptualisation suggests that there is no 'right' answer.

A natural reply is that this is only a concern when we use our concept over extended time. It is clear that 'universal suffrage' does not mean the same thing in 2016 as it did in 1890. Can this issue not be sidestepped by focussing only on one time-period? The following example suggests not. Crasnow describes how Alvarez, Cheibub, Limongi and Przeworski focus on contestation, rather than participation, in their analysis of democracy. Crasnow says that "According to Munck [they] argue that this is a legitimate preference since their concerns are really about the post-1945 era and so universal participation (universal suffrage) can be assumed" (Crasnow, 2015, pg. 4). In other words, there are two things that democracy can mean—contestation and participation. The authors focus on just one of these—contestation, because the other meaning—participation, is taken to be present in all the democracies under consideration in this study, because they have universal suffrage. This is only valid if 'suffrage' and 'participation' mean the same things, but it isn't clear that they do.

We can have universal suffrage in a population, but if certain groups within the population feel alienated from the political process, leading to their non-participation in elections, then 'participation' may be low, despite universal suffrage. This sub-group has the right to vote, but because they don't see their interests represented by any party, they don't participate in the political process. For me, 'participation' is not equivalent to 'universal suffrage'; it means not only the right to vote, but a belief that the political process is responsive to, and reflective of, the views of the population. If you agree with me about these different meanings, then Alvarez, Cheibub, Limongi and Przeworski are wrong to focus just on contestation in their analysis of 'democracy'. They have used a concept with many possible meanings without an awareness of this. Consequently, their research is problematic for someone not sharing their views about what participation means. For example, an alternative approach, motivated by the idea that participation does not just mean suffrage, would include participation in an analysis of post-1945 democracies. This example illustrates that it is possible to disagree about the relevant meaning of a concept, even when the focus is on a fixed time period.

The problem with Crasnow's proposal for refining social science concepts is that a solution is sought in specifying meanings of a concept in a particular context. This is taken to solve the problem of what a concept means. But, as I have argued, it has not done so. The concepts used at the second level, or, in other words, the aspects used to specify what concepts mean are in many cases no more precise than the initial concept. The implication of this is that the conclusions of research are limited, and sometimes unconvincing, for anyone who has other meanings of these concepts in mind. This becomes clearer if we think about the ways in which these concepts are Nomadic.

The concept 'democracy' has been discussed throughout this chapter, which has suggested that it is very Nomadic, the concept has many meanings which have different locations on the social landscape, these meanings are changing over time and are subject to boundary issues. Furthermore, many of these meanings are themselves Nomadic.

A natural reply is that, while this problem may apply to the analysis of 'democracy' by certain authors, it does not apply universally. This is true, but is not an objection. Just because 'democracy' is Nomadic, this does not mean that social science concepts are equally so. It is a mistake is to believe that all social science concepts are Nomadic in the same way, and to apply the same methodology for making them precise to them all. The following paragraph discusses an example from Ragin, which suggests that some concepts can be made more precise by using other concepts that do not have multiple meanings.

Ragin (2008) aims to show how social scientists can quantitatively assess the degree to which some phenomenon fits into a set. The degree of inclusion in a set can range from 0 (full exclusion) to 1 (full inclusion). His contribution to the debate is that social concepts can be characterised as fuzzy set relations. For example, we can think about developed countries as a subset of democracies, rather than just analysing the correlation between development and democracy. Thinking about relationships in terms of set relations can often reveal relationships that are not apparent from pure correlation data. This is a useful insight, but is not the concern here. What is of interest is the process by which Ragin establishes which countries count as developed, and to what extent.

In one of Ragin's examples 'high per capita national income' is used as an indicator for being a developed country'. In this example, he calculates the degree of membership in the set of 'developed countries' using per capita national income data (Ragin, 2008, pg. 86-90). However, the concept of interest is 'developed country', which can mean a number of things, of which high 'per capita national income' is only one. Given the wider spirit of Ragin's proposal he would no doubt agree with this, but his example is used only to motivate the following point. Other meanings of 'developed' could include a democratic political system, a relatively equal distribution of income, a high level of healthcare provision, or high educational standards. In this respect, his approach is the same as Crasnow's—this example concentrates on one of many possible meanings of the first level concept. In this case, it is debatable whether we are actually still talking about 'developed countries' as opposed to merely 'countries ranked by their level of per capita national income'. If per capita national income is what matters in our particular situation this is fine, but if we really are trying to talk about 'developed countries' it is unclear that just fixing on one meaning, and defining this relatively precisely, helps us to understand the wider concept.

Nevertheless, there is an important contrast with Crasnow. Setting aside our worries about the extent to which 'developed' means 'high per capita national income', the advantage of using 'high per capital national income' is that this concept is not plagued by multiple meanings in the way that 'participation' is. Doubtless we can still argue about what should be included in 'income', and even about whether 'per capita' should include transient populations, so the meanings of these terms are not *totally* fixed; however, we can, with some margin of error, calculate 'per capita national income' for most countries. The methodology for doing so is relatively clear. The contrast with 'participation' is that whereas 'participation' can mean a number of very different things (such as suffrage, or engagement with the political process), 'per capita national income' can mean a number of closely related things. 'Per capita national income' is a concept that is more synonymous with its numerical measure. The difference, of course, is a matter of degree.

There is something different, therefore, between the concepts 'democracy', 'contestation' and 'per capita national income'. This difference is important because when making the concept ‘democracy’ more precise we do not succeed in doing so by using an alternative concept, ‘contestation’ which is also Nomadic. But we do succeed in making the concept ‘developed country’ more precise when taking it to mean ‘per capita national income’. ‘Per capita national income’ is not Nomadic in the way that ‘participation’ is. This is important because if social scientists want to make concepts more precise, then understanding the different ways in which concepts can be Nomadic allows social scientists to assess when particular methodologies are appropriate, or not. It is unacceptable just to say that social concepts are ‘vague’ and note that they can be understood in a variety of ways. This illustrates the difference between Crasnow’s suggestion that ‘democracy is understood in terms of ‘participation’ and ‘contestation’ and Ragin’s suggestion that ‘development’ is understood in terms of GNP. A Nomadic concept can be made more precise by focusing on one or two meanings, or aspects, of a concept but only when these meanings are themselves not just other Nomadic concepts. Furthermore, when a concept is Nomadic partly because there is nothing that all instances of it share, different social scientists may not even agree that the chosen aspect, or meaning, of the concept really is a meaning of this concept. Understanding the structure of social concepts in greater detail should enable greater precision in social science research.

Chapter 3. Counterfactuals

Introduction

Counterfactuals are used in a variety of contexts. Tetlock and Belkin (1996) note that historians have been using counterfactuals for at least two thousand years. They say, "Counterfactuals fuelled the grief of Tacitus when he pondered what would have happened if Germanicus had lived to be Emperor" (Tetlock & Belkin, 1996, pg. 3). In more recent times, Fearon suggests that "When trying to argue or assess whether some factor A caused event B, social scientists frequently use counterfactuals" (1996, pg. 39). Tetlock and Belkin go further, suggesting that counterfactual analysis is "unavoidable in any field in which researchers want to draw cause-effect conclusions but cannot perform controlled experiments" (1996, pg. 6). Nevertheless, historians are suspicious of counterfactual analysis, believing that it is often speculative. Many of the attempts to make counterfactual analysis less speculative involve making use of evidence in assessing counterfactual claims, or making the terms used more precise. While this is a worthwhile endeavour, this chapter argues that there is an additional problem faced by those wanting to use counterfactual analysis in the social sciences. This problem is that counterfactual analysis is particularly problematic when dealing with Nomadic concepts. When dealing with Nomadic concepts, advice to be more precise is beside the point. The following section illustrates this problem in general terms, before moving on to a more detailed analysis of counterfactuals in the remainder of the chapter.

Yuen Foong Khong discusses the counterfactual: 'If Britain had confronted Hitler with the threat of war over Czechoslovakia, Hitler would have backed down.' This is coupled with the further consequent that 'World War II might not have happened'. (Khong, 1996, pg. 95). Britain's Prime Minister, Chamberlain, was hesitant to risk war with Germany in 1938. According to Khong, he had three reasons for this; firstly, the memory of World War I was still fresh, secondly, Britain was militarily unprepared for war in 1938, and thirdly, Chamberlain had a high belief in his own ability to find a diplomatic solution to Hitler's expansionism. Khong notes that Churchill advocated a much stronger response to Germany, a view that was shared by other politicians- Eden and Cooper. He concludes that it was very possible for Britain not to have appeased Hitler at Munich (Khong, 1996, pg. 99-105).

This counterfactual is therefore one that can reasonably be asked, because Britain might not have appeased Hitler in 1938. This is in contrast to counterfactuals that are less possible, such as those asking what would have happened at Waterloo if Napoleon had had bombers. The counterfactual about British appeasement is not just an interesting academic exercise; the belief that this counterfactual is true, that Hitler would have backed down if he had been confronted, influenced US foreign policy post World War II. Khong writes that "A recurrent theme on post-World War II American foreign policy is the necessity of avoiding another Munich". Examples include debates over US policy towards Korea, Vietnam and Bosnia. Khong writes:

> "From Harry Truman's equating inaction over North Korea's invasion of South Korea to a mistake of Munichlike proportions, to Lyndon Johnson's portraying the Vietnam War as a war to prevent future Munichs, to recent US mutterings about the need to distance itself from the Munichlike policies of Britain and France towards Bosnia, the Munich analogy has served as a major script of the likely course of events if the United States failed to do X." (Khong, 1996, pg. 96)

Asking whether Hitler would have backed down, and whether World War II would not have happened, is a question that it makes sense to ask because the antecedent could plausibly have happened. Moreover, deciding whether this counterfactual is true or false has important consequences for the political landscape post-World War II. The belief that appeasement would have led Hitler to back down, and possibly avoided World War II has influenced US foreign policy, and may do so again in the future.

Counterfactual analysis remains controversial, however. Fearon argues that, in many cases, we simply cannot know what would have happened in counterfactual cases. He advocates restricting counterfactuals where the antecedent and the consequent are close to one another in time, and linked by a few, well understood, causal links (Fearon, 1996, pg. 66). Speculating about the occurrence of World War II following a confrontation at Munich does not obviously meet this criterion. There is no well specified causal link between the events at Munich and World War II. However, in this case, analysing this counterfactual is useful, despite our inability to decide what would have happened. Khong outlines three counterfactual scenarios following an imagined confrontation with Hitler at Munich:

i. Hitler would have started a war in 1938.
ii. Hitler would have started a war in 1938, but Hitler's enemies within Germany, specifically those in the military, would have deposed him.
iii. Hitler would have backed down.

Khong concludes that it is unclear which of these scenarios would have occurred (1996, pg. 116). He is in agreement with Fearon that we cannot know what would have happened. Nevertheless, although we cannot pick between these scenarios, they are important because, he argues, acknowledging the existence of more than one scenario is sufficient to throw some doubt on the uncritical assessment of this counterfactual by post war political actors. He writes, "This claim raises serious questions for those—scholars and policy makers—wont to advocate standing firm as a general rule of diplomacy because history "teaches" that a more resolute England in the 1930s would have "certainly" caused Hitler to back down." (Khong, 1996, pg. 116). It seems therefore, that even when it is impossible to judge whether a counterfactual claim is true or false, giving reasons for thinking it true or false may be useful in and of itself, if only to reduce confidence in the truth of any particular counterfactual scenario.

To summarise, Khong suggests that there are three possible counterfactual scenarios following a counterfactually more confrontational Chamberlain. He also gives reasons for thinking that the plausibility of these different scenarios has implications for US foreign policy, because US foreign policy has often relied on a simplistic understanding of this counterfactual. These implications take some of the heat out of Fearon's worry that we simply cannot know what would have happened; it is sufficient to know that a number of different things might have happened, not all of which support more aggressive foreign policy in this, or other, situations. Counterfactual analysis may therefore be a useful exercise. In this particular case, the existence of a number of possible counterfactual scenarios is a virtue, because it throws doubt on a simplistic analysis of this counterfactual.

This chapter will show that generating multiple counterfactual scenarios is possible for many counterfactuals, not just the counterfactual about confronting Hitler. Specifically, when social scientists analyse counterfactuals which include Nomadic concepts, it is possible to include a wide variety of social phenomena in the scope of a concept, and thereby generate a large number of plausible counterfactual scenarios. While this is a virtue in the Munich example, in many other cases social scientists do want to determine what the truth value of a counterfactual is; particularly when they seek to apply Woodward's (2003) manipulationist account of causation to the social sciences.

Khong specifies three scenarios, but it is possible that there are many more plausible counterfactual scenarios because the concept of 'confronting' Hitler is a Nomadic concept. There are numerous ways in which a person can be confrontational. Even just restricting ourselves just to the situation stipulated in this counterfactual, the concept is Nomadic. Khong writes that there were three men who were willing to confront Hitler in 1938; Eden, Churchill and Cooper. The antecedent of the counterfactual requires just that "their recommendations fall on receptive ears and have a chance of being implemented." (Khong, 1996, pg. 113). This is most easily achieved if one of these three were Prime Minister, but can also be achieved if two of the three had been members of the cabinet in 1938. There are many things that confrontation can mean in this case. One of Churchill, Eden or Cooper becoming Prime Minister, or any two of them being in the cabinet and thereby influencing policy. Khong writes as if these two options are intersubstitutable. He discusses the option of two of them being in the cabinet, as opposed to one being Prime Minister, as an alternative formulation of the antecedent to allay worries that making one of the three Prime Minister is too great a change to make to the historical record. He is right that these are both ways of satisfying the antecedent of the counterfactual because Hitler would have been 'confronted'. He writes, "Had one of them been prime minister in September 1938, England would most probably have confronted Hitler with war over the Sudetenland" (1996, pg. 112) and later "Had two or more of the Churchill-Eden-Cooper trio been members of the Chamberlain cabinet in September 1938, the chances of Britain's confronting Hitler would have been greatly increased." (1996, pg. 113-114). However, depending on which of these options occurs, and which of the three are involved, Hitler may have been 'confronted' with war in different ways, which may have had further counterfactual consequences.

For example, De Mesquita suggests that, even in 1939, Hitler doubted that Britain and France were prepared to fight to protect Poland. Hitler believed not only that Chamberlain's government lacked the resolve, but that Britain's military was unprepared. An important factor in understanding the counterfactual is therefore not only a verbal 'confrontation' but another meaning of 'confrontation' suggested by military readiness for war. De Mesquita suggests that more rapid rearmament would have been a consequence of Churchill becoming Prime Minister, and that this would have affected Hitler's estimates of the likelihood of war. (1996, pg. 221). What is less clear is whether this rearmament would have happened, had two of Churchill, Eden or Cooper been members of the cabinet; or whether Cooper or Eden would have rearmed to the extent that de Mesquita believes Churchill was prepared to; or whether Churchill would have succeeded in his plans for rearmament in 1938. These different options may give rise to different counterfactual scenarios resulting from 'confrontation'. Each of these possible scenarios are compatible with a 'confrontation' occurring in 1938, because this concept is Nomadic and can encompass a wide variety of social phenomena, even within the context of this example. These different counterfactuals may each suggest that the counterfactual is true, or false.

De Mesquita also illustrates the potential that, apparently small details, have for the historical record. Chamberlain did confront Hitler in 1939. He made a statement in the House of Commons on March 31st, 1939 with which he intended to demonstrate his resolve to support Poland militarily. De Mesquita argues that it was unlikely to do so because Chamberlain referred to the threat to Poland's independence, rather than its territorial integrity. Indeed, de Mesquita cites Speer as writing that Hitler did not believe that Britain and France would fight, even after the official declaration of war (de Mesquita, 1996, pg. 219-220). Both de Mesquita and Khong provide evidence that Hitler viewed Churchill, Eden and Cooper in a different light to Chamberlain, and believed that they were willing to fight a war. Nevertheless, this example illustrates the importance of considering not only what is done: a 'confrontation', but how the confrontation is carried out. In other words, whether by 'confrontation' we mean a show of force that was unambiguous, or one that left room for doubt in Hitler's mind.

The concept of 'confronting' Hitler is therefore one that can be understood in a variety of ways, even just within the context of just this example. It is a Nomadic concept because it can mean many things. It can mean a confrontation of the kind Chamberlain gave in 1939, which might have made very little difference to Hitler's actions. De Mesquita highlights the importance of readiness for war in Hitler's calculations. This suggests that 'confrontation' may mean not just diplomacy, but more rapid rearmament. It is also plausible to think that if it had been Eden or Cooper in the cabinet the confrontation might have occurred in a different way than had it been Churchill. This is important because these differences lead to differences in the counterfactual scenarios that we believe possible. In turn, these different scenarios suggest that the counterfactual is true, or false. Nomadic concepts also appear in the possible counterfactual scenarios that Khong proposes. Specifically, they refer to a 'war', and that Hitler would have 'backed down'. There are, once again, a number of things that these concepts can mean, within the constraints of this example.

The remainder of this chapter outlines the problems Nomadic concepts pose for counterfactual analysis in the social sciences in more detail. Unlike in Khong's example, social scientists often want to know what the truth value of a counterfactual is. I will argue that, often, attempting to determine the truth value of counterfactuals leads to multiple, contradictory conclusions, which means that it is impossible to assign truth values to many counterfactuals in the social sciences. These multiple, contradictory, conclusions are, I will argue, enabled by the Nomadic nature of many social science concepts. This is a problem, insofar as social scientists seek to apply Woodward's account of causation to analyse causal claims in the social sciences. In order to demonstrate this, it is necessary to first review the philosophical background to counterfactual analysis and how it has been applied to the social sciences.

Outline of the remainder of the chapter

I begin by introducing Lewis's analysis of counterfactuals and two frequently discussed problems that counterfactual analysis faces—judging closeness, or similarity of possible worlds, and backtracking[4]. I argue that backtracking is not a problem for counterfactual analysis in the social sciences *per se* but that the Nomadic nature of many social concepts enables us to fill in backtracking scenarios in a variety of different ways. Following Reiss (2009), I argue that differences in how we decide to backtrack make assessing the truth value of a counterfactual difficult. However, extending this analysis, I argue that this flexibility is problematic even in cases where we do not need to backtrack because the same flexibility exists in assessing the consequences of a counterfactual change. I apply this to Woodward's (2003) social science examples and argue that they do

[4] Backtracking describes the changes we need to make prior to a counterfactual change to enable that counterfactual change to happen. For example, when considering the counterfactual that I drank orange juice rather than whiskey with breakfast, I need to imagine going back in time and changing the historical record prior to breakfast, such as by saying that I bought orange juice the day before, in order to make this counterfactual change possible.

not all work in the way he believes they do. Finally, I present a disagreement between two historians about a counterfactual to illustrate the significance of the Nomadic nature of concepts for real debates in the social sciences.

The standard interpretation of counterfactuals.

Most contemporary discussions of counterfactual theories of causation begin with Lewis, who observes that we think causation has something to do with making a difference. This difference is best thought of as what would have happened if a particular cause had not operated. For example, a rock hitting a window caused it to break. We relate the rock hitting the window with the window breaking as cause and effect, and we do so, at least partly, because we believe that had the rock not hit the window the window would not have broken. A rock hitting the window makes a difference to a world in which no rock hit the window. (Lewis, 1973, pg. 557). Lewis aims to show how this type of analysis works to distinguish causes and effects. At the outset, he limits his analysis to causation among events, such as "flashes, battles, conversations, impacts, strolls, deaths" (1973, pg. 558) etc. These are exactly the types of things that social scientists are often interested in.

Lewis's two analyses of counterfactuals

Lewis (1979) provides two analyses of counterfactuals. Most of the discussion of Lewis focusses on his second analysis, which is the more general treatment. However, I think it is worth reviewing his first analysis because he believes that it works for a large number of cases. Specifically, these types of cases are very like the types of cases that historians, and many other social scientists, consider.

Analysis 1

"These types of counterfactuals take the form 'if it were that A, then it would be that C', where A is entirely about affairs in a stretch of time t_a. Consider all those possible worlds w such that:

1 A is true at w.

2 w is exactly like our actual world at all times before a transition period beginning shortly before t_a.

3 w conforms to the actual laws of nature at all times after t_a

4 during t_a and the preceding transition period, w differs no more from our actual world than it must to permit A to hold.

The counterfactual is true iff C holds at every such world w." (Lewis, 1979, pg. 462)

In Analysis 1 we take the actual past up until just before we want to consider a counterfactual occurrence and make a few changes in this transition period to enable the counterfactual to occur. In doing this we avoid "gratuitous difference from the actual present" (Lewis, 1979, pg. 463). Then we make the counterfactual change and let the situation evolve according to the actual laws of nature. Lewis also notes that there may be a "variety of ways" in which the transition period can be filled out, so there may be no true counterfactuals that "say in any detail how the immediate past would be if the present were different" (Lewis, 1979, pg. 463). Analysis 1 fits a wide range of cases, but it does require details of a particular time.

Analysis 2

Analysis 2 is preferable because of its greater generality, but "…needs to agree with Analysis 1 over the wide range of cases for which Analysis 1 succeeds" (Lewis, 1979, pg. 464). The explicit references to particular time periods in Analysis 1 mean it can't be generalised to consider counterfactuals like 'had more countries been democracies in the twentieth century then there would have been fewer wars during that time'. This is a counterfactual that has been much discussed in political science. In this case there is no clear candidate for t_a. We might try to make t_a the whole of the twentieth century, however this poses problems for us in specifying the transition period. This is because creating a possible world in which more countries had been democracies requires multiple transition periods. We could try to reduce it to type 1 counterfactuals by saying 'had country x been a democracy at t there would have been fewer wars, and had country y been a democracy at t there would have been fewer wars, and had country z been a democracy…’. However, it isn’t clear that this makes it any easier to assess because there are still multiple interventions and multiple transition periods. Furthermore, the countries involved affect each other so it isn’t possible to assess each counterfactual in isolation. Had country x been a democracy at t, then this might affect the plausibility of country y being a democracy, as well as the likelihood of conflict between them, and other countries. In whatever way we choose to assess this counterfactual, it is problematic using Analysis 1. Analysis 2 makes this easier to assess because it replaces the ‘transition period’, with a miracle that changes the world in the ways we need without considering how this miracle came about.

Analysis 2 is based on comparative similarity of possible worlds. “A counterfactual ‘if it were that A, then it would be that C’ is (non-vacuously) true if and only if some (accessible) world where both A and C are true is more similar to our actual world, overall, than is any world where A is true but C is false” (Lewis, 1979, pg. 465).

A can be any supposition we want. Taking the example above, we get ‘If it were that more countries were democracies in the twentieth century, then it would be that there would be fewer wars’. There is no transition period and we just ‘assume’ more democracies, or, in other words assume some miracle has taken place that turns a number of non-democracies into democracies. This gets us around the problem of specifying t_a for a general counterfactual claim. This counterfactual is true only if all the possible worlds in which there are more democracies and fewer wars is more similar to our actual world than any world in which there are more democracies and not fewer wars.

‘Accessibility’ imposes two restrictions; nomological and historical. Nomologically accessible worlds are those which are subject to the same laws as our world. Historically accessible worlds are those which share the same history with our world up until a particular point in time. So, a “nomologically or historically accessible world is similar to our world in the laws it obeys, or in its history up to some time” (Lewis, 2009, pg. 8). This is important because it restricts the number of worlds that are considered ‘possible’. In the assessment of these possible worlds the question of ‘similarity’ arises. Lewis notes that the idea of overall similarity is vague and that this can be understood in different ways in different contexts. Therefore, this is just a skeleton analysis until it is supplemented with the appropriate notion of similarity in a particular context. I will discuss problems with judging similarity later in this chapter.

Historical counterfactuals and Analysis 1.

Analysis 1 only applies when we are dealing with particular time periods, but this is often the case with the counterfactuals that historians analyse. In addition, historians largely follow the same outline as Lewis's Analysis 1, but do not cite him as a source for their methodology. What, therefore, are the requirements for the successful analysis of a counterfactual using Analysis 1? Lewis says that "Analysis 1 is built for a special case. We need a supposition about a particular time, and we need a counterfactual taken under the standard resolution of vagueness." (Lewis, 1979, pg. 464). Lewis describes counterfactuals as "infected with vagueness, as everyone agrees" (1979, pg. 457) and believes that different ways of resolving the vagueness are appropriate in different contexts. Lewis's standard resolution of vagueness is based on the observation that we normally want to preserve the intuition that counterfactual dependence is asymmetric; that counterfactual changes affect the future but not the past. If true, then this suggests that backtracking is not permitted because when we backtrack, changes are made to the past because of a counterfactual change in the present. Analysis 1 preserves this asymmetry of causation because changes to the past are restricted to the transition period. Lewis says, "Analysis 1 guarantees the asymmetry of counterfactual dependence, with an exception for the immediate past […] We have the counterfactuals whereby the affairs of later rimes depend on those of earlier times" (1979, pg. 463). He describes Analysis 1 as an "Asymmetry-by-fiat strategy" (Lewis, 1979, pg. 464).

How we treat this transition period is critical for historians, and they impose strict conditions about what can happen in the transition period. Tetlock & Belkin (1996) give examples of counterfactuals that have been considered in historical and political debates.

> CF 1: 'If Yeltsin had followed Sachsian fiscal and monetary advice in early 1992, Russian inflation in 1993 would have been a small fraction of what it was.'

CF 2: 'If the US had not dropped two atomic bombs on two Japanese cities in August 1945, the Japanese would still have surrendered roughly when they did.'

CF 3: 'If all the states in the twentieth century had been democracies, there would have been fewer wars.' (Tetlock & Belkin, 1996, pg. 4)

The first two of these clearly refer to counterfactual occurrences at specific times—early 1992 and August 1945. We can therefore use Analysis 1 in these cases. However, we need Analysis 2 to assess his third example. Tetlock and Belkin call counterfactuals taking the form of his third example 'nomothetic' [i.e. law-giving] because they apply to well-defined theoretical or empirical generalisations, such as between democracy and war. However, the majority of the historical counterfactuals that are discussed in the literature are not general because they refer to specific hypothetical changes at specified times. I will therefore spend the rest of the chapter discussing the first kind of historical counterfactuals.

Tetlock and Belkin list six criteria that they believe encourage more rigorous thinking about counterfactuals. These are:

1 Clarity: the hypothesised antecedent and consequent must be clearly specified and unambiguous.

2 Logical Consistency: hypothetical events linking the antecedent and the consequent should be specified, and be consistent with each other and with the antecedent.

3 Minimal Rewrite Rule: antecedents should require altering as few 'well-established' or agreed upon historical facts as possible.

4 Theoretical consistency: connecting principles should be consistent with 'well-established' theoretical generalisations that are relevant to the hypothesised antecedent-consequent link.

5 Statistical Consistency: connecting principles should be consistent with 'well-established' statistical generalisations relevant to the antecedent-consequent link.

6 Projectability: testable implications of the connecting principles should be used to test whether the hypotheses are consistent with additional real-world observations. (Tetlock & Belkin, 1996, pg. 18).

The main criterion of interest here is the 'minimal re-write rule' because it addresses how historians should think about the transition period. It stipulates that antecedent conditions should alter as few 'well established' facts as possible. Setting aside Tetlock and Belkin's implied worry about what 'well established' facts are, there is ambiguity in how we are to understand the minimal re-write. Given a choice, are we better off changing one big fact, or lots of small ones? Which of these options is 'minimal'? Furthermore, how do we count the 'size' of a historical fact? For example, do we look at the consequences of a fact, or do we just focus on the point in time when a change is made? Historians are likely to disagree in the same way that philosophers are likely to disagree about similarity judgements. However, the distinction becomes clearer if we contrast a 'lesser re-write' with a 'greater re-write', as follows:

Consider the application of Tetlock & Belkin's methodology to counterfactual CF 2: 'Had the US not dropped two atomic bombs on two Japanese cities in August 1945, the Japanese would still have surrendered at roughly the same time'. We will hold actual history constant until just before the first atomic bomb was dropped on Japan. The way this counterfactual is phrased glosses over the fact that it is referring to two distinct events: There were two flights to Japan, which happened on different days. We have two options here: Either to treat them as one event, or to see them as two, closely related, events. I prefer the second approach because the absence of the first flight would, probably, have affected the likelihood of the second flight occurring. However, this complication does not really affect the analysis here, so I will refer to the counterfactual change regarding the 'first flight' and leave the second flight as part of the hypothetical, counterfactual, scenario. Depending on how we see the counterfactual playing out this flight may, or may not, happen.

There are various ways in which we can set up this counterfactual. Perhaps we consider the counterfactual situation in which a decision was made to cancel the first flight to Japan at the last moment. Or perhaps we consider the counterfactual situation in which the plane carrying the bomb did fly to Japan, but was called back at the last minute. Either way, we make as few changes to the historical record as possible (lesser re-writes). This is in contrast to the case where we consider a counterfactual in which there was a war between the US and Japan, but in world in which planes had not been invented (a greater re-write). This would require considerably more re-writing and should, according to Tetlock and Belkin, be avoided. They note that possible worlds should begin with the real world, as it was, before asserting a counterfactual turn of events.

They suggest not re-winding long stretches of history, because if the antecedent and the consequent are separated by large stretches of time, or a host of alternations to the historical record, it is difficult to see how they can be linked as cause and effect. For example, we could make changes to Japanese foreign policy going back to, say the 1890's, and then try to assess whether, in this counterfactual world, the US and Japan would have gone to war. However, in this case, the link between the purported cause- Japanese foreign policy, and the consequent- the Second World War, is convoluted and passes through several possible 'branchings', or events that could have been affected in a variety of ways by the earlier changes to the historical record. In such cases it is difficult to make a clear causal link between the first change in foreign policy and the final outcome of war. Additionally, there are many things we can take 'Japanese foreign policy' to mean so, using the analysis of Nomadic concepts presented in Chapter 2, we could go further and say that to suggest the 'Japanese foreign policy caused the Second World War' is not a precise claim at all. Tetlock and Belkin's minimal re-write criterion is therefore best seen as a stipulation of what not to do, rather than a way of discovering the scenario with the fewest re-writes amongst a selection of scenarios that do not make use of 'major re-writes'. This suggestion is clearly unsatisfactory as there will be some crossover point between lesser and greater re-writes; it is unclear how we can move beyond this. Tetlock and Belkin set out criteria for counterfactual analysis, but their stipulation that historians should only minimally re-write history in the transition period is difficult to implement in any precise way.

Ferguson (1997) imposes stricter limits on counterfactual analysis that constrain what can happen in the transition period even further. He suggests that counterfactuals in history should focus on alternative courses of events that were considered by actors at the time. He says, "We should consider as plausible or probable *only those alternatives which we can show on the basis of contemporary evidence that contemporaries actually considered*" (Ferguson, 1997, pg. 86. Italics in original). Later, he supplements this with the requirement that "we can only legitimately consider those hypothetical scenarios which contemporaries not only considered but also committed to paper, (or some other form of record)…" (Ferguson, 1997, pg. 87). This considerably restricts the possible worlds that are considered accessible because it essentially just involves considering a decision or course of action that someone could have executed in a number of ways. Assuming that the people in question have an accurate view of their circumstances, these are the alternative scenarios are seen as possible, given their, current, state of affairs. Ferguson's justification for this is two-fold. Firstly, historians need to construct "plausible alternative pasts", for him, and alternative past is only 'plausible' where it is based on "historical evidence" (Ferguson, 1997, pg. 87). In the main, historical evidence is written record. Secondly, he believes that historians should give greater weight to outcomes that contemporaries considered possible than to those present-day historians think possible. His constraints are an attempt to limit counterfactual analysis by using evidence and giving weight to the expectations of historical figures. I will discuss one of these types of counterfactuals later in this chapter.

So, at least some of the counterfactuals considered in history are of the type that is successfully analysed using Analysis 1. The 'transition period' in Lewis's analysis is the period historians alter in the light of their 'minimal re-write rule'. In other words, it is the timeframe in which changes need to be made to the historical record to make the counterfactual possible. Historians are at pains to make as few changes as possible during the 'rewrite'. How well does Analysis 1 do the job?

Let us pick up the US-Japan counterfactual again. Stating it in exactly the same way as Lewis does ("If it were that A, then it would be that C" where A is entirely about affairs in a stretch of time t_a) we get:

'If it were that the US had not dropped two atomic bombs on two Japanese cities in August 1945 then it would be that the Japanese would still have surrendered at roughly the same time.' The time period in this case is entirely about a few months in 1945.

1. A is true at w.

Suppose that w is a world such that 'The US did not drop atomic bombs on two Japanese cities' is true at w.

2. w is exactly like our actual world at all times before a transition period beginning shortly before t_a.

This condition holds if we postulate that the first plane was recalled at the last minute.

3. w conforms to the actual laws of nature at all times after t_a

There is no reason to suppose that this would not hold. If we speculate that the first plane was recalled, so no bomb was dropped. Consequently, the second plane also does not fly so it also does not drop a bomb. The alternative future of the US and Japan obey the laws of nature.

4. during t_a and the preceding transition period, w differs no more from our actual world than it must to permit A to hold.

Under the filling out of the scenario above this condition holds. We have rewritten as little as possible to enable the counterfactual to happen.

5. The counterfactual is true iff C holds at every such world w.

We may decide the counterfactual is true, or not, depending on how we judge the consequences of the antecedent of the counterfactual. The precise difficulties in doing this will be the subject of the second half of this chapter. For now, let me just note that, at first sight, this outline suggests that Analysis 1 is a good fit for the way in which historians think about some counterfactuals.

However, we have glossed over what Lewis calls the 'immediate past'. He says that the immediate past does not depend on the present in any 'definite way'. Lewis advocates leaving what happens in the transition period vague[5]. He says,

> "There may be a variety of ways the transition period might go, hence there may be no true counterfactuals that say in any detail how the immediate past would be if the present were different. I hope not, since if there were a definite and detailed dependence, it would be hard for me

[5] Lewis is not using the word 'vague' in the sense that it is used in the philosophy of language but, as Crasnow does, to indicate a lack of clarity, or uncertainty about what happens in the transition period. Where the word vague is used in this chapter it is used in the same way as Lewis does.

> to say why some of this dependence should not be interpreted- wrongly, of course- as backward causation over short intervals of time…" (Lewis, 1979, pg. 463)

In other words, the transition period may be fleshed out in a variety of ways. What Lewis seems to be saying here is just that, if we leave it unclear what happens in the transition period, with no specifics on how the transition period is filled out, then we avoid pointing to anything of which it could be said that it is being caused by our counterfactual change. This seems odd, for the very basic reason that although we may leave the transition period unclear, this does not mean that we don't have various alternative events in mind. Was the flight carrying the first bomb to Japan prevented from leaving US airspace due to technical failure, or was it cancelled, or was the plane recalled over the Pacific? If we opt not to choose, then, although the transition period is unclear, it still contains events which, potentially, could be seen as instances of backward causation. This is because we specify the counterfactual change at t, and then make changes because of this change at t-1, to enable the change at t. The remedy doesn't work because ignorance of what is happening in a transition period does not mean that we are not assuming, implicitly, that something specific did in fact happen. Furthermore, if we do not choose, we can't determine future consequences exactly. For example, if Truman changed his mind at the last minute and ordered the bomb not to be dropped, then the Japanese might have thought that Truman was losing his nerve and attempted to press on with the war more fiercely than before. Whereas, if the flight was cancelled due to bad weather, that isn't possible. When we turn to the critique advanced by Reiss, below, we will see that different ways of filling in the transition period do, in fact, make a difference to how we assess the truth value of a counterfactual.

Lewis's Analysis 1 fits many of the counterfactuals that historians consider, however, there are two clear problems illustrated by the discussion above. The first is judging closeness, or similarity, and the second is transition periods. The following section discusses each of these in turn and concludes that there is no way to resolve either.

Judging closeness of worlds

One clear question that arises even just with the definition of Analysis 2 is the difficulty in judging degrees of similarity between possible worlds, and between possible worlds and the actual world. This implies that we are able to compare possible worlds to our own and judge how much possible worlds differ from the actual world. Comparing worlds and judging their similarity to one another is an explicit component of Analysis 2, but is also a necessary, if implicit, component of Analysis 1. Point 4 above states that the transition period "differs no more from our actual world than it must to permit A to hold". We are asked to avoid making gratuitous changes to the actual world. In other words, the changes we make should keep the counterfactual world as close as to the actual world as possible.

How do we judge which changes during the transition period are closest to our actual world? We can think of cases where this is obvious. It is obvious when reconsidering the window example outlined on page 70 above. The counterfactual in this case is 'If a rock had been thrown, then the window would have broken'. Who do we think would have thrown the rock? A possible world in which I throw a rock at a particular window is closer to the actual world than a possible world in which a dinosaur throws a rock at a particular window. This is because we have to think of things being only slightly different to enable me to throw a rock at a window whereas we have to think of many things being different to enable a dinosaur to throw a rock at a window.

But this assumes that the number of changes we need to make determines similarity between worlds. If I throw a rock it seems that we just need to make me do something, while if a dinosaur is to throw the rock we need to make a number of changes to the history of life on Earth. This sounds intuitively true, but it depends entirely on how we describe the situation. At first sight we do only need to make one change to enable me to do the throwing—I just need to throw the rock. But I don't throw rocks. I have never thrown rocks and I struggle to think of anything that would make me actually throw a rock at a window. The seemingly small deviation from actuality turns out to be a very large change indeed. Perhaps we need to change aspects of my character. Perhaps we need to postulate some extreme event that would make me throw a rock. Civil unrest? A threat to my life? Sometimes, purportedly small deviations from the actual world may actually involve surprisingly large deviations. That said, we might still be confident that however many changes we need to make to my character, the possible world in which I throw a rock is closer to the actual world than one in which a dinosaur throws a rock, but I have suggested that such judgements depend partly on our description of a scenario. If we describe the scenario in terms of a person (me) who is capable of throwing rocks it appears that only a small change is required. On the other hand, if we describe the scenario including aspects of my character then many more changes will need to be made. The number of changes we need to make is not obvious. Furthermore, historians, and social scientists in general, are not often asked to assess scenarios involving as obvious a contrast as between a human and a dinosaur. A more realistic comparison is between me, or Person X, throwing a rock at a window. The discussion above is intended to illustrate that determining which scenario involves the fewest, or smallest, changes is difficult and may very well depend on our description. Person X might be a criminal, and therefore, at first sight, fewer changes will be needed to the actual world to enable Person X to throw the rock than would be required to enable me to do so. However, what if Person X, despite their violent criminality, hates breaking glass?

Lewis's use of miracles in Analysis 2 avoids this problem, because we simply assume that a counterfactual event happened, without describing a state of affairs that makes this reasonable. However, in the cases historians consider, the use of miracles isn't a viable strategy. The reasons for this will be discussed in more detail below. Lewis is aware of the difficulty in judging similarity but takes the idea of "comparative over-all similarity among possible worlds" as primitive (1973, pg. 559). He means by this that we are familiar with these sorts of relations of all over similarity between two possible worlds and the actual world. We make these sorts of judgements, according to Lewis, all of the time, such as in judging similarity between people (Lewis, 1973, pg. 559). Again, this is intuitive and obvious, when considering the breaking window example. We say that a world in which the window partially breaks (perhaps it cracks but the glass stays in place) is closer to the real world, in which the window is left alone, than a possible world in which the window breaks completely. However, how we extend this to slightly more complex situations is unclear. In the window case the counterfactual change relates just to the extent of a crack in a window pane. The larger the crack the further away from the actual world the possible world is. This is a measurable change. Judging similarity between people is not, obviously, measurable and involves comparisons between a large number of variables, which is considerably more difficult.

To illustrate this, take Lewis's example of judging similarity between people—am I more similar to my brother than my sister? My hair and eye colour are the same as my brother's but not my sister's, some aspects of my character are more similar to one of them, some with the other. With whom do I share more experiences? When many features (eye colour, character, hair colour) are involved and where differences in these features are not easily measured, judging similarity poses two problems. Firstly, which features are the salient ones and secondly, once we have decided which are salient, how do we weight them? Lewis is well aware of the 'vagueness' in the idea of similarity and says this will "not be entirely resolved. Nor should it be" (1973, pg. 560).

Lewis is right that we often do make similarity judgements between people. The fact that we can make any old similarity judgement doesn't help though. Lewis agrees, saying, "Do not assume that just any respect of similarity you can think of must enter into the balance of overall similarity with positive weight" (Lewis, 1979, pg. 465). His advice is that we should find the right sort of similarity judgement that combines with Analysis 2 to yield the proper truth condition of a particular counterfactual. But how do we decide which is the right similarity judgement? What if there is disagreement? I may decide that I am more similar to my sister, but my mother may judge that I am more similar to my brother. Which similarity judgement is right? It doesn't matter if we are just discussing similarity of people, but when we are trying to judge which transition period is closest to the actual world, or which possible world is closest to our own, we need our similarity judgement to be right. Even if we agree what the right metrics are for judging similarity, we may disagree about the aggregation of these. Indeed, Lewis explicitly rejects the possibility of quantifying similarity. This is because, in doing so, we are likely to assume that the degree of similarity of A to B is the same as the same as the degree of similarity of B to A. This would be unjustified if, in each world A and B, different things count more for the purposes of making similarity judgements. The example Lewis gives is colour, which may weight moderately highly in our world, but may weight very highly in another world (1973, pg. 50-52). In summary, Lewis provides little way of deciding between similarity judgements.

He does delve into the similarity problem by considering a historical example raised by his critics. 'If Nixon had pressed the button there would have been a nuclear holocaust'. We have a strong intuition that this counterfactual is right, and therefore want an account of counterfactuals to yield this conclusion. The problem with this counterfactual is that given any world in which the antecedent and consequent are both true it will be easy to imagine a closer world in which the antecedent is true and the consequent false. Lewis notes that there may be all sorts of worlds where Nixon presses the button. We must consider which of these differs least, under the appropriate similarity relation, from our world. He says that a tiny miracle (he describes this as a violation of the laws of nature) takes place that makes Nixon press the button.

The worlds following the pressing of the button by the insertion of a small miracle in the timeline should turn out to be more similar to our world, under the right similarity relation for this case, than any of the other worlds in which Nixon pressed the button. In other words, they should be more similar than other possible worlds which have a large amount of backtracking prior to the pressing of the button by the counterpart Nixon. Lewis says that "a lot of perfect match of particular fact is worth a little miracle" (Lewis, 1979, pg. 469). Similarly, worlds in which the counterpart Nixon pressed the button but which did not end up in holocaust also need to be less similar to the actual world than the world in which the counterpart Nixon presses the button and holocaust results. Initially, these worlds might appear to be very similar to the actual world; Nixon has pressed the button but nothing happens, perhaps because the button has faulty wiring. But Lewis thinks that this similarity could not last because some of the tiny differences, such as Nixon having had the intention to begin nuclear war, would lead to big differences as time goes by. Lewis says that many things would be different if Nixon pressed the button, but the wiring was faulty. Nixon's finger print would be on the button, Nixon would be agitated, the pressing of the button might be recorded, and Nixon might write his memoirs differently, at a later point. Lewis concludes that "I should think that the close similarity […] could not last" (1979, pg. 470).

There is a further possibility—that the counterpart Nixon presses the button but that a second miracle restores the world to the same state as the actual world. All the traces of the button pressing are removed through this miracle. This possible world also needs to turn out not to be more similar to the world in which Nixon presses the button and causes a holocaust. In summary, under the similarity relation we are after, the possible world in which there is a small miracle, Nixon presses the button and nuclear holocaust follows, must be closer than any of the other possible worlds. This includes those in which there is no holocaust, and those where other miracles occur.

This illustrates how determining the right judgement of similarity is critical. It also illustrates that the similarity judgement should be made with the aim of vindicating the conclusion we seek. Lewis says we are seeking "a similarity relation that combines with Analysis 2 to give the correct truth conditions for counterfactuals…" (1979, pg. 472). In other words, the truth conditional of the counterfactual that we seek, drives what we consider to be the appropriate similarity judgements. However, what if the purpose of the analysis is to determine the truth value of a counterfactual? Lewis compares the possible world in which Nixon presses the button and nuclear holocaust results, and the possible world in which he presses the button and no holocaust results, with the real world, and judges that the first is closest to the real world. His justification for this is that tiny differences, such as Nixon's intention to start war, would lead to large differences in the future. But how do we know this, really? It depends largely on our ability to think about differences in the possible paths of history stretching well into the future. But can we not also think about a world in which Nixon presses the button, a nuclear bomb does go off, but a few years down the line we think this world is actually pretty close to the actual world in 2016. Perhaps the bomb lands in a relatively uninhabited place so the effects of the nuclear explosion are limited, perhaps the bomb only partially exploded, perhaps we suppose that no other country would have retaliated. We can think about either possible world in a variety of ways and just like judging closeness or similarity between people, it isn't clear which features are relevant, or how we are to weight them. Further, different people may do this differently, and without a pre-existing commitment to a particular truth value of the counterfactual there is little way to choose between them.

In conclusion, judging similarity between possible worlds is difficult, especially when we have no pre-existing commitment to a particular truth value of a counterfactual. Difficulties with judging similarity complicate Analysis 1 because it is difficult to judge how different to the actual world alternative transition periods are. Tetlock and Belkin suggest re-writing the transition period 'minimally', but while it is unclear how we are to judge similarity between alternative transition periods and the real world it is unclear how 'minimal' various alternative transition periods are. Furthermore, in the cases that historians often consider, Analysis 2 is not a good model because they do need to consider what would have enabled a counterfactual change to happen. This is discussed in the following section.

Backtracking

Backtracking also poses a problem for counterfactual analysis. Lewis notes that the way the future is depends counterfactually on the way the present is. If I make changes in the present, I can affect the way the future turns out. However, if I make changes in the present I won't usually affect the past. Counterfactual dependence therefore usually works in one direction, but not the other. Lewis then describes a case where we do think that if the present were different the past would have to have been different to enable the present to be as it is. This is a backtracking argument, and in such cases we do think that if the present were different the past would be different too.

His example is from Peter Downing:

> "Jim and Jack quarrelled yesterday, and Jack is still hopping mad. We conclude that if Jim asked Jack for help today, Jack would not help him. But wait: Jim is a prideful fellow. He never would ask for help after such a quarrel; if Jim were to ask Jack for help today, there would have to have been no quarrel yesterday. In that case Jack would be his usual generous self. So if Jim asked Jack for help today, Jack would help him after all" (Lewis, 1979, pg. 456)

In other words: If Jim had asked Jack for help today, there would have had to have been no quarrel. Backtracking counterfactuals are characterised by a particular way of speaking. They are usually characterised by phrases of the form "If it were that… then it would have to be that…" (Lewis, 1979, pg. 458). As in 'if Jim had asked Jack for help today it would have to be that there had been no quarrel'. They arise because we think that in order to make the antecedent of the counterfactual possible, changes need to be made to the past to enable the change to be made. We saw this above with the US Japan example. If the US is not to drop nuclear bombs on Japan, then we need to make changes to the historical record to bring this about, even if these changes are sometimes very small. Backtracking counterfactuals of this sort are common in historical analysis.

Why do we need to do this? Why don't historians just assume some sort of miracle that stops the bombs being dropped? Because it isn't clear how we can assess the truth of a counterfactual without this backtracking. Suppose that our counterfactual is 'Had

Margaret Thatcher not been Prime Minister of the UK then the unions would be a much stronger political force today'. We can take Margaret Thatcher out of the situation, perhaps we assume she was never born, but what now? Can we think about the UK in a political context without a Prime Minister? We need to put someone else in her place. Can we assess the strength of the unions without any idea of who might have been in charge during the 1980s? In order to assess the truth of this counterfactual we need, at the very least, some sort of coherent picture of the UK political situation and this is not what is left just by miraculously subtracting Margaret Thatcher from it. So, we need to backtrack to fill in the alternative scenario before we can assess the truth of the counterfactual. Despite the seeming necessity for these types of counterfactuals Lewis says that he has analysed them "only so that I can ask you to ignore [them]" (1979, pg. 458). As I will argue below, ignoring them is difficult in many social science situations. The real worry with backtracking is that how we fill out the transition period affects how we assess the consequent. Reiss (2009) is one of a number of authors who have pointed this out. When backtracking is needed, and it often is, then how we fill out the transition period affects the assessment of the consequent.

Reiss begins by noting that historians' analysis of counterfactuals bears some, superficial, resemblance to Lewis's scheme. However, he says that in Lewis's analysis an antecedent is implemented by a miracle "a minimal incision that breaks all causal laws that have the antecedent event as effect and brings about the event without itself having causal antecedents" (Reiss, 2009, pg. 718). This, as we have seen above, is only the case if we are using Analysis 2. He further notes that historians are concerned with the reasonableness of antecedent conditions including the rewrites that need to be made to enable the counterfactual change. In other words, the changes historians make when they backtrack should be 'reasonable'. We will see below some suggestions historians have for understanding 'reasonableness'.

He says that almost all the historical counterfactuals he has analysed do backtrack. He discusses a counterfactual relating to the Cuban Missile Crisis. In July 1962, Krushchev and Castro agreed to construct missile bases in Cuba, precipitating a crisis that brought the world to the brink of nuclear war. Krushchev was emboldened to pursue this strategy partly by Kennedy's lack of resolve during the US Bay of Pigs invasion and the Berlin Crisis, both in 1961. An interesting counterfactual to consider is: Had Kennedy shown greater resolve prior to the Cuban Missile Crisis, Krushchev would not have sent missiles to Cuba. Citing Lebow & Stein, he says it does not make sense to just consider an alternative world in which Kennedy issued a warning, thereby showing resolve, prior to the missile build-up in Cuba because we need to consider "what conditions in the antecedent's past would have had to be in place in order for the counterfactual antecedent to appear possible or likely" (Reiss, 2009, pg. 717). In other words, we need to consider what conditions would need to have been present in order for Kennedy to have been 'more resolved'.

Backtracking is a problem for counterfactual analysis because once we do it there is considerable flexibility in how we change the past to enable the difference in Kennedy's attitude. Lewis accepts that there are a variety of ways in which the transition period can be filled out. But, as Reiss notes, differences in the way we think about this counterfactual scenario make the consequent either true or false. Lewis suggests that the transition period can be vague, but we have seen why it is difficult to assess historical counterfactuals without making the transition period specific. Reiss suggests that how, exactly, we decide to do this can make the counterfactual either true or false.

For example, if Kennedy issued a warning because he received advanced information of Soviet plans we might conclude that Krushchev would not have altered his plans because he saw this as a one-off, given Kennedy's lack of resolve in 1961. Alternatively, if we fill in more detail about how knowledge of the plans was obtained we might conclude that the realisation that the US had that level of access to Soviet plans would have led Krushchev to change his strategy. In other words, changes in the way we set up the counterfactual situation lead to changes in the conclusions we draw from the counterfactual analysis. Reiss notes that Lebow & Stein see this as a reason to reject the viability of this counterfactual.

The same argument goes for our Japanese example. If our counterfactual scenario is that the US did not drop the first atomic bomb on Japan because of a technical failure with the flight, or the bomb, then we may conclude that Japan would still have surrendered at the same time because, we might think, they thought that if they didn't surrender they would be bombed. The mere threat of atomic bombs, in this scenario, is sufficient to cause the Japanese surrender. If our counterfactual is instead that the US government changed its mind and cancelled the flight, we might conclude that the Japanese surrender would have occurred much later, if at all. Our reasoning in this case is that the Japanese were given no reason to surrender, and may even have seen the cancelling of the flight as a sign of US weakness. As with Reiss' example, the specifics of how we fill in the transition period affect the conclusions we draw from the counterfactual.

The backtracking problem is therefore not that historical counterfactuals backtrack *per se* because we can use Analysis 1, rather than the miracle as required by Analysis 2, and ensure that our transition period makes as few modifications to reality as possible. In theory, using Analysis 1 we can fill in the counterfactual history using the guidelines from Tetlock and Belkin, or Ferguson. The problem with backtracking is that this counterfactual history can be filled in in a variety of ways. I argue that Lewis is not correct in thinking that we can leave this vague because, depending on how we decide to specify this alternative scenario, we may alter the results of the analysis. This would not be a problem if we could decide which alternative transition period is closer to the actual world, or, which is 'more likely' (to use the historians' terminology) but as the discussion above has shown, there seems little way to decide this.

Woodward (2003) accepts these problems with transition periods and similarity judgements and believes that his idea of an 'intervention' overcomes these problems. This extends Lewis's analysis in important ways. Dealing firstly with transition periods. Woodward agrees that transition periods are problematic, and thinks that they are problematic for the reason argued for above. He says, "…once transition periods are countenanced at all, there may be a large number of possible transitions, none obviously closer to the actual world than any other." (Woodward, 2003, pg. 144). Secondly, as far as similarity judgements are concerned, he begins by noting that the criteria Lewis provides for judging similarity are relatively imprecise. He says:

> "Lewis is explicit that his standards for similarity are adopted simply because they make those counterfactuals come out true that we think, pretheoretically, ought to be true. The standards themselves are complex and not particularly intuitive" (2003, pg. 137).

For example, Lewis suggests that small miracles are preferred to large miracles, but Woodward notes that it is unclear how one should count miracles, or what counts as a small versus a large miracle.

Woodward proposes the idea of an intervention which overcomes these worries. His central thesis is that the concept of causation should be understood in terms of our ability to change things, or to make a difference. He claims that we can think about the differences we make through our manipulations in terms of counterfactuals—we assess the effects of our manipulations by considering counterfactuals like 'Had I not pulled the lever, the weight would not have changed position'. If this is true, then we have highlighted a cause of the weight's changing position. A manipulation breaks the causal chains leading up to the point at which an intervention is made, which means that there is no need for a transition period. We simply pull the lever. It is no longer relevant to ask how it came to be possible that we pulled the lever, we simply do so. This sounds almost like a miracle, and, in fact, Woodward believes that, in the cases where miracles are successfully invoked, these miracles work "like the notion of an intervention" (2003, pg. 135). The puzzling notion of 'similarity' is unnecessary because Woodward asks us to consider what would happen if we made a counterfactual change; similarity between possible worlds is not a consideration. Woodward applies his analysis to social science cases, and in the following section I argue that the use of counterfactual analysis in the social sciences is problematic because the problems presented by backtracking apply even in cases where we do not need to backtrack. Social scientists often make use of Nomadic concepts which enable the construction of multiple, plausible, and often contradictory counterfactual scenarios which we have no principled way of deciding between.

Counterfactuals and social science

Chapter 2 provided a framework for understanding how many concepts used in the social sciences are Nomadic. Nomadic concepts allow for a wide variety of social phenomena to be included within their scope. If the counterfactuals we consider use concepts that allow for the inclusion of a wide range of social phenomena, then there are many ways in which we can see the counterfactual playing out. These alternative counterfactual scenarios are often equally plausible, but contradictory, therefore making it impossible to assign a truth value to these counterfactuals.

Woodward (2003) explicitly rejects the suggestion that subjectivity is involved in assessing the consequences of counterfactual scenarios. Why is this? He says that in judging which counterfactuals to consider we need something like an idea of 'serious possibility' (Woodward, 2003, pg. 89). For example, we want to reject saying things like "the failure of a large meteor to strike me as I write these words will count as a cause of me writing them, and so on." Woodward counters the objection that the phrase 'serious possibility' is "unclear and subjective" by saying that, although it is true that an investigator's interests and purposes influence which possibilities are taken seriously, at least some of the considerations that go into such decisions are based on facts about the world and "seem perfectly objective" (2003, pg. 89). He says:

> "Once we fix which possibilities are serious […] there is no further sense in which the counterfactuals about the outcomes of the hypothetical experiments associated with typical causal claims are in some way dependent on human attitudes or beliefs." (2003, pg. 118)

In other words, subjectivity may be involved in judging which possibilities are 'serious', but, nevertheless, many of the judgements about what a 'serious possibility' is, are objective. Once we have decided which possibilities are 'serious', then there is no further subjectivity involved in assessing a counterfactual. We can think about the process of judging the consequences of a counterfactual change as a hypothetical experiment. He then (2003, pg. 119) asks us to consider the hypothetical experiment in which he steps in front of a speeding bus and assess whether he will be injured. Whether he will be injured does not depend on his beliefs and desires, as it is facts about the world that determine what happens to him.

As far as the social sciences are concerned, the claim that assessing the consequences of a counterfactual change is like a hypothetical experiment is largely mistaken. This mistake results from a failure to appreciate the differences between the concepts used in Woodward's example and the concepts used in many social science examples. In assessing the consequences of 'stepping in front of a bus' we are restricted to solely considering physical events and attributes. Woodward steps in front of a speeding bus. We have access to data from accidents and crash test dummies that enable us to identify a range of impact forces, and a range of damage he may suffer, given the speed of the bus, his weight, previous medical history, and so on. While we cannot be certain what injuries Woodward will suffer—whether he breaks his spine or suffers brain damage, we can be reasonably sure about the probabilities of various levels of injury he will suffer as a consequence of the collision.

By contrast, in the social sciences, we have issues about meanings to disentangle, meanings which may be unrelated, with unclear boundaries, and which change over time. This allows greater subjectivity in how we construct the counterfactual scenario under consideration. In the following, I argue that, for this reason, there are two related problems for counterfactual analysis in the social sciences. Firstly, that, despite Woodward's remarks to the contrary, it really is difficult to judge what counts as a 'serious possibility'. Secondly, that even when we fix what a 'serious possibility' is in a particular context, these 'serious possibilities' often involve Nomadic concepts which make difficult to judge the consequences of a counterfactual change. Given this, it is hard to see how we can find a principled way of choosing between divergent counterfactual scenarios identified as "serious possibilities". There are a variety of ways for people to be 'less ambivalent', or 'more resolved', and many different political set-ups can count as a 'democracy'. These ways need not be vastly different, but they can be different enough to significantly affect our assessment of a counterfactual.

What does a counterfactual claim mean?

I will use an example from King, Keohane and Verba (1994) to illustrate these two problems. They consider the effect of being an incumbent vs not being an incumbent Democratic candidate for the US House of Representatives, calculated in terms of the proportion of votes received. The difference between the votes received by the Democratic candidate as incumbent and non-incumbent gives us the proportion of the vote that is attributable to 'incumbency' status. They discuss re-running a specific election in New York in 1998, but I am interested in the more general process underlying their use of counterfactual analysis.

The reason for picking this example is that it appears, on first reading, to be a clear and precise counterfactual. We are asked only to consider the effect that switching one feature of a candidate (being the incumbent candidate) had on the number of votes they received. However, although it is clear that being the 'incumbent candidate' just means the candidate is the existing holder of the office, the concept of 'voting for' the incumbent candidate is Nomadic. The following paragraph will show why this is the case, and the consequences this has for assessing this counterfactual.

The counterfactual under consideration is 'what proportion of votes would the incumbent candidate have received, were they not the incumbent candidate?'. In assessing this counterfactual, we are asked to consider the effect on voter's behaviour of switching the candidate's incumbency status. In other words, would the proportion of voters voting for the candidate be different? The language used here is 'voting for', but it could just as easily be 'preferring', or 'supporting'. The analysis below applies equally well to any of these terms because they are all Nomadic concepts. Firstly, there are a number of things that 'voting for' the incumbent candidate can mean. In Woodward's language, there are a number of 'serious possibilities'. Taylor suggests that, as previously discussed, in voting for a proposal, a person may be saving the honour of their party or defending the value of free speech or any number of other things (1971/1994, pg. 190). This applies to the present case. Prior to beginning the analysis, we need to be clear that we understand what is really going on this situation. Voting for the incumbent candidate could reflect a simple preference for a particular candidate who has performed well, wanting to keep an alternative candidate out of office, a preference for the status quo, a lack of interest in politics, a fear of recrimination, or habit.

One response is that this is primarily definitional. Clearly, we can say that 'voting for an incumbent' can mean a variety of things but, once we have decided what the right interpretation is, we can move ahead with our analysis. This is Woodward's view. As discussed previously, in the discrimination case Woodward accepts that causal claims in the social sciences can mean many things and that therefore it is unclear which the relevant counterfactual is. As a brief reminder, the claim that "Being female causes one to be discriminated against in hiring and/ or salary" (2003, pg. 115) can mean a number of different things. As it stands, Woodward describes the statement as "unclear". However, when claims such as this are unclear we can either try to make the claim clearer, or, simply consider a number of interventions.

In the voting case, we may even be able to settle what we mean empirically by surveying the population to establish why they voted for the incumbent candidate, for example. This response is reasonable, but suggests that in many situations counterfactual analysis, prior to any empirical research that establishes the meaning of an action, is unenlightening. In any event, this is only a partial response because empirical evidence may not be available, or at least not easily attainable. Differences in how we understand what a 'more resolved' Kennedy means might lead to differing conclusions regarding the

Cuban Missile Crisis but it is unclear how we could clear up what 'more resolved' means empirically. Furthermore, as outlined in Chapter 2, specifying a meaning of a Nomadic concept in order to make it more precise only works when we do so using other Nomadic concepts. Nevertheless, it is difficult to find fault with the suggestion that we simply consider all the possibly relevant counterfactuals. Despite this, the following section shows that, even when we have fixed what we mean (or which counterfactual we need to evaluate a causal claim), or which scenarios are serious possibilities, there are still many ways of filling out these counterfactual scenarios.

Filling in counterfactual scenarios

Continuing with our analysis of the King *et al* example, suppose that research has determined that 'voting for the incumbent candidate' really does just mean that voters don't care about elections, or politics in general, and vote for whoever is currently in office. The reasoning behind this might be that voters don't pay that much attention to political issues and assume that whoever is in office must be doing an acceptable job. We can describe this as a lack of political engagement. The original counterfactual was 'would the proportion of voters voting for the candidate be different?'. But, after empirical investigation, social scientists have determined that 'voting for the incumbent candidate' doesn't mean that voters are giving the candidate their active and well considered support, but that it really means voters just aren't politically engaged.

This means that the answer to the question 'Would the proportion of voters voting for the candidate be different if they were not the incumbent candidate?' Is 'Yes, it would have been very much smaller'. However, given that it is lack of political engagement that is driving this voting behaviour, this doesn't tell us very much. There is no sense in which switching the incumbency status differentiates between differing levels of support for the candidate. It is simply that whoever is the incumbent gets the vote. Further analysis would yield a fuller picture of what is going on. To do this, the relevant counterfactual is 'Would the proportion of voters voting for the candidate be different, if voters were more politically engaged?'. We began with a Nomadic concept 'voting for', a particular candidate, which can mean many things, and, in this instance we have settled that it means lack of 'political engagement'. However, this is also a Nomadic concept. The following paragraph illustrates this.

Let us suppose that our hypothetical intervention in the case of voters who aren't politically engaged is requiring them to attend a short course on politics. Following this, how we see our voters behaving depends on what we take 'political engagement' to mean. Do they actively participate in political debates, do they perform thorough research on the main candidates, or do they care about some types of issues more than others? Entertaining this counterfactual state of the world is, primarily, an exercise in judging how a population that cares more about elections, and politics, might behave. Could we find out empirically what this population might look like? Empirical research might provide us with a range of possible forms of political engagement, but it is also likely to confirm the great variety of forms that political engagement can take. So, in this case, it isn't clear that we can be certain what our alternative state of the world looks like. Social scientists can include a wide variety of social phenomena in a counterfactual that uses the concept 'political engagement'. This allows us to draw different conclusions from the analysis.

To simplify this, consider the following two alternative states of the world. In the first scenario, voters are more engaged with the political process because (following their short Introduction to Politics course) they read all the material that candidates post to them, they watch the news, they have more nuanced political views and assess their local candidates' suitability on the basis of their positioning on these wider issues. In the second scenario, although they are aware of these wider issues, voters care much more about local issues, perhaps whether a hospital will be closed, or whether new houses will be built. Consequently, they care much less about how their candidate will vote on issues such as defence or foreign policy. Both these scenarios are plausible, and are, furthermore, plausible fillings-out of the notion of greater political engagement following our imagined intervention. In order to assess the implication of these counterfactual situations on votes for an incumbent let us suppose that we check his attendance record, voting record, history of supporting local issues, and so on, and find that he has a history of fighting local issues but often misses key votes on national and international issues.

In this case he is clearly likely to lose votes under the first understanding of greater political engagement, but this is less clear under the second understanding. The way in which we construct the counterfactual situation leads us to different conclusions and there seems little reason to endorse one conclusion over another, or to see one as more likely than another. Furthermore, it is difficult to specify what we need to do to choose between the two, given that both are equally plausible.

What, though, is the worry with this choice? Recall that, for Woodward, counterfactual analysis supports causal claims. If a counterfactual is multiply realised, as in the way envisioned above, can we still use it to support causal claims, such as that greater political engagement causes a population to vote for an incumbent candidate? Yes, we can, but only if each of the realisations points to the same conclusion, and there is no reason to think that this will be the case. As we saw above, although we might look for help from the idea of closeness or similarity to the real world this isn't much help because, given their equal plausibility, there isn't anything that can tell us which is further from our own world. And insisting that a 'fully specified' description of the situation really will enable us to judge which is the closest possible world is a reply that only a philosopher could love.

The assumption so far has been that we have two, easily distinguishable counterfactual scenarios that we can compare. But this is highly unrealistic. A more realistic cashing out is that social scientists would incorporate many ways in which a population might be more politically engaged. In this case, assessing the counterfactual scenario will involve balancing various conflicting consequences, which two social scientists are likely to do in different ways. They could conclude that the incumbent would get more votes, the same number of votes, or fewer votes. In this case, it will also be difficult to specify exactly how they have come to their differing conclusions. Without neatly distinguishable scenarios it is hard to see how we can begin to judge their closeness to each other, let alone their closeness to our world. This is very much like the problem of judging similarity between people. We can have information on as many dimensions of similarity as we like but we have nothing with which to decide between two opposing judgements about similarity.

The problem is this: A social scientist might think about a counterfactual situation in which voters are more politically engaged. They may think about this in a rigorous way, considering exactly how the population will change, and the effects this will have on the proportion of votes the incumbent will receive. However, 'political engagement' is a Nomadic concept. There are many things it can mean, and a great deal of social phenomena that can reasonably be included in the analysis of a 'politically engaged' voting population. In constructing a detailed, and plausible, alternative state of the world different social scientists are likely to construct a variety of possible, and equally plausible, alternative worlds using the same information. These alternative plausible worlds exist because of the latitude that the concept ‘political engagement’ gives us in including social phenomena. Where there is evidence, social scientists are likely to disagree on what evidence is relevant, and how it is to be weighted. If these alternative counterfactuals point to different outcomes we have plausible but contradictory conclusions. In this case it is difficult to see what we can conclude from the counterfactual.

A minimal social science example

Does this really apply to all social science examples? Can we think of one that is closer to Woodward’s bus example in the sense that no conflicting ways of assessing the consequences of the counterfactual exist? Some have thought so. Tucker (2009) provides an interesting example of a counterfactual for which we are almost certain of the consequent. “Had George Bush Sr died in 1990, Vice-President Dan Quayle would have become the 42nd president of the United States” (2009, pg. 230). Let us suppose that George Bush Sr stepped in front of a bus while taking a stroll outside the White House, and died. This counterfactual appears relatively unproblematic because Quayle was Vice President, the rules of the constitution stipulate the Vice President should take over if the President is unable to perform the functions of his job, and there is evidence that the political elites were acting in accordance with the constitution at the time. Tucker’s example is very different from our voting case because the outcome is determined by prior existing constitutional rules, which are intended to be clear and unambiguous. In this case, the concepts are precisely specified by constitutional rules. Furthermore, these rules are explicitly intended to determine the outcome of this type of event, because uncertainty at such a critical time could lead to political unrest. This is what makes it purportedly similar to Woodward’s bus example. Both counterfactuals ask us to inspect regularities—either regularities relating to the effects of impact injuries on humans, or regularities of behaviour when it follows the rules of the US constitution. In this case Woodward is right; once we establish that the counterfactual is a serious possibility we can run our thought experiment and see, with a high degree of confidence, what happens.

However, is Tucker’s counterfactual of any interest in the form presented here? Probably not, because it gives us no new information over and above inspection of the rules of the US constitution. As it stands, it just illustrates the rules. Furthermore, if we alter the example only slightly, the possibility of conflicting counterfactual scenarios reemerges. Perhaps Bush survived the impact but sustained slight brain damage. The medical verdict might be that Bush 'should' function normally but that there is a chance that evidence of cognitive impairment will present itself over the coming months. How do we assess whether Quayle will become the 42nd US President now? Section 4 of the 25th Amendment states that "Whenever the Vice President and a majority of either the principal officers of the executive departments or of such other body as Congress [...] submit their written declaration that the President is unable to discharge the powers and duties of his office, the Vice President shall [become] Acting President" (see www.constitutionus.com) Does slight brain damage with possible cognitive impairment count as inability to discharge the powers and duties of his office? There are a variety of ways we can see this counterfactual playing out, based on what we think other politicians and civil servants think the constitution really means in this context. Depending on how we decide this, Dan Quayle does, or does not, become the 42nd President (or acting president).

The relevant concept in this case is the 'ability to discharge the powers and duties of the office'. There are a number of things that this can mean. We could consider it to reflect a basic capability to fulfil the requirements of the job, or we could take it to mean the ability to fully deliberate on the problems facing a country. In other words, depending on what we think this means, we are happy with different degrees of help with the different aspects of the President's job. These meanings are relatively closely related because they just reflect differences in degrees of 'ability'; they do not reflect fundamental disagreement about what 'ability to discharge.... means'. However, these meanings don't have precise boundaries because it will be difficult to judge, in any particular situation, the extent to which he did 'discharge the powers and duties of his office'. The thought process leading to actions is not transparent—so it will be difficult to be sure of the extent to which the President really did 'discharge the powers and duties of his office'. Change over time is also relevant because our judgement of the President's ability is unlikely to remain constant as the months pass. So, in assessing the consequences of the President's injury we are free to consider a wide range of possible phenomena. It is not just 'we' who consider these things. Assessing the counterfactual depends on considering how other significant political players assess Bush's 'ability to discharge the powers and duties of his office'. The ways in which we think about these give us different conclusions about whether he would be replaced. This contrasts with Tucker's example where Bush dies.

Death is, as far as this example is concerned, an all-or-nothing matter and we are much more constrained in the social phenomena we can consider.

Whereas the unproblematic counterfactual wasn't illuminating, this version of the counterfactual is potentially interesting because it asks us to consider the consequences for the US political system of George Bush Sr's mental impairment. We can debate whether this really is interesting, but considering the implications of a mentally impaired US President is much closer to the work social scientists do than the unproblematic, first version, of this counterfactual. In other words, what is of interest to social science in cases like this is not just inspection of the rules, but consideration of what happens when the application of the rules is, potentially, unclear.

To summarise, Woodward argues that assessing the consequences of a counterfactual does not involve subjectivity and can be understood as a hypothetical experiment, which his bus example illustrates. However, this example is very different to most social science cases. We saw one social science counterfactual that works in the same way as the bus example, but it is also uninformative. As soon as we enrich the scenario with concepts that are Nomadic, we reintroduce the problem that there are multiple possible paths the counterfactual scenario can take, about which there is likely to be little agreement. These different paths affect what we conclude from the analysis.

The following example illustrates the problems introduced above in more detail. In this example two historians disagree about the truth value of a counterfactual. I argue that, in this case, no backtracking is required, and it is clear what the counterfactual claim means. It is clear, therefore, what the causal claim we are interested in is. Furthermore, the intervention is unambiguous. In this example, as far as possible, all complicating issues have been removed. This example illustrates that even when we have fixed which possibilities are serious, it is not the case that there is "no further sense in which the counterfactuals about outcomes of hypothetical experiments associated with typical causal claims are in some way dependent on human attitudes or beliefs' (Woodward, 2003, pg. 118). The outcome of this hypothetical experiment is impossible to judge.

Real historical counterfactuals

Historians, when they make use of counterfactual analysis, are at pains to demonstrate that the counterfactuals they consider rely on evidence. On the face of it, this focus on evidence should make counterfactual analysis less problematic. However, the following example illustrates that, even where considerable evidence exists, two historians can reach different conclusions from counterfactual analysis.

Adamson (1999) considers the counterfactual scenario in which Charles I of England did not decide to open negotiations with the Covenanters (Scottish rebels) in June 1639 but engaged them militarily. He argues that had Charles I engaged the Scottish militarily he would have won and the Civil War would not have happened. The two armies were ready for battle by June 4th & 5th in 1639. There is evidence that Charles considered both possibilities—military engagement and negotiation. Evidence also exists suggesting that, had a military conflict taken place, Charles would probably have won. For example, one of Charles' commanders overestimated the size of the opposing army, morale was failing in the Scottish forces, and they were struggling to feed and provision their army. In assessing the consequences of this counterfactual, Adamson aims to ground the construction of his alternative world in documentary evidence, which is clear as far as the short-term military consequences are concerned, but becomes less so the further into the 'alternative future' his scenario stretches. Still, the counterfactual scenario itself and the immediate consequences do seem well supported. Here, therefore, is Adamson's counterfactual:

Adamson's counterfactual: Had Charles I engaged the Covenanters (Scottish rebels) militarily in June 1639, rather than negotiating with them, he would have won, and the Civil War would not have happened.

This counterfactual is precise in a number of ways:

Meanings: It is clear what we mean. The two armies were prepared for battle, so a 'military engagement' simply means a physical conflict between these two armies. The Civil War is a real event, which is comprised of various events on various dates, and we are clear about the main actors involved in these events. The counterfactual claims that had a military conflict taken place, Charles would have won, and that the Civil War would not have happened.

Manipulation: The manipulation required is also relatively clear. Charles just needs to make a different decision; to fight rather than negotiate. This was an option he was already considering.

Backtracking: Backtracking is minimised, if not eliminated. Everything was in place for engagement or negotiation so the only manipulation required is a change in Charles I's decision and the historical evidence suggests that he considered both possibilities and that his decision could have gone either way.

Despite this, the two historians do disagree about this counterfactual:

Adamson believes that: Had Charles fought, he would have won, and there would have been no Civil War.

Tucker believes that: Had Charles fought, he would have won, but there would still have been a Civil War.

Tucker argues that Adamson underestimates the extent of the rifts in society which meant that "something similar to the Puritan Revolution" would still have happened (Tucker, 1999, pg. 235). Tucker and Adamson do agree that, had a military conflict taken place, Charles would probably have won. They agree about this based on the evidence about the relative strengths and capabilities of the armies. How we see a military conflict playing out is constrained by evidence of problems in the Scottish forces, amongst other things, which make it highly unlikely that Charles I would have lost. Furthermore, Jones analyses the history of warfare and argues that "Victors traditionally experienced less attrition (the casualties suffered) and the seriously disorganised vanquished more" (1987, pg. 639). This makes it sensible to assume that the larger army would probably have won. The role of evidence in this case brings us closer to the examples of Woodward and George Bush Sr colliding with buses. Just as there is a small chance that Woodward could walk away from the collision without injury, there is a small chance that Charles I would have lost, but data about collisions on the one hand, and evidence which we place in the context of previous military engagements make these possibilities unlikely.

What they disagree about is the consequences of this victory. As soon as we move away from this evidence, different interpretations of what the victory means lead to conflicting conclusions. There is an unhelpful ambiguity in Tucker's argument when he says that something 'similar' to the Civil War would still have happened. What counts as something 'similar' is difficult to specify—do we require full scale battles across the country, does it need to happen within a certain number of years, and so on. However, I think it is defensible to strengthen Tucker's position to mean that had Charles won the battles then *the* Civil War, the actual, real, historical event, would still have happened. This takes him to mean that had Charles won the battle it would have made very little difference, which I take to be in the spirit of his disagreement with Adamson.

They disagree about the consequences of the victory because Adamson believes that "looking beyond a hypothetical royal victory in 1639, the chances of Charles I being coerced by domestic rebellion or being forced to summon Parliament against his will would have been small" (1997, pg. 109). He points, amongst other things, to evidence of the declining influence of militant Catholicism, and the old age of many of Charles I's opponents. Tucker, on the other hand, argues that there were several independent causal chains leading to the revolution. A military victory would have made little difference to the occurrence of a revolution in England, it would just have been brought about in a different way. In other words, Tucker and Adamson can look at the same evidence and assess the consequences of the victory in very different ways. Even when we consider relatively simple counterfactuals, involving small manipulations that don't require backtracking, and where at least some of the consequences are supported by evidence, it is still possible to reach opposing, but plausible, conclusions which assign different truth values to the counterfactual.

To make this precise: This counterfactual asks us to consider the consequences for English history of Charles I winning a battle against the Scottish rebels on a particular day. Tucker and Adamson agree about the outcome of the battle, had Charles decided to fight, but disagree about the implications of this victory for English history. For Adamson, the victory would have meant that the opposition to Charles I had been defeated, while for Tucker, the opposition remained viable.

This is because in judging what winning the battle, and the subsequent military victory mean, the two historians can look at many aspects of the counterfactual social world. They consider such things as whether certain people would still have been motivated to rebel, or whether Charles would still have been coerced or forced to summon Parliament, or whether certain people would have been disinclined to demand more power for Parliament by the victory. 'Winning' is a Nomadic concept, because there are many things it can mean. Additionally, to the extent that the historians disagree about people's motivations and mental states, there seems little way to settle such disagreements empirically. For example, let us suppose that we disagree about whether a particular group of nobles would have been less willing to voice opposition to Charles, had he won the battle. Evidence often exists, about their economic circumstances, their parliamentary leanings, their fear of imprisonment, but the question of interest is whether a military victory for Charles would have led them to see these circumstances differently. This can be evaluated in any number of ways. Adamson's counterfactual about the consequences of Charles I winning a battle against the Scottish rebels benefits from clarity in what the claim means, it does not require backtracking, and the manipulation needed is clear and unproblematic. Despite this, the two historians disagree about the consequent. This does not seem at all like the hypothetical experiments that Woodward has in mind, because there is significant scope for in judging the consequences of a manipulation in different ways. Even when evidence exists, the historians disagree about what evidence is most relevant in constructing the counterfactual scenarios.

Why, then, do Adamson and Tucker agree about the outcome of the battle? Is this not similarly indeterminable? When they restrict themselves just to the question of military victory there is a rough rule, or pattern, that applies. The army with the most people is likely to win. This restricts the possible outcomes of the battle. However, although Adamson and Tucker agree about the likelihood of Charles's' victory it is arguable that this narrow, military meaning of 'winning' is more Nomadic that they suppose. For example, there are many ways of winning a battle and it might make a difference whether, for example, Charles's army slaughter every Scottish rebel they find and then go on a murderous rampage through the surrounding area. Or, alternatively, whether Charles wins narrowly and leaves much of the opposing army alive. Adamson and Tucker are not explicit about how they understand the military 'winning', but this may, in fact, mean various different things that might plausibly make a difference to the assessment of the counterfactual too. For example, had Charles's army gone on a murderous rampage we might think the Civil War would have been hastened. This suggests that there are even more ways in which we can assess the likelihood of the consequence of the counterfactual than the two presented by Adamson and Tucker.

In conclusion, this example is set up in ways that are similar to an experiment. The meaning of the antecedent and consequent of the counterfactual are clear. No backtracking is required. The manipulation is clear and unambiguous. Despite this, the two historians 'run' this hypothetical experiment and reach opposing conclusions because assessing the likelihood of the consequent involves a Nomadic concept 'winning'. These sorts of concepts allow social scientists to include a wide range of social phenomena in their scope. The two historians are therefore both able to construct plausible, but contradictory counterfactual scenarios.

Saving Woodward

It is possible to try to save Woodward's approach by questioning firstly, whether this counterfactual really does deal with a causal claim and, secondly, whether Adamson and Tucker are just assessing different causal claims, rather than disagreeing about the same causal claim, and thirdly, whether some of Woodward's later comments go some way towards accepting the criticism above. The response to the first worry is that, unhelpfully, Adamson and Tucker do not explicitly discuss causation, but instead discuss the contingency of historical events. Adamson's motivation for assessing this counterfactual is that people often assume that historical events were inevitable. He says, "the belief that Charles's experiment in government without Parliament was inherently unviable continues to enjoy currency" (1997, pg. 93). Adamson seeks to show that this was not the case, that "the critical moment was 1639.", and that Charles's failure to defeat the Scottish in 1639 "initiated the disastrous sequence of events which flowed from that failure." (1997, pg. 94). It seems reasonable therefore to take Adamson as saying that the failure to defeat the Scottish in 1639 was a, if not, the cause of the Civil War because he sees it as the initiator of a chain of events leading to the Civil War. Tucker's disagreement can also, therefore be taken to indicate a rejection of the causal claim that the events of 1639 did cause the Civil War.

The second worry with this historical example is therefore whether Tucker really is disagreeing about this causal claim, or assessing a different one. Are Adamson and Tucker just assessing different causal claims? It seems difficult to argue this, in this case. This is because they do not disagree about the set-up of the counterfactual. Tucker does not argue that Charles could not have fought, or that the manipulation needed is ambiguous, or about what the counterfactual means. Tucker accepts Adamson's set up, accepts that a military victory for Charles was probable, but disagrees about the consequences for the Civil War. This seems very much like two historians running the same hypothetical experiment and reaching different conclusions. To the question of whether A caused B, Adamson says 'yes' and Tucker says 'no'.

However, Tucker does not agree with Adamson about the causes of the Civil War. This does not mean though, that when assessing Adamson's counterfactual that they are assessing different causal claims. Tucker believes that, had Charles won the battle then this would have had little real impact because much of the actual historical record would have remained unchanged and that, therefore, other things would have led to the Civil War. Adamson believes that had Charles benefitted from a military victory, none of those other causes, nor anything else, would have let to Civil War. When there is disagreement about whether *A* caused *B* it seems reasonable to run an experiment where *A* is manipulated and to observe the resulting change (if any) in *B*. The problem with this historical example is that the historians can agree to the set-up of the experiment, but disagree about whether there is a resulting change in *B*. In a real experiment, whatever my initial views about what causes *B*, if I observe it changing in response to changes in *A* I will have to accept this. In the historical example, regardless of the strength of Adamson's argument, if I am committed to another interpretation of the Civil War, then this hypothetical experiment will not change my mind. I will always be able to argue that changing *A* really doesn't make any difference.

One further objection is that Tucker's and Adamson's counterfactuals aren't equally plausible; Tucker's is surely more plausible. This is a natural reading of the situation. However, although Tucker's counterfactual does seem more plausible this results from an asymmetry between Tucker and Adamson's alternative histories. This asymmetry makes it easier for us to think about Tucker's alternative historical timeline, but have nothing to do with the greater plausibility of his analysis.

In considering Adamson's views we need to imagine a world in which Charles wins a battle, then we have to imagine a world without a Civil War, while attempting to keep the rest of the historical record somewhat like the real historical record. In order to endorse Tucker's position we need to make the same counterfactual change to enable Charles to win the battle, but then we return to the actual historical record. It is easier to imagine Tucker's position because he is, essentially, saying that Charles winning the battle would have made little difference to the actual historical record. If Adamson is right, we need to think about a counterfactual world that differs greatly from the actual world. The

English Civil War was an important event in English history so considering what the history of England might have been like without it is difficult. It is easier to mentally construct a counterfactual scenario in which Tucker is right because we need to do very little mental work.

Finally, this discussion elaborates on the criticism in Chapter 1 of Woodward's (2016) attempt to outline some criteria for successful variable choice. In this paper, he appears to concede that some variables are better than others for counterfactual analysis, and hypothetical experiments, than others. He suggests seven criteria that should influence the choice of variables. The first two are important for the current discussion; these are:

> "Choose variables that are well-defined targets for (single) interventions in the sense that they describe quantities or properties for which there is a clear answer to the question of what would happen if they were to be manipulated or intervened on" (2016, pg. 1054)

> "choose variables that have unambiguous effect on other variables of interest under manipulation, rather than variables that have ambiguous or heterogeneous effects" (2016, pg. 1054).

If we take Woodward's advice seriously and only use variables for which there is a clear answer about what will happen if we manipulate them, this suggests that Adamson and Tucker's historical example is a poor example of a hypothetical experiment because the consequences of manipulating the battle are unclear. However, as argued in Chapter 1, in order for someone to select variables for which the results of manipulation are clear and unambiguous they must know that the results of the manipulation *really are* clear and ambiguous. The historical example illustrates that this is not often the case. Adamson believes that there is a clear answer to the question of what would happen if the decision to fight, or not fight, is manipulated, and that these effects are unambiguous. Tucker, on the other hand, believes that it is clear that Adamson's answer is not the right one. In other words, social scientists may think that Woodward's criteria are satisfied, but this doesn't guarantee that they are. Woodward is right to focus on the importance of choosing the right variables, but articulating this in terms of clear and unambiguous knowledge of the consequences is, I believe, not the right strategy. The preferable strategy is to understand exactly how variables differ, and how these differences make a difference to the viability of using counterfactual analysis. Chapter 2 has attempted to outline this.

Conclusion

This chapter has argued that counterfactual analysis is problematic in the social sciences. Problems arising from the need to judge similarity and backtracking have been well discussed. Woodward largely takes his analysis to have solved these problems, and believes that assessing the consequences of a hypothetical manipulation is like running a hypothetical experiment. However, this chapter has given reasons for thinking that the flexibility we have in filling in transition periods also applies to judging the consequences of manipulations. This problem arises from the use of Nomadic concepts with which social scientists often deal. As described in Chapter 2, when social scientists work with Nomadic concepts a wide range of social phenomena can be included within the scope of these concepts. When assessing the consequences of a hypothetical manipulation social scientists can construct multiple, but plausible, alternative counterfactual worlds which incorporate different social phenomena. Furthermore, there seems little principled way to choose between these counterfactual worlds.

Chapter 4: Intentions, Nomadic concepts, and predictive behaviour

The social sciences are usually concerned with understanding intentional behaviour. This is because a distinction is made between reflexive behaviour, and intentional actions. Purely physical responses, such as reflexes, are usually considered to be the subject of biology, or physiology, not the social sciences. It is no surprise that we might be able to formulate generalisations about the dilation of pupils, for example. But these are usually uninteresting from a social science perspective. However, taking an intentional perspective on behaviour often involves the use of Nomadic concepts. This chapter therefore begins with a discussion of the problem that Nomadic concepts pose for understanding intentional behaviour, before outlining the structure of the remainder of the chapter.

The standard analysis of agency and intentions stems from Davidson (1963/2002) and Anscombe (1957/2000), their view is that agency is just the performance of intentional actions. Further, that when people act intentionally, they act for a reason. A reason should be understood as a set of beliefs and desires that rationalise an action. For example, suppose that I go for a walk today. I usually assume that this is an intentional action, or, in other words, that I intended to go for a walk. In support of this, I can point to my beliefs, such as that walking is good exercise, and my desire to do some exercise today, and say that my belief and my desire caused me to go for a walk. It is important to note that this understanding of agency and intentional action is the 'standard' view and, although influential, is not universally accepted. However, there is no need to assess complications and concerns here (although later in this chapter some alternative analyses of intentional action will be reviewed) because the current claim is just that when social scientists take an intentional stance, this frequently requires the use of Nomadic concepts. Furthermore, social scientists frequently ask questions such as 'Why did those people do x?' and accept an answer of the form 'Because they wanted x, and believed that by doing y they would achieve x.'. Usually, such explanations are accepted without first requiring an analysis of agency, or intentional action, from a philosophical perspective.

To summarise, when trying to understand behaviour that we believe is intentional, we often seek to understand it in terms of beliefs and desires. However, once we elaborate on what these beliefs and desires are, we are likely to use Nomadic concepts. The problem is that these concepts allow for the inclusion of a great deal of social phenomena so social scientists are likely to find it difficult to agree about whether explanations given in terms of beliefs and desires are, in fact, explanatory. The following example illustrates this. Suppose that Bill spends his family's savings on a sports car. His wife, Mary, is less than impressed. One social scientist might explain this behaviour in terms of Bill's desire to enhance his status among his friends. His friends all own sports cars, so, they say, he believed that buying a sports car would improve his status. Another social scientist might explain his behaviour in terms of his desire to annoy Mary, because he suspects her of having an affair, and his belief that buying a sports car would annoy her. Bill's behaviour can fall under the scope of a number of Nomadic concepts, in this case, 'annoying' Mary, and 'enhancing status'. In other words, there are a number of possible intentional explanations of his behaviour.

We could ask Bill why he did it, and he might say that he did it purely to annoy Mary. However, the first social scientist might ask why he bought a car when he could have bought many other things that would also have annoyed Mary? Buying a sports car does enhance his status, so even if Bill doesn't give this as an explanation of his action, it really does explain it. Similarly, Mary might say that Bill couldn't have bought the car to annoy her because Bill believed her when she said she wasn't having an affair. Mary thinks Bill bought the car because he's depressed, and thought that buying it might make him feel better. Mary might also say that Bill is only saying he bought it to annoy her because Bill doesn't want anyone to know that he's depressed. The problem arises because we are explaining Bill's behaviour in terms of beliefs and desires, and these beliefs and desires are expressed using Nomadic concepts, such as 'depression', and 'enhancing status' and 'annoying'. A wide variety of behaviour falls under the scope of such concepts, which also do not have precise boundaries. This makes it difficult to determine whether Bill's behaviour is an example of 'depression' or 'enhancing status', or 'annoying'. Furthermore, because these concepts lack precise boundaries it isn't clear exactly which aspects of Bill's behaviour we should restrict ourselves to analysing. We are trying to explain why he bought a car, but, in doing so, we look at Bill's wider behaviour; for example, who his friends are, what happened with his wife, whether he was depressed. It is unclear how much, or how little of this behaviour is relevant to judging the correctness of the explanation of his buying the car. Moreover, it is not just unclear how much is relevant, but how we would go about deciding on relevance. For example, I might seek to explain Bill's behaviour in terms of childhood influences. It seems that we can bring an almost endless amount of Bill's past experience to bear on the question of why he used the family's savings to buy a sports car.

An objection to this is that some explanations might clearly be wrong. The example above is not intended to show that any explanation is as good as any other. For example, an explanation that Bill bought the car to impress his mistress would just be wrong if it turns out that Bill doesn't have a mistress. This example is intended to show that when we seek to explain behaviour in terms of beliefs and desires and express these beliefs and desires in terms of Nomadic concepts, there may be a number of plausible explanations that we are unable to choose between. This is because the concepts with which we are analysing behaviour are Nomadic.

Following the analysis of Nomadic concepts in Chapter 2, we can offer a number of competing, but plausible explanations of the same behaviour, and it is difficult to agree about whether a specific behaviour is an example of a generalisation about behaviour. For example, suppose that it is widely accepted that men of Bill's age go through a mid-life crisis. Despite our agreement that this generalisation is true, we can still disagree about whether Bill's buying the car is an instance of a man, of a certain age, going through a mid-life crisis. Despite it being one explanation, I may put more weight on Mary's view that Bill is depressed (but not about his age). Another person may simply believe Bill, and say that he bought the car purely out of a desire to annoy his wife.

One response to the worry that explanation in terms of beliefs and desires is problematic, is to suggest that, perhaps, if beliefs and desires are redescribed, explanations will be more precise. The remainder of this chapter takes this suggestion as its starting point. The first section outlines McIntyre's (1994) argument for thinking that the social sciences will only find regularities in behaviour if beliefs and desires are redescribed. McIntyre's concern is with regularities, but his argument is equally relevant to a discussion of explanations. Furthermore, McIntyre discusses regularities relating to individual behaviour, not regularities of the sorts discussed in previous chapters (such as whether inequality leads to revolution). I conclude that McIntyre's argument is only partially convincing. He is right to think that we should explain behaviour using concepts other than beliefs and desires, but he is wrong to think that beliefs and desires are redescribed.

McIntyre takes it as given that the behaviour he analyses is, in fact, intentional. The second section of this chapter takes issue with this, and demonstrates that the aspects of the behaviour he describes are intentional, in an ordinary sense. To show this, I discuss my own example of regular behaviour in order to discover what really is going on in cases such as these. Attali and Bar-Hillel (2003) describe 'middle bias'. Middle bias is the tendency for people to bias the middle option when presented with a finite set of linearly arranged options. The attraction of analysing middle bias is that this phenomenon appears in a variety of different contexts, such as when people are asked to 'pick a number between one and ten', when they are asked to make a choice from among a row of objects, or when they are answering multiple choice questions. Understanding middle bias is important because it appears to be an example of a stable behavioural regularity. If this is true, then this is a first step towards formulating regularities which describe this behaviour. This section of the chapter aims to discover firstly, why this behaviour is regular, and secondly; whether, by understanding this behaviour we can uncover reasons which allow us to identify why most social phenomena in the social sciences are not regular.

The third section of this chapter argues that middle bias is not intentional in the way intentional behaviour is usually understood. The middle bias phenomenon is fully characterised, and I show that this does not fit within the usual framework of intentional action. I describe the characteristics middle bias does have and I suggest that we call this type of behaviour 'derivatively intentional'. The final section of this chapter shows that derivatively intentional behaviour is describable without using Nomadic concepts, which makes it considerably easier to find regularities. The regularities exhibited by these types of behaviour are the closest we can come to achieving the aim of formulating reliable generalisations in the social sciences.

The aim of this chapter seems slightly at odds with the focus on *social* science concepts in previous chapters. This is because the discovery of a regularity describing individual behaviour does not necessarily imply anything about the discovery of regularities describing social behaviour. The justification for analysing middle bias is that it is a good example of a regularity concerning human behaviour, so it makes sense to understand it fully before considering what this example implies for the search for regularities in *social* behaviour. In other words, although middle bias is not quite what we are after, it is worth beginning with the best example we have.

McIntyre's argument

McIntyre's (1996) suggests that redescription of the subjects of social science research is central to finding regularities in the social sciences. He is not alone in thinking that this is the case. For example, Scriven's claim that if physicists sought to explain the trajectories of individual falling leaves they would not have succeeded in finding laws was introduced in Chapter 1. McIntyre gives a number of examples of regular behaviour. He writes that his intention is to "give the reader purchase on what I mean by redescription. I hope to clarify and illuminate the concept of redescription […] by providing examples that present *candidates* for social scientific laws" (McIntyre, 1996, 105). He is explicitly *not* suggesting that the analysis of these cases is complete. The examples he gives are the tendency for people to re-evaluate historical choices they made so as to convince themselves that they made the right choice, even in the face of evidence to the contrary; the ability to manipulate people's responses to questions by getting them to learn word pairs, position effects that influence the choice between identical products, and the effects of suggestions under hypnosis. What unites all these examples, according to McIntyre, is that the beliefs and desires of the subjects involved in these experiments are not explanatory because the subjects are unaware of the real reasons for their behaviour. McIntyre does not discuss these examples at length so I will focus on one of these, the word association case, purely to illustrate his argument. I will argue that his example, while successfully highlighting a regularity in behaviour, does not show that any redescription is taking place in the successful explanation of the phenomenon. However, his suggestion that the regularities suggested by this behaviour need to be expressed without reference to beliefs and desires, is correct.

McIntyre's central point is that some form of redescription is needed in the social sciences before they can find reliable generalisations (McIntyre, 1996, pg. 104). He argues that generalisations in the social sciences need not always make reference to things which we are used to talking about: Wars, revolutions, and democracies. One of his examples of a regularity is word associations. He describes this type of behaviour as "a good one for our purposes because it deals with intentional and freely chosen action that is nonetheless highly regular" (McIntyre, 1996, pg. 106). He describes a 1977 experiment by Nisbett and Wilson in which subjects were asked to memorise lists of word pairs "some of which had been previously found to elicit specific target responses with high frequency" (McIntyre, 1996, pg. 106). Subjects learned a word pair such as 'ocean, moon' and when asked to name a detergent said 'Tide' more frequently than control groups. The reasons subjects gave for their choice included things like familiarity with the brand, rather than the word list they had learned.

The significance of this type of regular behaviour is that explanation in terms of the beliefs and desires of the agent are not, obviously, explanatory because whatever is going on is not something that agents are conscious of. People reported that they had said "Tide" for all sorts of reasons, such as, to repeat, familiarity with the brand. McIntyre notes that explanation of such behaviour in experimental literature may not make much sense to the actor. Whatever is going on in cases like this is going on outside the agent's awareness. McIntyre concludes that this shows "a role for redescription and suggest the value of nomological models in the explanation of human behaviour" (1996, pg. 106). In other words, while we may have originally thought that the explanation for saying 'Tide' had to do with subjects' beliefs about the brand 'Tide', these beliefs do not in fact explain why subjects made the choice they did. The real explanation is that they were 'primed' to say this by learning certain word pairs. Other literature supports the idea that people's self-reports about why they are doing things are often unreliable. For example, one recent study suggests that people are often not aware of their reasons for buying certain goods (Gram, 2010). Gram's particular study was on family food shopping. She notes that parents know that children influence their parent's decisions about food, but neither parents nor children agree about the degree of influence. Gram concludes that observational studies are beneficial in determining what is going on in situations such as these because of the unreliability of self-reporting.

McIntyre's focus is on the importance of redescription in identifying regularities. He suggests that, part of the reason, we struggle to find laws in evolutionary biology and social science, in contrast with natural science, is that we often suppose that the natural sciences refer to objects or kinds in laws that match what we ordinarily think about. A simple example is that astronomy makes generalisations and predictions about planets and tides. Planets and tides are also the things we initially set about to enquire into. He says, "To a large extent the "natural kinds" that serve as the referents of the laws in these cases match the natural kinds we already use to think about the phenomena in question." (1996, pg. 96). This, McIntyre believes, has led social scientists to, wrongly, suppose that laws in the social sciences must refer to concepts with which we regularly deal. He also notes that this is an overly simplistic view of science because scientists often redescribe these terms, or, "redefine the phenomena in such a way that we come up with new referents altogether, and thus may arrive at a law-like connection that gives us new insight into what the phenomena most basically *are*." (1996, pg. 100. Italics in original). In other words, social scientists often begin the search for laws with common, everyday categories which were sometimes a fruitful place to begin in many scientific endeavours. However, as in many scientific cases, redescription is required to uncover what is really going on.

McIntyre believes that evolutionary biology illustrates what he means by 'redescription'. I will discuss this biological example purely because it is the one McIntyre uses, not because I have any particular views about evolutionary biology. The critique of this example will be applied to McIntyre's social science examples later in the chapter, but there would be a loss of clarity were I not to initially discuss his biological example because he believes it provides a model for how to deploy the strategy of 'redescription' in the social sciences. As will become clear over the following paragraphs, one of my objections is that it is unclear what McIntyre means by redescription. While noting that there is considerable disagreement among biologists about whether there are laws in evolutionary biology McIntyre takes Dollo's Law as the best example of a law in this field because it has "mountains of empirical confirmation, virtually no exceptions, is theory based, and seems to be attempting to be genuinely explanatory" (McIntyre, 1996, pg. 86). Dollo's Law states that

> "A structure (i.e. a complex part of an organism or the entire organism itself) can never undergo complete reversal, so that it perfectly attains a previous ancestral state. Secondary convergences can always be recognised morphologically by preservation of some trace of an intermediate stage" (McIntyre, 1996, pg. 85)

McIntyre reviews the dispute about the nomological status of this law. He cites Gould as arguing that the different possible interpretations of complexity empty Dollo's law of empirical content (1996, pg. 86-90). The worry is that the stipulation that a *complex* part of an organism cannot undergo complete reversal leaves the debate open to equivocation about which parts of organisms are 'complex'. If any reversal of a structure is found it can be rejected as a counterexample because it is not 'complex'. This is a particularly significant problem because Dollo did allow for the reversal of simple structures and reversal of functions. However, McIntyre notes that Gould does accept the principle of irreversibility because evolution is an historical process; we can't rewind it. McIntyre takes this as Gould's tacit acceptance of irreversibility as a law and reiterates his argument (reviewed in Chapter 1) that uniqueness is no barrier to the formulation of laws.

It is McIntyre's assertion that redescription is needed before laws can be formulated that is the central concern here, so controversies about whether Dollo's Law is a genuine law, or worries about McIntyre's criticisms of Gould, will be set aside. Even if we accept the nomological status of Dollo's Law it is unclear that its formulation results from redescription of any sort. The importance of redescription, for McIntyre, is that scientists may need to redescribe phenomena in unfamiliar terms before the phenomena begin to yield law like regularities. This is not a new idea, Weber suggests this possibility when he says "In short, (the sociologist) imposes a totally different sense on the legal concept, which he makes use of because of its precision and familiarity." (1978, pg. 17). The 'legal concepts' he has in mind include 'state', 'nation', and 'family'. This is a particularly problematic assertion for Dollo's Law. Dollo's Law refers to wholes, or parts of, organisms. These are precisely the things people have historically thought about when considering the natural world. He says "we do permit science to cut nature at different perceptual "joints" than we do in our everyday experience" (McIntyre, 1996, pg. 99). However, the scientific and everyday 'joints' of Dollo's organisms and parts of organisms match up, quite literally. In his defence, McIntyre also cites Williams (1980), who defends the existence of laws in evolutionary biology and argues that part of this process is defining a new vocabulary. However, McIntyre's example, that of Dollo's Law, does not demonstrate that redescription is required.

McIntyre's suggestion is therefore that we need to redescribe the subject matter of the social sciences in order to formulate reliable generalisations. Unfortunately, his social science example works no better than his biological one, although for different reasons. The word association example above is an attempt to demonstrate how redescription might work. He says this, and his other example, illustrate "what I mean by redescription" (McIntyre, 1996, pg. 105). What is being redescribed here though? McIntyre says that the social sciences are often taken to be concerned with intentional description "and at this level of interest, we are just going to be talking about things in terms of the 'beliefs' and 'desires' of the agent" (1996, pg. 105). His examples are dealing with intentional behaviour, but "only upon redescription of talk about 'beliefs' and 'desires' do we begin to be able to explain the regularities manifested in our social behaviour" (1996, pg. 105). Specifically, any attempts to formulate generalisations about what people say when asked to name a detergent that are based on their beliefs and desires about detergents are an unpromising place to begin. By contrast, generalisations may be more successful if we explain people's statements about detergents in terms of the effects of word association. In this specific case, that after learning certain word pairs, people are likely to pick a specific response. Admittedly, generalisations about word associations fall short of what is often hoped for in the social sciences, as such generalisations are little consolation for those hoping to generalise about wars and poverty, or those hoping to predict future states of society.

The word association example reveals behaviour that is highly regular. Furthermore, it is certainly true that the explicit beliefs and desires of agents do not play a part in explaining this regular behaviour. Indeed, agents are mistaken about the motivations for this behaviour. In the word association case the subjects 'invent' reasons for naming a particular product. What is less clear is how this illustrates 'redescription' of beliefs and desires.

McIntyre is not entirely explicit about exactly what redescription is supposed to mean. He says we need to "redescribe the beliefs and desires of the subjects" (1996, pg. 108). His example does not show that desires and beliefs have been redescribed because beliefs and desires simply do not appear as redescribed terms. However, elsewhere he notes the "explanatory value of *alternative levels* of description" (1996, 106). Chapter 1 noted a certain ambiguity about what 'levels' are, but the suggestion that the phenomenon itself is redescribed, or, perhaps more specifically, that the phenomenon is analysed using a different approach, or different conceptual framework, is right. Phenomena that we might once have explained in terms of beliefs and desires is now explained in terms of something else.

To summarise, McIntyre suggests that we are more likely to find successful (that is, empirically verifiable) generalisations if we focus on theoretical or other entities which may not be the ones we are most used to thinking about. He suggests that the social sciences are usually concerned with explanation in terms of beliefs and desires of agents, not because we have particularly good generalisations stated in terms of beliefs and desires which underwrite these explanations, but simply because we are used to thinking about beliefs and desires. Once we redescribe beliefs and desires in other terms, reliable generalisations may become easier to find. He then provides a number of examples, one of which is the regularities found with word associations. In each of these cases the behaviour is intentional, but cannot be explained in terms of beliefs and desires; thus, he would argue, we need a different conceptual framework to discuss these regularities. Davidson makes a similar point in "Actions, Reason and Causes" where he says that even though we believe that particular reasons cause certain actions and that these relationships are subject to an underlying law, this does not mean that the law is formulated in terms of reasons and actions. The law may describe neurological, or chemical, or physical facts. (Davidson, 1963/2002, pg. 17). McIntyre says that "if one wants a better explanation of the regularity manifested in their behaviour, one will have to redescribe the beliefs and desires of the subjects in terms they may not understand or have access to" (1996, pg. 108). As with the biological example, I do not think his examples demonstrates that redescription of beliefs and desires is required, but he is right that the behaviour needs to be analysed without reference to beliefs and desires.

McIntyre takes it as given that the behaviour he describes is, in fact, intentional. He says, when describing various experiments showing word associations, that "We are still dealing with freely chosen and intentional behaviour […] if one wants a better explanation of the regularity manifested in their behaviour, one will have to redescribe the beliefs and desires of subjects in terms they may not understand or have access to" (1996, pg. 108). However, later he writes "we are often interested in going beyond explanations that are given only in terms of "Intentions"" (1996, 108) and that "the intentional perspective on human events is certainly important, but it does not alone provide a complete picture of human action" (1996, 109). He provides no justification for the view that this behaviour is intentional, and his later comments suggests that understanding of intentions may need supplementing. In the word association example I take this to mean that the intentional behaviour of the subjects needs to be combined with an understanding of the biases they show, which is not explained by referencing subjects' beliefs and desires. This is important because if this word association behaviour is not intentional, being like a reflex, or a biological response, it is not obviously the kind of behaviour in which social scientists are interested. Also, McIntyre ties himself to one notion of intentional action because he only discusses beliefs and desires. There are other proposals for understanding what makes actions intentional. As the discussion below will illustrate, it is possible that what makes action intentional is not causation by beliefs and desires in the right way, but the role actions play in plans, or the right kind of bodily awareness.

The purpose of the following section is therefore to explore in more detail how intentions are relevant for understanding the word association behaviour. I firstly review the main proposals for understanding intentional action. This portion of the chapter is largely descriptive and my aim is not to endorse any particular view. This review is nonetheless necessary because in order to understand whether the regular behaviour that McIntyre describes is intentional, it is necessary to show how this behaviour can be analysed from a variety of approaches to intentional action. I will argue that McIntyre's regular behaviour is not intentional in any ordinary sense. Demonstrating this solely by adopting the desire-belief approach to intentional action would leave open the possibility that the behaviour is unproblematically intentional from the perspective of a different account of intentional action. It is therefore necessary to take as broad an approach as possible. With this overview in hand I will conclude that an important distinction needs to made between the behaviour of subjects in experiments and the specific part of this behaviour that is regular and predictable. This distinction is critical to understanding what is happening in these cases.

What are intentional actions?

Discussion of intentional action often begins with the intuitive distinction that there is a difference between things we do and things that happen to us. Anscombe asks, "Why is giving a start or gasp not an 'action', while sending for a taxi, or crossing the road, is one?" (2000, pg. 10). A further example is distinguishing between raising my leg because someone has hit me on the knee, thereby activating a reflex, and raising my leg to kick someone. Kicking is described as an intentional action while my leg rising because of the activation of a reflex is not. We can try to make the distinction in terms of the mental states that an agent is subject to in the two cases. When I raise my leg to kick someone we can talk about my intention to kick someone, my plan to get revenge by kicking them, and my ability to control my behaviour to ensure that I kick them where it hurts. These psychological descriptions are not available to us when describing the activation of a reflex. If a reflex is activated, intuitively, there is no mental conceptualization taking place of which the agent is aware.

There are three main approaches to the analysis of intentional action.

- The desire-belief account, which seeks to identify intentional action through the desires and beliefs which caused the action (Davidson 1963, 1978; Garcia 1990; Pitcher 1970; Davis 1984).

- The planning account, which focuses on the role actions play in plans that agents formulate. Both these approaches are 'causal' because they emphasise the causal history of an action (Bratman 1979, 1999).

- In contrast, non-causal accounts of intentional action see the causal history of an action as, largely, irrelevant (Frankfurt 1978; Grunbaum 2007, 2010; Castaneda 1982, 1992).

The following discussion of the analyses of intentional action may appear lengthy, given the narrowness of the question addressed, namely, whether McIntyre's word association example is an example of intentional behaviour. However, the arguments presented in this section are also necessary for the analysis of middle bias, so, to avoid repetition they are included here.

Desire-belief account

This is the account that McIntyre has in mind. Davidson writes that "if someone acts with an intention, he must have attitudes and beliefs from which, had he been aware of them and had the time, he *could* have reasoned that his action was desirable (or had some other positive attribute)" (1978/ 2002, pg. 85. Italics in original). In other words, for Davidson, when people act with intention they have attitudes and beliefs about their action that make it desirable, or reasonable, or prudent, or some other positive reason for performing the action. One other key aspect of Davidson's philosophy that will be relevant for the following discussion is the importance of how an action is described. Davidson says, "Action does require that what the agent does is intentional under some description, and this is turn requires, I think, that what the agent does is known to him under some description" (1971/2002, pg. 50). This suggests that it is not simply action that is intentional, but action *under some description*. Description is important because suppose that I intend to make a cup of tea because I am thirsty. In pouring the water out of the kettle I fill my cup, but also spill water on my book. It would be right, according to Davidson, to say that I intentionally filled my cup with water, but it would be wrong to say that I intentionally spilt water on my book. Filling my cup and spilling water on my book are both descriptions of my pouring water from the kettle, but it is only when my action is described as 'filling my cup' that it is intentional because, in this case, I have positive attitudes towards filing my cup with water that make it intentional. I have no positive attitudes towards spilling water on my book.

What matters for Davidson therefore is the causal history of an action, or, in other words whether beliefs and desires have caused an action 'in the right way'. In my example, that my desire for tea, and my belief that pouring hot water from a kettle into my cup will help me to make some tea. The reason for saying that beliefs and desires cause the action in 'the right way', and not simply that they cause them, is because deviant causal chains are a problem for this account of intentions. Davidson gives the example of making stew. I am making stew and believe that adding sage will improve the taste. I hold the spice jar over the stew, intending to put some in, but then my cat knocks my hand, thereby causing sage to fall into the stew. In this case, what I intended did in fact happen, but it did not happen because of my intentions. Davidson summarises the problem as follows, "an agent might have attitudes and beliefs that would rationalize an action, and they might cause him to perform it, and yet because of some anomaly in the causal chain, the action would not be intentional in the expected sense, or perhaps in any sense." (1971/ 2002, pg. 87). This is why these types of causal chains are called deviant and they make it difficult to decide whether an action is intentional or not. The beliefs and desires seem not to have caused the result 'in the right way'. Such examples have motivated the inclusion of criteria stipulating that an agent does something intentionally only if events follow the agent's plan sufficiently closely. Spelling out what 'sufficiently closely' means is going to be difficult. A further stipulation that has been suggested is that the outcome must not depend too much on luck, whether this luck has been foreseen or not (see Kieran, 2015 for a summary of these approaches).

Garcia (1990) sidesteps the debate about whether actions caused by intentions 'in the wrong way' are intentional by arguing that actions can be 'partially intentional'. He defends the principle that for a person to do something intentionally is for them to do it *in the way* that they intended. 'In the way' they intended means along the causal path "along which an intention comes to produce the events of the types which figure in its propositional content" (Garcia, 1990, pg. 192). It is important, for him, that I kill my friend by shooting him, not by running him over. When we do achieve our ends in a different way Garcia says that we have more or less achieved what we intended to achieve and acted more or less in the way we intended. Garcia's analysis of these cases is that I *to some extent* killed my friend in the way I intended, and that my killing him was therefore *partially* intentional. He says this proposal is more appealing than saying either that I did or did not kill him intentionally (1990, pg. 192).

Garcia distinguishes between the concept of intending, and acting with intention, (also known as acting intentionally). For him, the concept of intending is more basic. Garcia's example illustrates this distinction. Bernice intends to shoot the king and fires her gun at him, but she misses and hits the branch of a tree, which falls on him and kills him (Garcia, 1990, pg. 192). Bernice intentionally fires her gun, but unintentionally hits the branch. Garcia argues that cases like this are difficult for Davidson, for whom Bernice firing her gun and shooting the branch are "one and the same action" (Garcia, 1990, pg. 193). This is a consequence of a Davidson's non-standard concept of an event. For him, events are "a fundamental ontological category", analogous to material objects (Davidson, 1969, pg. 307). This view of Davidson's is of no further consequence for the following discussion. If they are one and the same action then both are intentional, or not. For Garcia, when we say 'Bernice's firing of her gun was intentional' this is to say that 'Bernice's firing her gun was done in the way she intended to fire her gun'. This is true. In contrast, when we say 'Bernice's shooting the branch was intentional' this is to say that 'Her shooting the branch was done in the way she intended to shoot it'. She never had such an intention, so this is false. Her shooting the branch was not intentional. For Garcia, therefore, we can have actions, such as shooting the branch, that are not intentional. To summarise, an agent has intentions, and consequently acts with intention. However, because not all of these actions follow the path the agent intended not all of these actions are done intentionally.

Although appealing, Garcia's account isn't entirely satisfactory. If Garcia is right, and I 'more or less' killed my friend in the way I intended, then we need a way of deciding to what extent my plan is more or less the plan I intended. Although Garcia suggests that it is a matter of degree, he provides no criteria for judging the extent of the degree of difference. This is a problem because it makes it difficult to judge whether what happened really is, more or less, what I intended.

A further problem for deciding what intentional actions are is how we deal with side effects. Pitcher (1970) discusses an example first proposed by Chisholm, viz. the example of intending to kill the king. We intend to kill the king, but in doing so, also step on an ant. Does a person do both intentionally? Pitcher notes that part of the confusion surrounding such debates is that it seems wrong to say that the agent did them unintentionally, because, in most examples, the agent knows they are doing them. This observation holds particularly well in Bratman's example. Bratman comes home in the dark and considers whether to turn on the light. Doing so will enable him to see but will also wake Susan. Were he to turn on the light, causing Susan to wake up angry and ask "Why did you wake me?" responding "I didn't do it intentionally" isn't satisfactory. Bratman did intentionally turn the light on, after all. Pitcher's way out of this is to say that we can deny that a person performed an action unintentionally, while also not accepting that they did it intentionally. We can do something neither intentionally nor unintentionally. This differs from Garcia's suggestion because, using Garcia's framework, we can say that turning on the light was an intentional action because Bratman turned on the light in the way he intended to turn on the light, but waking Susan was not an intentional action because Bratman had no intention to do so. In contrast, Pitcher just opens up a gap between intentional and unintentional action. Some action may be neither, but this does not mean that there are degrees of intention.

Pitcher also notes that, although examples in the philosophical literature often direct our attention towards cases where the 'side effects' are large. "The standard—and we might as well say, the paradigm—cases of expected side effects, as the name itself implies, are relatively trivial things that we know will inevitably accompany what we intend to do." (Pitcher, 1970, pg. 667). The importance of this observation is that any concerns over labelling actions neither intentional nor unintentional are alleviated if they are insignificant. However, this creates a degree of unease when we consider significant, particularly morally significant, side effects. Are we really happy to say that I did not intentionally kill a child who I ran over while drunk driving? Davis (1984) bites the bullet in cases like this, arguing that unwanted, but expected consequences are not intended. He discusses playing the piano, which he intends to do, but which will also annoy the neighbours. He notes that saying 'I intend to annoy the neighbours' sounds wrong. What he intends to do is practise the piano; he does not intend to disturb anyone. He says "no one ever expresses an intention to bring about expected but unwanted consequences" (1984, pg. 47). He gets to the heart of the concern about morally significant side effects by suggesting that part of the reason for believing that expected consequences are intended is that we are responsible for them. However, responsibility and intention to not line up. He may be responsible for disturbing the neighbours, but this is not the same thing as intending to disturb them. We can therefore accept that I am responsible for the death of the child, and can be punished, but this does not mean that I did it intentionally.

Davis (1984) also has a slightly different understanding of the causal account of intention by emphasising motivation. His view is that we only intend to do something if we are motivated to do it. He defines it this way: "S intends that p iff S believes that p because he desires that p and believes his desire will motivate him to act in such a way that p" (1984, pg. 51). He believes that including motivation is important because having the right beliefs and desires are not sufficient, in and of themselves, for making an action intentional. He discusses Maura, who has a track record of having wanted the winning World Series team to have won. In other words, the team she wanted to win, did, in fact win. She knows this and, in the current series, wants the Red Sox to win. This isn't a case of wishful thinking, given her track record, but neither is it a case of intending, despite the fact that she desires that they win, and believes, quite reasonably (from her point of view), that they will win. Davis argues that the reason there is no intention here is that Maura knows that her desires will not bring about the result. If the Red Sox do win it will not be because Maura wanted them to. Motivation is what links beliefs and desires to other desires and actions. It is motivation, on this account, that is missing in the wayward causal chains. I did not expect my desire to motivate me to hit the branch. I therefore did not intend to hit the branch. Davis' example is problematic though because, given Maura's highly improbable track record, it isn't clear that Maura is not at least partially justified in believing that her wanting the Red Sox to win will lead to their victory. She may be at a loss to point to the causal process linking her desire to a team's success. But it isn't clear to me that she isn't justified in believing that there might possibly be such a process. In other words, it isn't clear when motivation is justified.

The philosophers discussed so far provide different ways of characterising intentional action. For Davidson, description is central to distinguishing intentional action from bodily movement. However, for Garcia and Pitcher actions may be unintentional or intentional. Garcia argues that actions may have degrees of intention, based on how closely they match what agents intend to achieve. Pitcher argues that we can do something neither intentionally nor unintentionally. We are now in a position to assess whether the word association behaviour is intentional.

It seems clear that it can be intentional under the desire belief account. The subjects' behaviour can be described as 'learning word pairs', and 'answering questions', or 'participating in an experiment'. The subjects would presumably recognise these descriptions of their behaviour. However, things are not as simple as this. A lot hinges on how we understand the 'behaviour' of the subjects. Their participation in an experiment is intentional. By contrast, it isn't clear that the subjects answered questions in the way they intended to. In Garcia's terms, it is clear that the subjects had a number of intentions, and therefore were acting with intention, but their intentional actions did not follow the path they intended that they should. They intended to name a detergent, but the name they said (surprisingly often) had nothing to do with their intentions. They also had no intention of being manipulated into giving a particular response. It isn't clear, therefore, that the regularities in their responses is intentional, even though the more general behaviour (participating in an experiment) is intentional. An obvious suggestion is that the regular behaviour is a side effect. But this doesn't seem obviously true either. The subjects did intend to give specific answers to the questions they were asked, and they did do this. A side effect, as described by Pitcher, would be something like stepping on the scientist's foot, or dropping something. In our case, the regular behaviour doesn't appear to be a side effect because it is inseparable from the intentional behaviour of question answering. Judging whether this regular behaviour is intentional is therefore more complicated than McIntyre suggests. At this point it is necessary only to note this, and to realise that there is a distinction between the behaviour in which subjects are engaged and the specific aspects of this behaviour that are regular. This distinction will be discussed in more detail when we turn to the example of middle bias because I do not think that it can be thoroughly resolved using McIntyre's example. Before doing this it is necessary to assess whether McIntyre's example is intentional using the other main analyses of intentional action. Recall, this is because, if my characterisation of this behaviour is right, then it should be right regardless of which approach to intentional action we take. If the word association behaviour is unproblematically intentional from the perspective of another account of intentional action, then McIntyre is vindicated, and he has highlighted an example of intentional behaviour that is regular and predictable.

Planning account of intentional action

Bratman (1999) rejects the desire-belief model of intention, focussing instead on the importance of planning in guiding action. We form many plans which guide present and future action. Intentions are the "building blocks of such plans; and plans are intentions writ large" (1999, pg. 8). Bratman's central reason for believing this is that under the belief-desire model of intention we ignore plans. Future directed intentions are not embedded in our present actions. In the same way, if we just focus on intentions we miss the fact that present action is often part of a prior plan. On this view, seeing my shooting of my friend just as an expression of my present desire to kill him, and beliefs about how I can achieve this, misses the historic plans of which this action is a part. For example, my desire to kill him may form part of a larger plan to make him change his will so I inherit his house when he dies; killing him is the culmination of this plan. Furthermore, Bratman notes that plans are often "formed, retained, combined, constrained by other plans, filled in, modified, reconsidered, and so on" (1999, pg. 8). These, he believes, are essential to understanding intention. Just focussing on beliefs and desires misses this. In other words, intentions need to be embedded in plans: Long term, short term, changing and interacting plans. Concentrating just on the beliefs and desires relevant to any particular action does not do justice to the complexity of what is going on when we act.

Bratman begins with the observation that our common sense understanding of intention allows us to understand intentions as states of mind while also allowing us to characterise actions as done intentionally (1999, pg. 3). Bratman says that the "main tradition in recent philosophy of mind and action" (1999, pg. 5) has been to give intention in action methodological priority. In other words, the focus has been on analysing the intentions behind action. This leads naturally, according to Bratman, to the view that intentional action is to be understood in terms of an agent's desires and beliefs, and the relations between these desires and beliefs and the action. On this view, future directed intentions (intentions about actions in the future) can be reduced to desires and beliefs. For example, my intention to get to London by train tomorrow is just for me to desire to get to London tomorrow, and believe that a train will get me there. There is no separate state of intending to act; desires and beliefs are sufficient for understanding intentional action.

Bratman explicates his notion of intention as a distinctive attitude, which is not reducible to desires and beliefs, by distinguishing between various aspects of intention. He writes that future directed intention involves commitment, which has two aspects. The first, volitional, aspect concerns the relation between intention and action. Intentions, in this sense, control action, whereas desires do not. Bratman says that intentions control action, whereas desires are just potential influencers of action. For example, I may have a desire to buy blue shoes this afternoon. However, I also weigh this desire against my other desire to finish this chapter. My desires potentially influence what I do. However, if I intend to buy blue shoes this afternoon I will simply execute this intention. The second aspect of commitment is the role it plays in making plans for the future. Bratman calls this the reasoning centred aspect of commitment. My intention to buy blue shoes this afternoon will influence the formation of further intentions. These two aspects, working together, "help explain how intentions play their characteristic role in supporting coordination, both intrapersonal and social" (Bratman, 1999, pg. 17). These two aspects of commitment help to explain why, when someone says that they 'intend to x' we generally believe that they will do x. This is a brief sketch of Bratman's reasons for believing that planning is central to understanding intentional action. However, it is sufficient for analysing McIntyre's example.

Is the word association McIntyre describes intentional according to the planning account of intentional action? It is reasonable to suppose so, although the plans involved in this behaviour are not ones we know about, given his description of the situation. For example, the experiment may have formed part of a variety of plans; to help out with a scientific research project, to fill an otherwise empty afternoon, or to answer the questions quickly so as to not miss the bus home. It is therefore possible to see the subjects' behaviour as influenced by plans. However, as with the desire-belief account, these plans do not seem, obviously, to relate to the behaviour in which McIntyre is interested; the frequency with which responses to questions can be influenced by prior learning of word pairs. This is because the subjects are unaware of the manipulation of their responses. Again therefore, although plans relate to the wider behaviour in which subjects are engaged, the specific aspects of this behaviour that are regular do not feature in any plans. There is one remaining account of intentional action that we can turn to.

Non-causal accounts of intentional action

In contrast with both Bratman, and the desire-belief model, Frankfurt (1978) argues that the causal history of an action is unimportant in understanding intentional action. What matters is the extent to which the action is under an agent's control. Importantly, he is not rejecting the idea that actions have causes, just that "it is no part of the nature of an action to have a prior causal history of any kind" (Frankfurt, 1978, pg. 157). In other words, prior causes are not the right way to distinguish between what people do and what just happen to them. If causal history is primarily what matters, then intentional and mere bodily movements do not differ in how they are experienced by the agent once they are happening. This is clear from the discussion above. Observing the physical event of my leg rising does not, necessarily, tell us whether this event is an intentional action, or not. The proposal above is that we should look to the prior history. However, if this is primarily what matters, this suggests that the leg moving should be experienced in the same way in either case, but it is not. Our leg feels very different when it has been activated by a reflex than when we raise it ourselves. Frankfurt concludes from this that a person exhibiting this behaviour only knows whether it is intentional or not by introspection at the time the behaviour is occurring. In deciding whether we are behaving intentionally we examine what is happening, not the prior causal history of the behaviour. The decisive issue, for Frankfurt, is a person's "connection with the movements of his body" (1978, pg. 158).

Later, he highlights the importance of whether the physical behaviour is "under the person's guidance" (Frankfurt, 1978, pg. 158). He concedes that when something happening under our guidance it is "under the guidance of an independent causal mechanism, whose readiness to bring about compensatory adjustments tends to ensure the behaviour is accomplished" (1978, 160) but he resists the suggestion that this amounts to a causal theory of action because these causal mechanisms co-occur with the bodily movements. Looking at things that occurred before the bodily movement began tell us nothing about how the action feels, or whether it is under our guidance. He reserves the description 'intentional' for behaviour that is purposive and under the guidance of an agent. His account therefore takes issue with the assertion that the physical act of raising my leg to kick someone and it rising because a reflex has been activated are physically indistinguishable. For Frankfurt they are primarily distinguishable because they feel different for the agent, not because of the different causal history. This difference in feeling, he believes, is critical to understanding intentional action. The critical difference between an action and mere behaviour, or kicking and a reflex, is that the former is under the agent's control or guidance. Intentional action is not only under the agent's guidance, but is also purposive; it is undertaken deliberately or self-consciously.

In Frankfurt's account, a person may also engage in physical behaviour that is not an action, but which becomes an action, once he brings it under his control. He describes a car coasting down a hill under the force of gravity. The driver may be happy with the course and speed the car is taking and may therefore do nothing to affect the motion of the car. This does not demonstrate that the car is not under the driver's guidance. This is because the driver is ready to intervene if he believes he needs to.

This view has support from Grunbaum (2008), who defends the idea that an agent's experience of moving "has an epistemic place in the agent's awareness of her own intentional action" (2008, pg. 243). He begins with the observation that the idea that our experience of moving our bodies is "involved in our awareness of what we are intentionally doing when we act intentionally" (2008, pg. 244) is intuitively appealing. For example, when I raise my leg to kick someone I feel that I am raising my leg. Were I not to feel this I would probably stop the action believing that there is something going wrong. Despite this intuitive appeal, Grunbaum notes that most philosophers deny that bodily awareness plays a substantive role in an agent's awareness of whether their behaviour is intentional. This is because it is often descriptions of behaviour which underpin notions of intentional action. We describe things people do, such a painting a wall, in ways that make the behaviour intentional. Such descriptions include 'pleasing my wife', 'decorating my house', 'painting the wall' etc. Grunbaum says "It would be at this level that the agent could be said to be aware of her own action, and at this level that she understands and is able to enter into the exchange of asking for and giving reasons for action […] Awareness of how one is moving is simply at the wrong level with respect to awareness of one's own action as intentional under some description." (Grumbaum, 2008, pg. 246). In other words, I am intentionally painting my wall because I have an awareness of myself 'painting the wall', or 'getting through my chores'. My belief that I am acting intentionally has little to do with the feeling I have of lowering and raising my arm, or raising my feet to go up a ladder.

Grunbaum's proposal is that when we intentionally do something, our awareness of doing it intentionally is partly grounded on our awareness that the behaviour is under our control. Were I to engage in some behaviour over which I have no control, then I would have no awareness of it being intentional. He illustrates this with the example of 'anarchic hand syndrome', whereby people's hands perform apparently goal directed movements, such as grabbing things. (See Della, Marchetti & Spinnler, 1994) This may look like an intentional action, but it is not, because the person with the syndrome has no control over their hand; they cannot stop or alter the movement. It is the sense of control that makes sense of the difference between anarchic hand syndrome and normal action. Grunbaum is not arguing that this is all that matters to intentional action, but that bodily awareness of performing an action is normally an important part of intentional action.

Grunbaum also argues that our experience of active movement comes in degrees. There is a continuum going from "premeditated movements and passing by habitual actions, fast movement reactions to our surroundings, and conscious blinking of the eyes to truly passive movements such as reflexes." (2008, pg. 249). At least at the level of bodily perception, actions can feel more or less intentional. The movement of my leg can be ranked in terms of awareness of control. When I am kicking, I have the most awareness of control, when my reflex is activated I have less, but not as little as when someone is raising my leg for me. One further philosopher will be reviewed before assessing McIntyre's example from this, non-causal, perspective.

Castaneda

Castaneda (1992) also emphasises the importance of bodily movements. There are different levels of description that we can use when describing an action. If I plan to learn algebra I do not even think of how I will open the algebra book, how I will lay out the pages to write the exercises, or how I will fold the pages when I move on to another exercise. However, I do learn algebra intentionally (Castaneda, 1992, pg. 443). When describing my action, I will say that I am 'learning Algebra', or 'studying this book', I will not describe it in terms of 'turning pages', or 'making marks on paper'. However, by drawing our attention to small physical actions he introduces the idea that the intentional action may bring with it other, unintended, actions, which nevertheless form part of the larger, intended action. He describes these bodily movements as 'intersubstitutable' because they are external to the wider project (1992, pg. 444). He means by this that, although it is necessary for me to turn the pages of the algebra book somehow, it doesn't matter how, exactly, I do turn them. The bodily movements constitutive of page turning are substitutable because they can be replaced by other, similar actions. I can replace the bodily movement of turning pages at the top with the bodily movement of turning pages at the bottom. He says of these bodily movements that "typically they seem to be performed in a reflex or habitual way. They are, when part of a complex project, intentional only by virtue of the project" (1992, pg. 444). In other words, such actions are not intentional in and of themselves, but intentional only as part of a larger, intentional, action.

Intention, on this view, is conferred 'from above', by the end point of a process, which is intentional. He does not go into much detail about how this happens. He says, 'bodily movements gain any intentionality they may have precisely because they are the initial segments of processes whose terminal points, and some intermediate ones, are intentional." (Castaneda, 1992, pg. 444), and later, 'the central point is absolutely clear: *the integrity and unity of a plan of action also comes from above* […] i.e., from the terminal point.' (1992, pg. 445. Italics in original). He believes this interpretation matches our ordinary language talk of intentions. He considers the following case: Francesca wants to kill Romeo. After considering various ways of doing this she decides on shooting him. She buys a gun and shoots Romeo. Although we can describe "Francesca's shooting of Romeo as having been done with the intention of shooting him: this intention is not a terminal point in Francesca's plan of action. We must say, appropriately, that she shot Romeo with the intention to kill him, not merely to shoot him." (Castaneda, 1992, pg. 445). I take it that this example is intended to illustrate the importance of the end point, in this case the killing of Romeo. The bodily movements leading up to this, including buying the gun, and firing the gun at Romeo, are intentional because the end point is intentional. Furthermore, when speaking precisely, we do say that Francesca shot Romeo with the intention of killing him and further, we only know this because we know what Francesca's intended end point was. However, exactly how this end point makes bodily movements leading up to the end point intentional is unclear.

A concern may also be raised about the extent to which such bodily movements are substitutable and permitted to vary. Accepting that pages can be turned at the top or the bottom is simple, but not turning pages at all, or turning ten pages at a time would change our understanding of the intentional action that is taking place. In such cases we might question whether we are still 'learning algebra'. Castaneda says that varying degrees of deviations from a plan are allowed, and that the deviations that do not "destroy the intentionality of the actions… may vary depending on the interests the community of the individuals involved have in those states" (1992, pg. 445). He expands on this later, saying that the degree of variation we can allow from a plan, while still considering it intentional, are not precise. They depend on the "moral and social gravity of the agents' purposes and actions" (1992, pg. 446). Presumably, what he means here is that if the end goal is murder, for which we believe the culprit should be punished, the variations we will allow are less, than if we are assessing someone's success in buying our shopping. Relating this to the present discussion, the degree of variation that we allow from a plan to learn algebra while still considering it intentional is, according to Castaneda, dependent on the normative framework within which we are operating. We might be willing to give more leeway to a person learning algebra for their own benefit than someone learning algebra in order to teach algebra to students. In the first case, if a person intends to learn algebra for their own benefit then we might accept that turning over two pages at a time does not change our view that they are intentionally learning algebra, if, for example, it is an easy textbook and they are very good at maths. By contrast, if the person aims to teach students, turning over two pages at a time may lead us to question whether they are intentionally learning algebra. It is important to note that Castaneda says that such judgements are not precise.

He also suggests that no project in which we engage is bounded by, or totally constrained by, our bodily movements, unless those bodily movements are the goal which we are pursuing. Bodily movements are critical if our goal is to balance a football on our nose, while they are substitutable if our goal is to learn algebra. Given this, he argues that bodily movements are only causally, not intentionally, relevant to intentional courses of action. I take this to mean that the bodily movements that form part of an intentional action are not, themselves, intentional because they are substitutable. However, they have a causal role in the intentional action of which they form a part.

We are now in a position to assess whether the word association example McIntyre describes is intentional according to the non-causal account of intentional action. In one sense it is because the subjects have, presumably, a bodily awareness that they are answering questions and learning word lists. These behaviours are under the control of the subjects. However, this question is problematic because although subjects feel as though their behaviour is under their control the frequency of certain answers (the response 'Tide' after learning the word pair 'ocean, moon') is not under their control. The response is elicited by manipulation by the experimenters. If subjects were told that the answers they give result from the learning of word pairs then they may feel a lack of control over this behaviour. Again, this highlights the importance of pinpointing the specific part of the behaviour in which we are interested.

Summary

McIntyre suggests that social scientists need to redescribe the phenomena they are interested in in different terms before they will successfully formulate generalisations. He believes that cases of word associations illustrate this process. Behaviour that we would once have described in terms of beliefs and desires are now described in terms of something else. As argued above, it isn't true that 'beliefs and desires' have been redescribed, but it is clear that they do not feature in descriptions of the situations of which word associations are an example. In his argument McIntyre relies on the assertion that this behaviour is intentional, but he provides no argument for this.

This section has reviewed the belief desire account of intentional actions, and two further accounts, the planning account and the non-causal account. The purpose of this has been to determine whether McIntyre's word association example really does describe intentional behaviour. Although a simple case can be made that the subjects involved in the word association experiments were acting intentionally, because their behaviour can be explained by beliefs and desires, it resulted from their various plans, and they had a bodily awareness of engaging in this behaviour, it isn't clear that this gets to the critical point of the example. The important point is that the subjects were not aware that their answers to questions had been, to some extent, predetermined by their learning of word lists. They had no awareness of this aspect of the experiment, so beliefs and desires, plans and bodily awareness are not relevant for understanding it. This suggests that this particular aspect of their behaviour may not be intentional. Furthermore, given that beliefs and desires, plans and bodily awareness do not feature in an explanation of what was happened it is difficult to make sense of McIntyre's suggestion that redescription is central. McIntyre's example has raised some interesting issues, most notably, that although subjects were engaged in intentional behaviour (participating in an experiment, or answering questions) the behaviour that is of interest (the frequency of certain responses) is not uncontroversially intentional because it might be a type of side effect, or simply something unintentional. In the following section I will discuss a different example, middle bias, to pin down exactly how we should understand this type of regular behaviour in the social sciences.

Middle bias behaviour

McIntyre's example, although interesting, is unnecessarily complicated due to its reliance on self-reporting by agents. A more recent, and simpler, example is taken from Attali & Bar Hillel (2003), who describe a 'middle bias' in the placement of correct answers in multiple choice questions. In a four answer choice the two middle positions should, if answer positions are genuinely randomized, contain the correct answer 50% of the time. However, a variety of studies suggest that the correct answer is in the middle positions between 70% and 80% of the time. Attali & Bar Hillel note that the 'middle bias' is present in both test setters and test takers. Importantly, they note that:

> "None of the debriefed acquaintances suspected what the point of the little exercise was. They were quite surprised to hear its true purpose- and doubly surprised at the results. They admitted no insight as to why they had placed their answers where they did and indeed seemed to have none. When inventing their questions, position was the last thing on their mind, and correct answers were positioned with little if any deliberation" (2003, pg. 111).

They also note that test takers are also unaware of their tendency to pick middle answers when they do not know the correct answer to a question, giving reasons such as "C just grabbed me" (2003, pg. 112). They also suggest that 'middle bias' is a feature of other phenomena, such as picking products on a supermarket shelf, or selecting which toilet stall to use. Middle bias is also observed when people are asked to 'pick a number between 0 and 10' (2003, pg. 115).

Middle bias partly illustrates the characteristics that McIntyre believes important. The behaviour is intentional, it is regular, it is regular across contexts, and people exhibiting this bias are not, seemingly, aware of it. When asked to give reasons for their behaviour they do not give the, apparently correct, reason. As McIntyre suggests, this type of behaviour is a candidate for a reliable generalisation, such as: When making a single choice between linearly arranged options, middle positions are favoured at a ratio of between 3 and 4 to 1 (Attali & Bar Hillel, 2003, pg. 1).

Prior to reading Attali & Bar Hillel's research we might have described question placement in terms of the beliefs and desires of agents. We might have said something like 'question setters want correct answers to be, roughly, allocated equally to different positions', or we might have suggested that they have no beliefs or desires relating to answer placement. However we think about the beliefs and desires of question setters, these beliefs and desires do not explain the pattern of answer placement. In order to explain the placement of answers we need to talk about a new phenomenon—-middle bias. McIntyre is right to suggest that a new theoretical framework might be needed to find regularities. In this case, this is exactly what has happened. Experiments demonstrate that people show a preference for middle positions, and this explains a variety of phenomena. But this still isn't a redescription of beliefs and desires. I will argue that:

1. Middle bias behaviour does not fit within the framework of intentional behaviour. It is a new category of intentional behaviour, which I will call "derivatively intentional".
2. Derivatively intentional behaviour is describable using concepts that are not Nomadic.
3. Reliable generalisations can be formulated for derivatively intentional behaviour.

1. Middle bias behaviour does not fit within the framework of intentional behaviour. It is derivatively intentional.

Intentional action can be analysed in terms of beliefs and desires, plans, or bodily awareness. However, as the discussion above illustrates, there is little consensus about firstly, whether there can be action which is not intentional, and secondly how to distinguish behaviour that is intentional from behaviour that is not. Nevertheless, there is no need for a unified theory of intentions to use these theories to assess middle bias, primarily because it does not fit within the bounds of any of these accounts particularly well.

It is difficult to describe middle bias as intentional in a standard sense because the agents involved have no psychological, conscious attitudes, whether beliefs or desires, or plans, regarding this behaviour. Nor does it fit within Bratman's planning account, because biasing the middle forms no part of any plan. Neither does it fit well within a non-causal account concentrating on bodily awareness. Following Frankfurt's account, there is no awareness of performing this action which can be interrogated to assess the level of control an agent has over it. The lack of awareness suggests that there is no control whatsoever over the phenomenon, as ordinarily understood. Using Grunbaum's example it fits at one end of the continuum of bodily awareness; it is more like a reflex, over which we have no feeling of bodily control than a premeditated action. Question setters know that they are writing or typing, but do not have an awareness of the placement of what they are writing. Indeed, they express surprise when they are told that they are 'biasing the middle'. It does not even fall within Davis's focus on motivation to act, because middle bias does not motivate anyone to act, at least not in an obvious sense. Middle bias does not fit within any of the accounts of intentional action reviewed here.

However, at the same time, this behaviour is the action of 'self-consciously active human agents'. The agents in question are engaging in behaviour that, when taken as a whole, is intentional. They are setting questions for exams or quizzes, or trying to answer these exams and quizzes, or making a, seemingly random, selection from a linearly arranged set of options. They intend to set or answer questions and they intend to make a choice.

If it isn't intentional, as ordinarily understood, is 'middle bias' a side effect, or unintended consequence of intentional action? I argue it is neither. Pitcher and Bratman's discussion highlights that side effects or unintended consequences are usually things that we are aware of doing— waking Susan, or annoying the neighbours. Much of the focus on side effects also concerns the allocation of responsibility. Neither of these features are relevant to the middle bias example. People are unaware of their tendency (at least until they are told about it), and given the incidental nature of this behaviour, issues of responsibility do not arise. However, the discussion of side effects does provide a useful insight. The problem with side effects is that they do not, obviously, fall into either the category of 'intentional action' or 'unintentional action'. Garcia introduces the idea that actions can be partially intentional. But his proposal is problematic for the middle bias example as it is based on the idea of us achieving 'more or less' what we intended. In the case of middle bias, how closely we achieved what we set out to do does not affect the extent to which our behaviour illustrates the middle bias phenomenon. For example, we might intend to set an exam with 50 questions, but actually set one with 40 questions. We have 'more or less' done what we intended, but the extent of the middle bias phenomenon does not vary with how closely we match our original plan. Nevertheless, there does seem something right about describing the middle bias behaviour as partially intentional. It isn't entirely unintentional, because it just is part of setting questions, or picking answers, which is intentional.

The account which matches the middle bias phenomenon most closely is Castaneda's. He discusses bodily actions that are only intentional by virtue of the action of which they are a part. The way we turn pages when reading an algebra book, for example. However, there is an important difference here. Castaneda's physical actions are substitutable. We can turn the pages of our algebra book at the top, or the bottom. We will need to turn them in some way or another, but it doesn't matter to our learning algebra how we do it. The difference with the middle bias phenomenon is that the physical action of placing answers might be theoretically interchangeable, because we could place correct answers in any place, but in fact, it follows a regular pattern. In an important sense it is not substitutable because we have a tendency to do it with a particular bias. The fact that the middle is biased affects the way that test takers and test setters behave. Page turning, when learning algebra might be described as a 'neutral' phenomenon, which has no effect at all on the end goal. By contrast, biasing the middle may affect the outcome—if test setters and test takers both show this tendency we can expect a higher pass rate than if test setters followed a policy of randomly selecting answer positions. On the other hand, it is possible that page turning follows a pattern. Were we to study it we might find that people turn 70% of pages at the top, and that this is a regularly observed pattern. But, in this case, it would still have no effect on the outcome. We would not be more successful at learning algebra if we moved from turning 70% of pages at the top to turning 50% of pages at the top.

However, the other aspect of Casteneda's description of intentional action applies directly to the middle bias phenomenon. This is that these actions are only intentional because of the behaviour of which they are a part. Castaneda says that the intentional action may bring with it other, unintended, actions, which nevertheless form part of the larger, intended action. This matches the middle bias phenomenon. I suggest that there are two aspects to middle bias as it applies to the placement of answers. The agents intend to set an exam question and in doing so their intention is directed at wording the question appropriately, writing down the correct answer, and thinking about false answers that can plausibly be included in the set of possible answers. However, in engaging in this intentional action they are also engaging in other behaviour, which is unintentional, and this is the bias to place correct answers in the middle of the list of answers. The two types of behaviour are interwoven because the unintentional part of the event only happens when the agent engages in the intentional action. In contrast with Castaneda's example, the middle bias phenomenon may affect the outcome of the intended task. Nevertheless, the unintended action depends for its existence on the intentional action.

There is a further important distinction. Following Frankfurt's suggestion that we can direct our awareness towards a bodily movement and 'take control' of it, although the tendency to bias the middle is not intentional, it might become so. For example, if I were to set an exam with the knowledge that I have a tendency to bias the middle and that my students have a tendency to pick answers in the middle I might well alter my tendency by self-consciously placing more correct answers at the edges. Therefore, although this behaviour is not intentional, it has the potential to become intentional once I am aware of it, and decide to have intentions about it.

To summarise; middle bias behaviour does not fit neatly into the theories of intentional action. The behaviour of setting multiple choice exam questions is intentional for the agent setting the questions. We can describe this as intentional regardless of which analysis of intentional action we choose to accept. Setting exam questions could be caused, in the right way, by agent's beliefs and desires about setting exams to test students, and their beliefs about how to go about this. Setting exam questions forms part of various plans, for example, to execute a task that we have been asked to do, to fulfil the obligations of a job, and to get through the day. While setting exam questions agents also have a bodily sensation of control over this behaviour. Questions are brought to mind, possible answers thought of and selected, final ideas are recorded. However, this intentional behaviour brings with it other behaviour—that of biasing the middle. This behaviour is only intentional because of the larger intentional action of which it is a part. In Castaneda's analysis, this intention is conferred from above. However, were the agent to become aware of this behaviour, and have intentions about it, then this behaviour could become fully intentional.

O'Shaughnessy & Sub-intentional actions

At this point, an insight from O'Shaughnessy is useful. O'Shaughnessy discusses a very similar behaviour to middle bias, which he calls 'sub-intentional' (1980 vol 2, pg. 59). He begins by accepting Davidson's general point that ascription of intention is description relative; but he rejects Davidson's idea that all actions are intentional under some description. He says, "I am convinced that there are exist acts that are intentional under *no* description… there *does not exist* a description under which they emerge as intentional" (1980 vol 2, pg. 59). These are sub-intentional acts. Examples include the movement of a person's tongue, or fingers, or toes, while they are performing other actions, such as listening to music or talking to a friend. O'Shaughnessy's reason for believing these actions to be intentional under no description is our unawareness of their occurrence. While we may have beliefs, such as that our limbs were in a particular position, we are not aware of these beliefs. Further, that belief does not form part of the reason for moving the limb, because there was no reason for doing so.

What determines whether an act is sub-intentional is "the character of its origin" (O'Shaughnessy ,1980 vol 2, pg. 59). By this he means whether 'the faculty of reason' has played a causal role in the generation of the act (1980 vol 2, pg. 60). However, he also says "…in the final analysis it is determined by the fact that no intention finds expression in the act". His first explanation suggests that sub-intentional acts are those we haven't reasoned about, but the second suggests that they are simply things we have not intended. These two descriptions are at odds with one another because, as the discussion about side effects suggested, there may be acts that we haven't reasoned about, but which we can describe as intentional.

What he means becomes slightly clearer with an example. O'Shaughnessy argues that when we think about what our bodies are doing at any point in time we will notice actions that we are performing, but which we did not intend to perform, such as drumming our fingers. The difference between this type of behaviour and the movement of our pupils or heart, for example, is that we are responsible for the first, but not for the second. When we notice the drumming of our fingers we become aware of an action that we are performing. It is something I am doing. This is not the case for the movement of my heart. Furthermore, before I notice the movement of my fingers I am unaware of it.

Once I do notice the movement of my fingers the action can become intentional, which is an interesting parallel with Frankfurt. I can decide to move my fingers differently, or stop moving them. So, an action that begins as a sub-intentional one can become intentional. For Grunbaum, bodily awareness is an important part of describing an action as intentional, and this bodily awareness is what is lacking in sub-intentional actions. We are not aware of sub-intentional actions, and therefore we do not reason about them in any way. Once we notice them, we may bring them under our control.

Sub-intentional action and middle bias

How, then, does this help our analysis of middle bias behaviour? As with middle bias, sub-intentional actions are things that we do, that are not unintentional, in the way that reflexes are, because we can bring them under our control. They are both things we are unaware of doing, while we are doing them. Sub-intentional actions are intentional under no description, and this is arguably the case for middle bias too. A number of descriptions are available for question setting behaviour. We could describe it as 'putting pen to paper', 'trying to tax the students', 'fulfilling the obligations of our job', 'making up for an easy test last time' etc. However, alternative descriptions, that might make the behaviour intentional are just not available for the middle bias aspect of the event of question setting. The phenomenon just appears. O'Shaughnessy says that this is because such actions do not involve any 'reason', in the sense of conscious mental machinery. Although this leaves open the problematic issue of how we are to understand 'reason', his account captures an important aspect of the middle bias behaviour. We are not aware of it, do not reason about it, and therefore it is difficult to see how it can be described as anything other than what it is—a tendency to place answers in the middle of the answer set. Under this one description it is not intentional, and there are no other descriptions available under which it might be made intentional.

However, there is an important difference with O'Shaughnessy's theory of sub-intentional action. The middle bias phenomenon is derived from other, intentional, behaviour. There is something intentional going on that is bound up with the middle bias behaviour. This is not the case for O'Shaughnessy. He describes physical actions, like tapping fingers occurring while other actions are taking place, such as talking to a friend or reading a book. The sub intentional, physical, action is not directly related to the intentional events that are occurring at the same time. Finger tapping could continue even after the music has come to an end, or after the conversation has finished. Middle bias comes to an end as soon as question setting is finished because it is a part of this larger action. Following Castaneda, we can confer some degree of intentionality on middle bias because it forms part of an intentional action. In his words, intention is conferred 'from above'.

A natural question here is whether side effects are also sub-intentional actions. I argue that they aren't. In Bratman's (1999i) example he arrives home late and wonders whether to turn on the light. Doing so will allow him to see, but will also wake up Susan. He decides to turn on the light, intending to see better, and expects, but does not intend, to wake Susan. Bratman gives three reasons for thinking that he does not intend to wake up Susan. Firstly, he is not inclined to pursue other means to wake Susan up, if it turns out that the light isn't working; second, he doesn't rule out other potential actions because they are incompatible with waking Susan; thirdly, he doesn't see himself as faced with a problem of how to wake Susan (1999i, pg. 258). This means that, for Bratman, only the belief about turning on the light does the causal work in motivating action. In other words, "I have my turning on the light as a goal, but I do not have as a goal my waking Susan. It seems to me that this difference is reflected in a difference in intention" (1999i, pg. 261). Using Pitcher's framework, side effects fall through the gaps in the distinction between intentional and unintentional action. However, this means that they aren't unintentional. Whereas, in contrast, sub intentional actions are unintentional because, where they are physical or mental, they are unrelated to our intentions. Furthermore, as O'Shaughnessy describes them, we are unaware of sub-intentional actions, whereas we are, usually, aware of side effects, in the way that Bratman suggests. We are now in a position to understand the phenomenon of middle bias:

- It is not entirely unintentional, because, following Castaneda, it forms part of a larger intentional behaviour. So, it is, to some extent, intentional by virtue of being part of question setting, or picking answers. From O'Shaughnessy we can take the additional reason that this behaviour is something that, when we become aware of doing it, is something that we are doing. It is not something that happens to us. Following Frankfurt and

Grunbaum, it is also something that we can bring under our control, and make intentional.

- It is not a side effect of intentional action, because side effects are usually understood as things we are aware of doing. To the extent that side effects can be things that we are unaware of doing, middle bias is a type of side effect. However, because middle bias behaviour is tied in with, and an integral component of intentional behaviour it is not a side effect, as ordinarily discussed in the literature. These side effects are usually not 'tied in' with the intentional action. They include stepping on an ant, waking Susan, hitting a tree. None of these side effects form a part of the intentional behaviour as a matter of course. In Bratman's case, for example, turning on the light only wakes Susan because she happens to be asleep. In other circumstances he could turn on the light without this side effect. Middle bias occurs whenever questions are set.

- It is not intentional, because, following O'Shaughnessy, no conscious mental machinery is involved. We are totally unaware of this phenomenon (until someone tells us about it). Placing this within the usual understanding of intentional action reviewed earlier, it forms no part of plans, it is not caused by any beliefs and desires, and we have no physical awareness of it.

- It is not sub-intentional, because, although it shares many attributes with O'Shaughnessy's examples, middle bias is 'tied in' with intentional behaviour in a way that O'Shaughnessy's examples are not.

I will use the term 'derivatively intentional' to describe middle bias behaviour and other, similar, behaviour. Derivative, because it is derived from, and dependent on, some intentional action. However, it can be analysed as an independent phenomenon. The reason for picking the word 'derivative' is that financial derivatives are a good metaphor. Derivatives, such as options and futures depend, for their very existence, on the underlying asset on which they are based. You can only have an option on Microsoft shares, when Microsoft shares exist. However, these options on Microsoft shares trade as independent entities, and are analysed as such. Nevertheless, the price of the options is related, albeit in a complex way, to the price of Microsoft shares. In the case of middle bias, the underlying behaviour is the question setting or answering, or the picking of an item. This is analogous to the Microsoft shares. Along with this behaviour is derived behaviour, which is the tendency to bias the middle. This is analogous to the options on Microsoft shares. The derivative behaviour only happens when the underlying behaviour takes place, in the same way that options on Microsoft shares only exist when Microsoft shares exist. Nevertheless, in the same way that option prices move independently of the movements in Microsoft's share price, the derivative behaviour can be analysed separately to the underlying behaviour. This combination of dependence and independence seems to fit the middle bias behaviour quite well. The key disanalogy is that options are intentionally created whereas as derivative behaviour is discovered. Derivatively intentional behaviour is therefore defined by the following necessary, and jointly sufficient, conditions:

1. Derivatively intentional behaviour is a subset of intentional behaviour.

2. Derivatively intentional behaviour occurs whenever the intentional behaviour (of which it is a subset) is undertaken.

3. An actor is, initially, unaware of this derivatively intentional behaviour and has no beliefs, desires, plans, or other deliberations relating to the derivatively intentional behaviour.

4. If an actor becomes aware of the derivatively intentional behaviour they realise that it is something they are doing. It is not something that is merely happening to them.

5. If an actor becomes aware of the derivatively intentional behaviour, it can be brought under the actor's control, if they choose to do so.

2. Derivatively intentional behaviour is describable using concepts that are not Nomadic.

The previous section outlined O'Shaughnessy's argument that sub-intentional actions are intentional under no description. Although I want to resist tying derivatively intentional actions to the belief-desire analysis of intentional action, this point is critical to linking this discussion about intentional action to the discussion about social phenomena is Chapter 2. This section will show that the possibility of redescribing actions is related to the possibility of attributing different meanings to action.

One of Davidson's examples of redescription is the following: The same physical action can be described as 'flipping a switch', 'turning on a light', 'illuminating a room', 'alerting a prowler that I am home'. He uses these alternative descriptions to argue that while I may intentionally have 'illuminated the room' I did not intentionally 'alert the prowler'. When we give a reason for an action, such as flipping the switch, we say we did it because we wanted to illuminate the room, we redescribe the action, finding a place for it in a pattern of what we believed and desired. However, this is not my concern here, because I do not think that accepting the possibility of redescription necessarily ties one to accepting a belief-desire account of intentional action. Consider a person driving down the motorway. We can use a non-causal understanding of intentional action and say that the person is engaged in an intentional action because they have a bodily awareness of control over the car. They can press the accelerator and increase the speed, or turn the steering wheel to alter the direction. However, we can describe this bodily awareness as being about different things. It could be bodily awareness of 'driving the car', or it could be bodily awareness of 'driving to work', or even bodily awareness of 'moving my foot down'. We can redescribe the behaviour without referencing beliefs and desires. Therefore, the possibility of redescribing actions does not seem, necessarily, tied to the desire-belief model of intentional action. It is possible just to redescribe behaviour in various ways.

The reason for labouring this point is that there is a parallel between saying that behaviour can be redescribed and saying that behaviour can have many meanings. To return (again) to Taylor's example of raising my hand. In raising my hand, I may 'cast a vote', 'save the honour of my party', 'return a favour', 'act in accordance with my beliefs' etc. Taylor points to one physical act and notes that it could mean many different things. In many critical respects this is the same thing as saying that behaviour can be redescribed.

In Davidson's example we can ask what flipping the switch meant for the person performing the action. By flipping the switch, they could have meant that they 'wanted to see what is in the room', 'wanted to turn on the light' or 'wanted to alert the prowler'. In Taylor's case, Davidson's analysis seems to apply uncontroversially. We can redescribe the behaviour as 'raising my hand', as 'saving the honour of my party', 'returning a favour' and so on.

However, these examples do suggest one important difference between working out what an action means and redescribing it. When we try to work out what an action means we, usually, assume that we are getting at what the person meant. It is impossible, in Davidson's example, for the person to have meant to 'alert the prowler', because they were unaware of the prowler's presence. When redescribing in Davidson's sense we are limited only to things of which the person is aware. With Taylor's example it would not make sense to say that the person raising their hand meant to 'wave to their friend' if they hadn't noticed that their friend was looking at them from across the room. When redescribing this is not a limitation—we can redescribe 'flipping the switch' as 'alerting a prowler' because this is one of the things that happened, but we can't say that by 'flipping the switch' the person *meant* to 'alert the prowler' because for a person to have meant something by their behaviour we, usually, assume that they were aware of what they meant.

What this suggests is that there is a symmetry between the ability to redescribe behaviour, and the ability to attribute different meanings to behaviour. However, the symmetry breaks down when we use redescription to point to something of which the agent was unaware. One way to bring the two into alignment is to draw an equivalence between meaning and intentional action. 'Alerting the prowler' was not what the actor intended by 'flipping the switch', nor was it what the actor meant by it. In other words, the different meanings we can attribute to behaviour line up with the possible descriptions that make behaviour intentional. This excludes descriptions of which the agent is unaware; alerting the prowler in Davidson's example. To clarify, 'what was meant' means what the action meant to the actor, not what the action meant as some sort of symbolic, or communicative, role. This latter notion of meaning does not apply in this case at all. The concern with suggesting that there is a symmetry between intentions and meanings in this way is that bringing intentions in, in this way, involves a commitment to a causal understanding of intentional action. Fortunately, this is a problem that can be sidestepped because, for the purpose of this argument, it can be established that *just for* derivatively intentional behaviour neither redescription, nor attributing multiple meanings is possible. These two points will now be argued for in turn.

It is not possible to redescribe derivatively intentional behaviour for the following reasons. Firstly, it is important to understand exactly what middle bias is, although this is difficult given the current state of research on this topic. Attali and Bar-Hillel review a range of data on middle bias, which yields slightly different numbers. However, they say that in single, isolated, questions test setters and takers prefer middle positions to extreme ones in a ratio of up to 3 or 4 to 1 (Attali & Bar-Hillel, 2003, pg. 1). In a later paper Bar-Hillel reviews a range of experiments yielding slightly different results. For example, significant differences are found when options are presented serially versus simultaneously, and it is often difficult to isolate middle bias from other position biases (Bar-Hillel, 2011). Nevertheless, in a commentary on her paper Rodway *et al* comment that "middle preference is surprisingly robust and widespread, having been found in a wide range of perceptual-motor tasks" (Rodway, Schepman & Thoma, 2016, pg. 1). Middle bias appears to be a robust phenomenon although it is not yet fully understood, nor comprehensively described. In advance of this it is difficult to formulate a precise definition. However, given the current state of play, let us take middle bias just as it relates to multiple choice questions and take it to indicate 'a preference for the middle options in a ratio of 3 or 4 to 1'.

How is it possible to redescribe this? We could describe it as 'an aversion to options on the edges in a ratio of 3 or 4 to 1', but this is simply a different way of saying the same thing. In multiple choice sets choices are either 'middle' or 'edge' so there is no fuzzy 'intermediate' category. It could be objected that 'illuminating the room' and 'flipping the switch' are similarly just different ways of saying the same thing, but I dispute this. These two descriptions are not simply synonymous because each description tells us something different about the event in which we are interested. If we are just given the description 'flipping the switch' this leaves open what the function of the switch is, it might have turned on the radio or opened a trap-door. Even when we know that the switch illuminated the room the use of the word 'flipping' indicates that the light is neither dimmable nor motion sensitive. These descriptions give us more information. The redescription of middle bias gives no additional information. Therefore, I argue, behaviour that is derivatively intentional cannot be redescribed in ways that provide additional information about the behaviour.

It is not possible to attribute multiple many because such behaviour, before we become aware of it, has no meaning at all for the actor. Asking a person displaying middle bias behaviour what they mean by it makes no sense. They do not mean anything by it, because they are unaware of it. This is primarily because no conscious cognitive faculties have been brought to bear on the behaviour. We can ask what they mean by the larger behaviour, that of question setting, because that is what they are thinking about doing, that is, the task on which they are concentrating. Middle bias behaviour just comes along with it and we have no attitudes towards it, let alone meanings. It is the same for Castaneda's algebra example, where there is nothing we meant by turning the pages of our book at the top rather than the bottom. Furthermore, there aren't multiple meanings we, as observers, can attribute to actors.

So, not only is there a connection between the ability to redescribe behaviour and the ability to attribute many meanings, but the reason we can't attribute many meanings, in this instance, is that the behaviour involves no conscious mental machinery that motivates us to act. At first sight, this suggests a commitment to the belief and desire model of intentional action. However, I want to deny this by suggesting that the reasoning is the other way around. It is not particular descriptions that, necessarily, make behaviour intentional; it is the presence of motivation and conscious 'mental machinery' that lead both to the possibility of multiple descriptions and multiple meanings. Using Bratman's terminology, there is no intention, as a mental state, involved in the behaviour that is derivatively intentional.

However, one obvious objection is that it is possible to ask what causes middle bias behaviour. Bar-Hillel proposes that middle bias is explained by the middle option being easier to reach, either physically or mentally (Bar-Hillel, 2011). This is easiest to see in physical tasks where people may reach for the middle object more easily. Rodway *et al* dispute this, arguing that a better explanation is a decision heuristic suggesting that the best, and most popular items are positioned in the centre. To substantiate this they cite experiments showing that middle bias is much less prevalent when people are asked to select from a set of identical bad paintings or ugly faces. Furthermore, when people were asked to pick which item they preferred from a set of identical objects middle bias was present, but when they were asked which they least preferred, middle bias disappeared.

(Rodway, Schepman & Thoma, 2016, pg. 3) So, it is possible to ask whether middle bias *means* that people subconsciously think that middle options are the 'best' option, or whether it *means* that people think middle options are easier to reach. However, this just reflects our current ignorance. Presumably, at some point, the question will be answered definitively and there will be only one thing that middle bias means about human beings.

Exactly how we are to understand the relationship between meaning and redescription when they are possible can be dealt with at a later date. The conclusion that for derivatively intentional behaviour neither informative redescription nor multiple meanings for the person engaging in the behaviour are possible allows us to bring together the analysis of middle bias with the analysis in Chapter 2. The argument is as follows:

i. Derivatively intentional behaviours are describable using concepts that are less Nomadic.

ii. Therefore, we can make reliable generalisations about derivatively intentional behaviour.

i. Derivatively intentional behaviours are describable using concepts that are less Nomadic

Chapter 2 explained that some concepts are better for making generalisations. Concepts that are better for formulating generalisations are those where there is little possibility for disagreement over which phenomena fall within the scope of a concept. In other words, concept that are less Nomadic. Central to this is the number of meanings people can have in mind when using a particular concept. Other characteristics are the the precision of the boundaries of these different meanings and the extent to which these meanings change over time. If a concept has only one meaning, and this meaning is precise, then it should be a promising one for using in generalisations.

I have established that derivatively intentional behaviour cannot be redescribed in such a way as to make it recognisably intentional from the point of view of the person performing the action, and it does not have multiple meanings for the actor. It remains to link derivatively intentional behaviour to the concepts with which we analyse it. To do this we need to distinguish between analysing the behaviour of a single actor and analysing at a more general level, making use of concepts.

- Derivatively intentional behaviour and the single agent: An agent setting multiple choice questions engages in the derivatively intentional behaviour of biasing the middle. Consequently, this behaviour has no meaning for him, and, when analysing his behaviour there aren't multiple descriptions of this behaviour that are not synonymous with the original description. This follows directly from the discussion of derivatively intentional behaviour above.

- Derivatively intentional behaviour and concepts: When analysing the phenomenon of middle bias we still say that the behaviour has no meaning for the agents, but we, as external observers, clearly mean something when we talk about middle bias. But it means only one thing. When we talk about middle bias we mean the behaviour people show when picking items from linearly arranged options. This is all we can mean. Once we understand what causes this behaviour we will also be able to say what it means about human beings.

What is the connection between the lack of meaning for an actor and the existence of only one meaning when an observer talks about middle bias as a concept? The reason why some concepts used in the social sciences can have many meanings is that we are aware of the things agents may have meant when engaging in the behaviour that we include within this concept. Consider again an example from Chapter 2. The concept of a 'democracy' can mean many things. These include 'participation' and 'contestation' to list just two. Why do we think the concept of 'democracy' can mean these different things? At least in part, this is because it's plausible to think of the people who are engaging with a political process characterised by 'contestation' or 'participation' as thinking of themselves as members of a 'democracy'. In other words, the actions associated with engaging with a political process are intentional in an ordinary sense, be it resulting from beliefs and desires in the right way, forming parts of plans, or bodily awareness of control over these behaviours. These behaviours do involve conscious mental machinery which motivates agents to act in various ways. We know that these behaviours are intentional in an ordinary way because, when we ask people about what they are doing they will talk about what they intended, or wanted, or hoped to do. Because of this the things people do are amenable both to redescription and multiple meanings. When this behaviour is analysed, we are aware of these possible descriptions, and meanings actions may have for agents, and therefore describe them using concepts which import this possibility of multiple meaning and description.

Furthermore, when behaviour is intentional in an ordinary sense people consciously think about their behaviour, be it in terms of plans, beliefs and desires, or the awareness of how they are executing a task. This means that each instance of a 'democratic process' will differ, in some respects, from other instances of a 'democratic process'. Agents are usually aware of how other people have behaved in similar situations, reason about their intended behaviour, and alter their behaviour accordingly. The importance of this is that even when agents intend to create a 'democracy', or a 'democratic process' they are likely to do it in a different way to other people. When analysing 'democracy' we are therefore faced with social phenomena that are not uniform. Derivatively intentional behaviour, on the other hand, is relatively uniform because we are usually not made aware of it. When we are made aware, this behaviour will be less uniform.

ii. Therefore, we can make reliable generalisations about derivatively intentional behaviour.

Let us consider a generalisation about middle bias.

When making a one-off choice between linearly arranged options, middle positions are favoured at a ratio of between 3 and 4 to 1 (Attali & Bar-Hillel, 2003, pg. 1)

The relevant concepts here are 'choice', 'linearly arranged options', 'middle positions', 'favoured'. Within the context of the cases we are applying this generalisation to, there is not much, if any, scope to question what we mean by these concepts. We could debate what a 'choice' is, but, when presented with a number of options and asked to 'choose one', this concept is uncontroversial. We might also question what it means to be in the 'middle'. When picking between four options the middle is clearly the middle two. However, when picking between five, do we count just the central position as the middle, or also the positions either side of this central position? Similar questions arise as we include more items. However, unlike many of the concepts discussed in chapters 2 and 3, we can easily settle on our meaning. Once we have done so the concept does not have boundary issues. Additionally, if you disagree about my specification of the 'middle' you can use the same data and widen what counts as the 'middle'. So, the concepts used in this generalisation are better suited for formulating generalisations than many of the concepts discussed in prior chapters.

The purpose of this section has been to show that the phenomenon of middle bias is describable using concepts that make formulating reliable generalisations relatively uncontroversial; these concepts are not Nomadic. This is because middle bias behaviour is derivatively intentional, which means that it is not subject to multiple descriptions or meanings. The importance of this is that, when analysing this behaviour, we can do so using concepts that do not have multiple meanings, and therefore the concepts cover a small, and relatively precise, section of phenomena.

Conclusion

This chapter began with a discussion of McIntyre's argument that the social sciences are in need of redescription before we can make reliable generalisations. In particular, he argues that we need to redescribe beliefs and desires in other terms. This chapter has sought to demonstrate that this is not true. Behaviour that is regular and predictable is behaviour in which beliefs and desires do not feature at all. The second part of the chapter analysed McIntyre's assertion that the regular behaviour he describes is intentional. Taking each of the main interpretations of what makes actions intentional: the desire belief account, the planning account and non-causal accounts, I argued that, although *prima facie* this behaviour is intentional, it is difficult to describe the aspect of this behaviour that is regular as intentional. The final part of the chapter analysed a different example of regular behaviour, middle bias, to explore in more the detail the characteristics that this type of regular behaviour has. The conclusion is that the middle bias example illustrates that regular and predictable behaviour is derivatively intentional. It is not intentional in and of itself, but derives intentionality from wider, intentional, behaviour of which it is a necessary part. I propose that this is an example of a regularity in the social sciences and that other behaviour sharing these characteristics is likely also to be regular.

Additionally, a link was made with Chapter 2, by arguing that behaviour that is derivatively intentional is describable using concepts that are not Nomadic. Derivatively intentional behaviour is therefore the best candidate, so far, for formulating reliable generalisations in the social sciences. The worry that this type of behaviour is not what the social sciences are about is a concern that will be addressed in the following chapter.

Chapter 5: Clarifications and further work

This chapter attempts to anticipate some clarifications that may be called for after reading the prior chapters. Firstly, if many concepts that social scientists use really are Nomadic, this suggests that formulating policies in the social sciences is difficult. This chapter discusses an example of social science research and shows that working with Nomadic concepts is possible, and that recognising that many concepts social scientists use are Nomadic may make social science research more successful. Secondly, if it is true that behaviour that is derivatively intentional can be analysed using concepts that are less Nomadic, then the less Nomadic concepts in Chapter 2 should also be derivatively intentional. This chapter shows that the concept of 'demand', which was described as less Nomadic in Chapter 2, does share some characteristics of derivatively intentional behaviour. Finally, the middle bias example may be seen as uninteresting from a social science perspective. This chapter discusses some other, potential, examples of derivatively intentional behaviour.

1. The idea of Nomadic concepts has troublesome implications for social policy.

If policies are formulated using Nomadic concepts it appears, from the analysis in Chapter 2, that it will be difficult for social scientists to agree about whether a policy is the right one, and how to decide whether it has been successful. The discussion in Chapter 2 illustrated these problems with the poverty example. Social scientists disagree about what the relevant meanings of 'poverty' are, and even if they agree about *this*, how they should measure changes in 'poverty'. However, the purpose of the Nomadic concepts framework is not to suggest that working on social policy is a waste of time. The following example is intended to show that working with Nomadic concepts is possible and, when done well, can help with the formulation of successful policies.

The Institute for Fiscal Studies (IFS) recently published a report on the provision of breakfast clubs at schools (IFS, 2017). The Magic Breakfast project was a randomised controlled trial in the UK that aimed to assess the impact of providing of a free breakfast before school on academic attainment (which was measured by teacher assessed performance in English and Maths tests). The trial ran for one academic year and focussed on schools in relatively deprived areas. The concept of ‘academic attainment’ is synonymous, in this study, with performance on national tests in English and Maths taken at the end of Year 2 and Year 6. The trial provided evidence that providing breakfast improved the results in Year 2 tests by approximately 2 months’ worth of progress. The effect on the Year 6 test results was positive, but less significant.

The concept of ‘academic attainment’ is a Nomadic concept, and it is debatable whether it just means ‘how children perform in national tests in English and Maths’. While English and Maths are, quite reasonably, the core of the school curriculum, there are other ways in which children can demonstrate ‘academic attainment’, for example in music, or science. Furthermore, a child might perform well in a relaxed classroom setting but struggle with formal tests; this is particularly plausible with Year 2 children, who are 7 years old. We might also think that ‘academic attainment’ might be improved without an increase in test scores, if children are more interested in class, ask more questions, and get more work done. Improvement is test scores might lag such improvements in ‘academic attainment’ beyond the length of the study. ‘Academic attainment’ can therefore mean a number of different things. One of these meanings, the test scores, is precise, but the others are not. Assessing the level of interest in class, or greater concentration, is a subjective matter. After all, for one child showing greater interest in classes might mean reading more, while for another it might mean not hitting other children in lessons.

At first sight, the prior paragraph suggests that this trial focussed only on one meaning of ‘academic attainment’. Although performance in tests was the primary focus of the trial, other data was collected to assess how providing breakfast affected test scores. Teachers completed surveys on ‘Classroom behaviour and concentration’, ‘attendance at school’, and data on ‘health’ was also gathered (this was primarily weight and height data). Children and parents were also asked to comment on the trial. In summary, while the trial focussed on the effect of breakfasts on test scores the surveys provide information relevant to the other things that ‘academic attainment’ can mean. This other data proved informative because points 3 and 4 of the summary document say:

> "3. The findings suggest that it is not just eating breakfast that delivers improvements, but attending a breakfast club. This could be due to the content of the breakfast itself, or to other social or educational benefits of the club.
>
> 4. Pupil behaviour, as measured by a teacher survey, improved in breakfast club schools. This is interesting because it shows that breakfast clubs may improve outcomes for children who do not even attend breakfast club, by improving classroom environments." (Education Endowment Foundation, EEF Projects, 2017, pg. 1)

The information relevant to these other meanings of 'academic attainment' was relevant to making sense of the trial. The teacher surveys were conducted because it was believed that they would help to discover the mechanisms by which breakfast provision influenced test scores, not because improved behaviour and concentration were seen as things that 'academic attainment' can mean. This is a relevant difference. If it had turned out that despite improvements in behaviour and concentration test scores didn't improve, the trial would have shown no link between breakfast provision and academic attainment. Recognising 'academic attainment' as a Nomadic concept would lead social scientists to use more diverse criteria for judging success in trials like this.

Point 3 also highlights that different breakfast clubs functioned in different ways. Specifically, in some cases the clubs had social benefits rather than nutritional benefits. The report says, "the case study evidence suggests that staff perceived the social environment of the breakfast club as positive for children's relationships and independence." (Education Endowment Foundation, Magic Breakfast, 2016, pg. 38). The breakfast clubs often played games, encouraged interaction between children in different academic years, and developed relationships with teachers outside the formal school regime (pg. 57-8). It is possible that in one school the breakfast itself improved academic performance because the children did not receive sufficient food, or food of the right sort, at home. In this case the extra calories made the difference. In another school, the children may have already been well fed but the breakfast club allowed them to chat with teachers about problems and difficulties they were having that negatively affected their academic performance. In this case it isn't the provision of breakfast that made the difference, but the chance to chat informally with teachers. Providing breakfast may mean very different things, and there are a number of ways in which breakfast clubs can positively affect academic attainment. It would be very difficult to formulate a policy that specifies all these things precisely.

In conclusion, 'academic attainment' is a Nomadic concept. The breakfast club trial suggests that when social scientists are dealing with Nomadic concepts, a number of meanings of concepts may be relevant to assessing the results of a trial. Recognising this can lead to the inclusion of a wide range of evidence when assessing improvements in academic attainment, rather than just focussing on the most obvious evidence provided by test scores. Also, when dealing with Nomadic concepts, an 'intervention' such as breakfast provision may be successfully implemented in many different ways. Breakfast clubs can mean different things for different schools, ranging from a provision of food, to a chance to mix with teachers on an informal basis. This heterogeneity is only a concern if social scientists are seeking to formulate prescriptive policies, but the success of the Magic Breakfast trial suggests that heterogeneity in breakfast clubs allows a variety of pupils to benefit in different ways. The Nomadic concepts framework does not suggest that social policy is difficult to formulate, but it does suggest that a wide range of evidence is relevant in assessing such concepts, and that effective policy is likely to be heterogeneous. A successful policy may be implemented very differently, so much so that two examples of the same policy have little in common.

2. How does derivative intentionality fit in with the regularities discussed in Chapter 2?

Chapter 4 suggested that behaviour that is derivatively intentional can be analysed using concepts that are less Nomadic. Chapter 2 also gave examples of concepts that are less Nomadic. If derivative intentionality really is a useful concept, then concepts that are less Nomadic should also be derivatively intentional.

The following section argues that this is true, and can be demonstrated by re-visiting the law of supply and demand introduced in Chapter 2. As a brief reminder, the concepts used in formulating this law are 'demand' and 'supply'. Page 38, above, showed that the concept of 'demand' is less Nomadic than the concept 'social exclusion'. This is because 'demand' means the amount of a good or service that a customer is willing to buy at a particular price. Although there are two different meanings of 'willingness to buy', these are closely related, and precisely defined. Furthermore, 'demand' in an economic sense, is unlikely to change over time.

How does the idea of derivative intentions fit into this description? The following paragraph will show that the concept of 'demand' fulfils some of the criteria of derivative intentionality. These criteria will be discussed in turn.

a. **It is behaviour that is a subset of intentional behaviour.**

   'Demand' for a good or service is something that interests economists, but it is difficult to see how 'demand' is something about which people have intentions directly. For example, we may each demand a number of things at various prices, and we may intend to buy them, now, or when the price is right. What is less clear is that our intentions are about the level of 'demand' for these goods and services. There are some cases when we do care about the level of demand; cases where we intend that the demand for a particular good be at a particular level. This could occur if we want to stop people buying something, such as animal fur. In this case we have intentions about reducing the level of demand for this product, but in most cases, we have intentions about buying certain items; the level of demand is not something that concerns us. In many cases, therefore, 'demand' is a subset of the intentional behaviour of wanting certain products and services.

b. **This behaviour occurs whenever the intentional behaviour, of which it is a subset, is undertaken.**

   When people want to buy things, there will be a level of 'demand' for these good and services. It is difficult to see how this does not happen as a matter of course.

c. **An actor is, initially, unaware of this derivatively intentional behaviour and has no beliefs, desires, plans, or other deliberations relating to the derivatively intentional behaviour.**

This is true when people just want to buy a variety of things. They have beliefs, desires, plans and other deliberations relating to the things they want, and prices they are prepared to pay for them, but usually, do not reflect on the demand for these goods. The obvious exceptions to this are cases where people fear a shortage, in which case they will pay attention to the level of demand, where these is a predictable connection between the level of demand and prices, or when (as above) they have views about demand for goods, either because they think there should be no demand, or because they think the level of demand should be higher (such as wanting to increase the level of demand for children's car seats).

d. **If an actor becomes aware of the derivatively intentional behaviour they realise that it is something they are doing. It is not something that is merely happening to them.**

This criterion works rather differently for 'demand' than for middle bias. This is because, in the case of 'demand', each person's contribution to the level of demand is likely to be negligible. So, it is something that they are 'doing' only in a very limited sense. Nevertheless, it is something that their actions form a part of.

e. **If an actor becomes aware of the derivatively intentional behaviour, it can be brought under the actor's control, if they choose to do so.**

Again, this is not the case in the current example. Only in exceptional circumstances can one person control the 'demand' for a good or service. However, they can form intentions about it, and try to influence it. For example, they could begin lobbying, or writing articles in the hope of influencing other consumers. They can also bring their, limited, contribution to the demand for a good under their control, for example by no longer buying it.

Is the concept of 'demand' derivatively intentional? It does not satisfy all the criteria because it is difficult to see how the demand for a good is something that people are doing. In the middle bias example biasing the middle is something that people are doing, even if they are the only person doing it. When economists are concerned with the demand for a good, they are usually concerned with the aggregate effect of many individual decisions. Consequently, demand is not something that can be brought under a person's control in any clear way. Nevertheless, it is true that the demand for a good is a phenomenon which arises, as a matter of course, out of the intentional behaviour of individuals, and about which they normally have no intentions.

The concept of 'demand' is less Nomadic than many other concepts, such as 'happiness' and 'welfare', but neither is it entirely precise either. Regularities using the concept 'demand' have *ceteris paribus* clauses, and many exceptions. While it would be incorrect to say that the level of demand for a product is derivatively intentional, it does meet some of the criteria for being derivatively intentional. This is an area worth researching further. Specifically, it is possible that when demand is derivatively intentional it behaves more in accordance with the laws of supply and demand. However, this is best assessed using empirical data and is an interesting project for further research.

3. Middle bias isn't really 'social' science.

One concern about Chapter 4 is that philosophers of social science have ambitions about generalisations in social science that go beyond formulating generalisations about biasing middle options. Instead, they hope to manage economies, or successfully foresee the impact of policy decisions. The middle bias example is no consolation, nor is it even 'social' in an obvious sense. The following section describes some other examples of behaviour that may be derivatively intentional which may alleviate this concern. These examples are intended to be purely illustrative at this stage and some of them may be shown to be mistaken. They are listed simply as interesting possible examples of derivatively intentional behaviour that are worth further research, and which may be more relevant to a more standard understanding of what social science is than the example of middle bias discussed above.

Other possible examples of derivatively intentional behaviour

Sentence structure in novels

This example closely matches the characteristics of middle bias. Drozdz *et al* (2015) discovered a surprising pattern in the variety of sentence lengths in a sample of international novels. They studied the sentence lengths in 113 English, French, German, Italian, Polish, Russian and Spanish books, and included books from a variety of time periods. Sentence length variability (SLV) was modelled and this points towards "power law long-range temporal correlations in SLV- thus to its fractal organisation- and indicates that it balances randomness and orderliness, just as it does for music, speech, heart rate, cognitions, spontaneous brain activity, and for other 'sounds of nature'" (2015, pg. 6). In other words, variability in sentence lengths in these books appears to conform not only to a regular pattern, but to the same regular pattern as other phenomena, such as music.

We might question whether the sample is large enough, and the extent to which the specific books chosen are representative of 'literature', or 'good writing'. Nevertheless, if we take the study at face value the fact that variability in sentence length is subject to regularities, is interesting. The authors of the books were engaged in intentional behaviour while writing, and they were also, presumably, paying attention to each sentence to ensure that it sounded right, or expressed their thoughts in the right way. It is, almost, certain that they did not have any clear intention that their sentence lengths should conform to this pattern. Indeed, they might find it surprising that it does. Reviewing each of the criteria for behaviour to be derivatively intentional in turn:

It is not entirely unintentional, because it forms part of a larger intentional behaviour; writing a novel. Where the authors to become aware of doing it, is something that would be doing. It is not something that simply happens to them; they are writing sentences with a specific structure. It is also something that they can bring under their control, either by checking that their new writing conforms to this pattern, or by ensuring that it does not. It is tied in with intentional behaviour, so it is not a side effect, as side effects are ordinarily understood because it is tied in with the intentional novel writing behaviour.

It is not intentional, because, following O'Shaughnessy, no conscious mental machinery is involved. Authors, and readers, are totally unaware of this phenomenon (until someone tells us about it). Placing this within the usual understanding of intentional action reviewed earlier, it forms no part of plans, it is not caused by any beliefs and desires, and we have no physical awareness of it.

This appears to be a potentially good example of derivatively intentional behaviour. Although it is, arguably, more interesting than middle bias, this example is also unlikely to satisfy the original motivations for finding regularities in the social sciences. However, the suggestion that sentence length variability conforms to the same pattern as other phenomena, such as music, does in my view, make this discovery important for understanding human behaviour. Clearly this is a matter about which opinions will differ.

Laws of internet searches

In 1998 Huberman *et al* proposed a 'law of surfing', which "determines the probability distribution of the search-depth that is, the number of pages a user visits within a web site." (1998, 95). They suggested that this formula could accurately predict page hits. The internet was in its infancy in 1998, but the belief in the existence of regularities in search patters has persisted. For example, Halvey *et al* (2006) analysed patterns in web surfing on mobile devices. They set out to discover whether, what they refer to as the "Universal Law of Web Surfing" (for which they cite Huberman *et al* as their source) applies to search patters when people are using mobile devices. They conclude that their results "confirm the generality of the Universal Law of Web Surfing" (2006, pg. 77).

The regularities they found "have been characterised by an inverse Gaussian distribution of surfing behaviour that helps determine the probability a user will click through a succession of pages (search to a given depth) in a surfing session" (Halvey *et al*, 2006, pg. 77). The data suggests that most internet users only select two or three links in a session and are unlikely to pursue searches further than this. They expected that differences between mobile and computer technology, and the difference in the quality of material available on mobile devices in 1996 would have led to differences in patterns of web surfing between computer and mobile based searches. However, the patterns originally derived from data on computer based internet searches applied to internet searches on mobile devices too.

The use of the description 'universal law' makes this example worth discussion, even though this description is undoubtedly inaccurate. The described pattern to internet searches is possibly another example of derivatively intentional behaviour. It is not entirely unintentional, because it forms part of a larger intentional behaviour; searching for information on the internet. Where people to become aware of doing it, it is something that they are doing. It is not something that simply happens to them; they are deciding whether to click on certain links or not. It is also something that they can bring under their control, either by consciously noting that their behaviour is regular, or by trying to change their behaviour (although it is unclear why people would want to change their behaviour in this instance). It is not a side effect because it is tied in with the intentional internet searching behaviour. However, this behaviour is not intentional, because, following O'Shaughnessy, no conscious mental machinery is involved. Exhibiting predictable patterns in website searching forms no part of plans, it is not caused by any beliefs and desires, and they have no physical awareness of it.

Understanding regularities in internet searching behaviour has a number of potential applications. It may help to understand in more detail how information is transmitted, and may indicate ways in which ideas can be communicated more effectively, both through website design as a whole, and by trying to increase the likelihood that people will increase their search depth. Huberman *et al's* original model, which is also the one used by Halvey *et al* derives the probability of an internet search clicking on a link by assuming that the decision to click on a link reflects the searcher's estimate of the probability that the next page will contain useful information (given the subject matter of many internet searches 'useful' and 'information' can presumably be very loosely understood). Understanding this in more detail might show how these probabilities can be changed.

Location predictability

The prevalence of mobile phones has enabled the collection of enormous amounts of data on people's locations over long periods of time. Song *et al* (2010) studied data on people's locations based on their mobile phone data and conclude that people's mobility patterns are highly predictable, and these patterns remain predictable even when people travel over large distances on a regular basis. Song *et al* analysed the mobility patterns of 50,000 people over a three-month period. They did not have access to real time, exact, locations but received information on which mobile phone masts the people were closest to whenever they made a call. After collecting data they sought to predict the location (which mobile phone mast they would be closest to) where a person would appear next. They concluded that location is 93% predictable.

However, things are not quite so simple as this. Predictability was, unsurprisingly, highest at night when most people are, predictably, at home. The least predictable time periods where when people were moving between locations, such as work and home. However, people did show a tendency to move predictably between locations. In other words, they do not often decide to walk a different way to work. Differences in age, gender, language groups, urban vs rural locations, did not affect predictability significantly. Importantly, the study did not find that people's locations were less predictable at the weekends, compared with weekdays. The authors conclude that "…regularity is not imposed by the work schedule, but potentially is intrinsic to human beings… despite this inherent population heterogeneity, the maximal predictability varies very little… and we see no users whose predictability would be under 80%" (Song *et al*, 2010, pg. 1021).

This behaviour also fits the criteria for derivatively intentional behaviour. Location predictability not entirely unintentional, because it forms part of a larger intentional behaviour. In this case, the intentional behaviour is not as well defined because at different times people's intentions will be different. For example, on a Monday morning many people may simply be intending to go to work, while on a Saturday night their intentions might be to go out to dinner, or to meet friends, or to start a fight.

Were people to become aware of the patterns generated by their behaviour, it would be something they are doing. It is not something that simply happens to them; they are going about their daily lives. It is also something that they can bring under their control, some people might be horrified that their seemingly spontaneous lives conform to this locational predictability and make a conscious effort to be less predictable. It is tied in with intentional behaviour, so it is not a side effect, as side effects are ordinarily understood, because it is tied in with the intentional behaviour.

It is not intentional, because, following O'Shaughnessy, no conscious mental machinery is involved. People are not aware of this phenomenon, until they find out about it. Placing this within the usual understanding of intentional action reviewed earlier, it forms no part of plans, it is not caused by any beliefs and desires, and we have no physical awareness of it.

This, therefore is another example of behaviour that may be derivatively intentional, with one caveat. In this case the intentional behaviour is more difficult to define, because there is no one activity in which people are involved, except if it is something extremely general like 'going about their daily lives'. Furthermore, the predictable behaviour is not tied in with one instance, or even one type of intentional behaviour. The people going about their daily lives are, in all likelihood, intentionally doing a number of very different things, including going to work, taking some exercise, and visiting friends. In this case we can understand the regular patterns in people's locations as derivative on many, intentional, decisions about where to go next. It is unclear whether this makes any significant difference to our understanding of this regular behaviour. This is an aspect of this type of regular behaviour that needs to be analysed further.

Song *et al* suggest that their research is relevant for urban planning, public health and traffic management. One potential use for this type for research is understanding patterns of disease diffusion and the spread of epidemics. However, Ferguson (2007) notes that analysing the possible spread of epidemics using travel and locations data from people's behaviour in normal times is potentially useless. He notes that, while people do not usually significantly change their behaviour when they have a cold, news of a pandemic is likely to make them avoid certain locations, and flock to others. This will make such models highly unreliable.

In conclusion, this section has tentatively proposed some examples of derivatively intentional behaviour that may be more relevant to social science research than middle bias. Even if these examples turn out not to be examples of genuinely regular behaviour they suggest the types of situations where regularities in behaviour may be found in the future. Middle bias is a disappointing example of a regularity in behaviour as far as standard social science research is concerned, but it is possible that other regularities may be discovered that are much more relevant for social science research; the examples above suggest areas where it is worth looking.

Chapter 6: Conclusion

The social sciences are often thought to be inferior to the natural sciences because they have not discovered many convincing examples of laws. This thesis has suggested a new reason why this is so, which is that the social sciences often deal with Nomadic concepts. This thesis has therefore proposed a new way of thinking about the concepts that social scientists use. The advantage of describing concepts as Nomadic is that it allows a number of methodological problems faced by the social sciences (extensive *ceteris paribus* conditions, reflexivity, and multiple realisability) to be subsumed within a common framework. Concepts are Nomadic when:

1- A wide variety of social phenomena can be included within the scope of the concept. This results from these concepts having many possible meanings, unclear boundaries, and changing over time. These characteristics are not an all or nothing matter because these concepts can vary in the number of meanings they have, how unclear their boundaries are, and the extent to which they change over time.

2- The characteristics outlined in criteria 1 mean that disagreements about Nomadic concepts, and arguments making use of them, are difficult to resolve with academic analysis. Over time, analysis of a Nomadic concept leads to the incorporation of different social phenomena.

There are two advantages of adopting the Nomadic framework for analysing concepts used in the social sciences. Firstly, viewing concepts in this way brings together a number of strands in the philosophy of social science literature. The concepts used by social scientists have been described in a number of ways, for example, that they are *Ballung* concepts, cluster concepts, essentially contested, multiply realisable, and vague. These descriptions are illustrated with different examples. To revisit some examples from Chapter 1, 'happiness' is an example of a *Ballung* concept, 'market economy' is an example of a cluster concept, 'democracy' is an example of an essentially contested concept, 'money' is an example of a concept that is multiply realisable, and 'bald' is an example of a concept that is vague. The problem with this is that these descriptions are of little use when trying to decide, for example, what sort of concept 'happiness' is. Furthermore, it is possible that 'democracy' is essentially contested, and multiply realised, and that it is a cluster concept, but it is unclear what we should conclude from this, or what implications this has for working with such concepts. The advantage of describing concepts as Nomadic is that it allows social scientists to discuss concepts using a single framework.

Nomadic concepts have many meanings, unclear boundaries, and change over time, which means that disagreements about these concepts are difficult to resolve. These characteristics allow the characterisation of the types of concepts social scientists often discuss. For example, *Ballung* concepts are those which have many meanings that do not all overlap at a single location. Cluster concepts have many meanings which do overlap at one or more locations. Essentially contested concepts have many meanings which change over time; these meanings may, or may not, overlap. Concepts that are multiply realisable have many meanings, some of which may overlap, it was argued in Chapter 2 that these meanings often do not have boundary issues.

Woodward (2016) argues in favour of a way of determining whether a concept is a good one for the purposes of analysis. Chapter 1 argued that his list is problematic for the social sciences because it relies on being able to judge whether a variable is able to do the work required of it. Within his manipulationist framework this means providing a clear answer to the question of what would happen when they are manipulated. Unfortunately, in the social sciences, this is difficult to know in advance of extensive research. Woodward says that 'obesity' is a bad variable, but does not say what it is about 'obesity' that makes it a bad variable. The Nomadic framework allows us to see why this might be so. There are a number of things that 'obesity' might mean, and although we might decide on one meaning- BMI, this will also have boundary issues. Depending on the research we want to do with 'obesity' these different meanings may have to be taken into consideration. The Nomadic framework allows social scientists to think about concepts in a structured way, and to make comparisons with other concepts.

The second advantage of adopting the Nomadic framework for analysing concepts used in the social sciences is that it brings together a number of methodological problems with the social sciences. The social sciences are subject to reflexivity, which describes the way behaviour changes as a result of interaction between people and the concepts that are used to analyse them. Chapter 2 argued that reflexivity is one way in which concepts can change. Changes of this sort might also give rise to new concepts describing similar phenomena. The existence of extensive *ceteris paribus* conditions should come as no surprise when social scientists are working with concepts that have many meanings, which may have boundary issues, and which change over time. In such a case, the concepts include shifting and changing aspects of social phenomena. Multiple realisability can also be seen as a different way of highlighting the many meanings that concepts can have.

Implications of adopting the Nomadic framework

There are three implications of adopting the Nomadic framework for analysing concepts used in the social sciences. These are difficulties with formulating generalisations, making concepts more precise, and counterfactual analysis. The first is that it demonstrates why generalisations have been so difficult to formulate in the social sciences. When we are dealing with concepts that are very Nomadic there is unlikely to be agreement about generalisations, or laws, using such concepts. Chapter 2 discussed the example 'income inequality makes revolutions more likely'. 'Income inequality', 'revolution' and 'more likely' are all Nomadic concepts. This means that anyone analysing this regularity can consider a number of meanings of these concepts and, consequently, reach different, but reasonable, conclusions about this statement. Chapter 2 contrasted this with a generalisation using less Nomadic concepts. The example was 'If the price of good x rises, then demand will fall'. The concepts of 'price', 'demand' and 'fall' are less Nomadic concepts because they do not have many meanings, or change over time. This means that social scientists analysing this relationship will focus on the same social phenomena. That is not to say that these concepts are entirely precise because they are not. Chapter 2 argued that these concepts are less Nomadic than concepts like 'income inequality'.

The second implication of adopting the Nomadic framework for analysing concepts used in the social sciences is that many proposals for making concepts more precise are problematic. Chapter 2 reviewed proposals from Crasnow (2015), Cartwright & Bradburn (2011), Goertz (2006, 2008) and Adcock & Collier (2001) which focus primarily on specifying what is mean by a concept, in a particular case. The example discussed at length attempted to make the concept 'democracy' more precise by stating that it means 'contestation' and 'participation'. Chapter 2 argued that 'contestation' and 'participation' are also Nomadic concepts, because there are many things these concepts can mean, which have boundary issues, and change over time. This was contrasted with a proposal to make the concept 'developed country' more precise by saying that it means a certain level of 'GDP per capita'. 'GDP per capita' is less Nomadic than 'contestation' and Chapter 2 concluded that simply seeking to specify what we mean by a concept is not sufficient for making it more precise. We also need to consider whether the meaning we have decided on is also a Nomadic concept.

The third implication of adopting the Nomadic framework is that is makes counterfactual analysis problematic in the social sciences. Woodward (2003) proposes that counterfactual analysis is like running a 'hypothetical experiment'. Chapter 3 argues that this view is difficult to endorse in the social sciences. Where social scientists are using counterfactuals which are stated using Nomadic concepts it is possible to construct multiple, plausible, counterfactual scenarios which there is little principled way of deciding between. Chapter 3 demonstrated this with a counterfactual that appeared to be a good candidate for a hypothetical experiment. This counterfactual asked whether the English Civil War would still have happened if Charles I had fought, and won, a battle with the Scottish rebels in June 1639. Adamson believed that the Civil War would not have happened, while Tucker believes that it would still have happened. This is a good candidate for a hypothetical experiment because the two armies faced each other and were prepared for battle. The backtracking that historical counterfactual often face is therefore minimised in this example. It is also clear what we mean by this counterfactual, and the manipulation required is small- Charles I just has to make a decision that he was already considering.

Despite this, although the two historians agree about the immediate outcome of the battle, they disagree about whether the Civil War would still have happened. They disagree because 'winning' is a Nomadic concept. Beyond the narrow military victory there are many things that Charles I's 'winning' could mean. It could mean that the opposition lost the will, or the courage, to stand up to the king again (if you are Adamson), or it could have meant that the pressure for change continued (if you are Tucker). The concept 'winning' has many meanings, and in assessing the consequences of this hypothetical experiment the historians can include a wide variety of social phenomena within the scope of this concept, which allows them to reach opposing conclusions. Chapter 3 concludes that where counterfactual analysis involves the use of Nomadic concepts unambiguous conclusions will be difficult to find.

The Nomadic framework and intentional behaviour

The social sciences are usually concerned with understanding intentional behaviour. Chapter 4 began with the observation that taking an intentional perspective on behaviour usually means explaining behaviour in terms of beliefs and desires. These beliefs and desires are often expressed using concepts that are Nomadic. This means that there will often be a number of plausible explanations of behaviour which it is difficult to decide between. One possible response to this is to seek to redescribe beliefs and desires in other terms. McIntyre's (1994) argument in support of this was reviewed, but found unconvincing. Despite this, there are a number of examples of behaviour that are regular, and predictable. Chapter 4 analysed the example of middle bias which, although an example of regular behaviour, is not intentional in an ordinary sense.

Chapter 4 proposed that this type of behaviour should be described as derivatively intentional. Derivatively intentional behaviour is defined by the following necessary, and jointly sufficient, conditions:

1- It is behaviour that is a subset of intentional behaviour.
2- This behaviour occurs whenever the intentional behaviour, of which it is a subset, is undertaken.
3- An actor is, initially, unaware of this derivatively intentional behaviour and has no beliefs, desires, plans, or other deliberations relating to the derivatively intentional behaviour.
4- If an actor becomes aware of the derivatively intentional behaviour they realise that it is something they are doing. It is not something that is merely happening to them.
5- If an actor becomes aware of the derivatively intentional behaviour, it can be brought under the actor's control, if they choose to do so.

Derivatively intentional behaviour is describable using concepts that are not Nomadic because there are not many things that derivatively intentional behaviour can mean. Furthermore, as the middle bias example demonstrates, what middle bias means is not subject to boundary issues, nor does it change over time. Chapter 4 concluded that derivatively intentional behaviour is the best example, so far, for formulating generalisations.

Chapter 5 discussed whether concepts that are less Nomadic are also derivatively intentional. Using the example of 'demand', from the law of supply and demand, it appears that this less Nomadic concept does share some characteristics with derivatively intentional behaviour, but that it does not meet all the criteria for being derivatively intentional. However, the law of supply and demand does also not describe behaviour that is as regular and predictable as middle bias. This suggests that exploring the link between Nomadic concepts and derivative intentionality is a possible area for further research.

Bibliography

Adamson, J. (1999) "*England without Cromwell: What if Charles I had avoided the Civil War?"* In 'Virtual History' edited by Niall Ferguson. Picador. 91-124

Adcock, R.; Collier, D. (2001) *"Measurement Validity: a Shared Standard for Qualitative and Quantitative Research."* American Political Science Review 95.3. 529-546

Alexandrova A. (2009) "*When Analytic Narratives Explain"* Journal of the Philosophy of History 3.1. 1–24

Anscome, G.E.M. (1957/ 2000) *Intention*. Harvard University Press.

Armstrong, D. M. (1983) *What is a Law of Nature?* Cambridge University Press

Attali, Y.; Bar-Hillel, M. (2003) "*Guess Where: The Position of Correct Answers in Multiple-Choice Test Items as a Psychometric Variable*", Journal of Educational Measurement Summer 40. 2. 109-128

Audi, R. (1999) *The Cambridge Dictionary of Philosophy* Cambridge University Press. Second Edition

Axford, N. (2010) "Is Social Exclusion a Useful Concept in Children's Services?" The British Journal of Social Work 40.3 737-754

Ayer, A.J. (1956/1998) *"What is a Law of nature?"* in 'Philosophy of Science- The Central Issues' edited by Martin Curd & J. A. Cover. Norton & Company. First Edition. 865-877

Bellomo, N.; Dogbe, C. (2011) *"On the modelling of traffic and crowds: A survey of models, speculation, and perspectives"* Society for Industrial and Applied Mathematics Review 53.3. 409-463

Bengtsson, B.; Ruonavaara, H. (2017) *"Comparative process tracing: making historical comparison structured and focused"* Philosophy of the Social Sciences 47.1. 44-66

Bergenholtz, C.; Busch, J. (2016) *"Self-fulfillemnt of social science theories: Cooling the fire"* Philosphy of the Social Sciences 46.1. 24-43

Bohman, J. (1994) *New Philosophy of Social Science- Problems of Indeterminacy.* Polity Press

Boyd, R. N. (1991) "*Realism, Anti-Foundationalism, and the Enthusiasm for Natural Kinds".* Philosophical Studies 61.1. 127-148

Boyd, R. N. (1999) "*Kinds, Complexity and Multiple Realization*". Philosophical Studies 95.1. 67-98

Boyd R. N. (2010) "*Homeostasis, Higher Taxa, and Monophyly".* Philosophy of Science 77.5. 686-701

Bradburn, N. (2011) *"Social science constructs".* In 'The Importance of Common Metrics for Advancing Social Science Theory and Research: A Workshop Summary' Washington DC: The National Academies Press. 53-70

Bratman, M (1979) *"Simple intention"* Philosophical Studies. 36.3. 245-259

Bratman, M. (1999) *Intention, Plans and Practical Reason*. CSLI Publications

Bratman, M. (1999i) *Faces of Intention: Selected Essays on Intention and Agency* Cambridge Studies in Philosophy

Brown, R. (1984) *The Nature of Social Laws- Machiavelli to Mill.* Cambridge University Press

Carney, P., Miller, V. (2016) *Vague Spaces* in "Strange spaces, Explorations into Mediated Obscurity" edited by Andre Jansson and Amanda Lagerkvist. Routledge. 33-55

Cartwright, N (1979) "*Causal Laws and Effective Strategies*" Noûs. Special Issue on Counterfactuals and Laws 13.4. 419-437

Cartwright, N. (1998) *"Do the laws of physics state the facts?"* in 'Philosophy of Science- The Central Issues' edited by Martin Curd & J. A. Cover. Norton & Company. First Edition. 865-877

Cartwright, N. (2005) *The Dappled World- A Study of the Boundaries of Science.* Cambridge University Press

Cartwright, N. (2007) *Hunting Causes and Using Them: Approaches in Philosophy and Economics.* Cambridge University Press

Cartwright, N.; Bradburn, N. (April 2011) *"A Theory of Measurement"*. The Importance of Common Metrics for Advancing Social Science Theory and Research: Proceedings of the National Research Council Committee on Common Metrics. 53-70

Cartwright, N.; Montuschi, E. (2014) *Philosophy of Social Science- A New Introduction*. Oxford University Press

Cartwright, N.; Runhardt, R. (2014*) "Measurement"* in 'Philosophy of Social Science- a New Introduction' edited by Nancy Cartwright & Eleonora Montuschi. Oxford University Press. First Edition. 265-287

Castaneda, H-N. (1982) *"Conditional intentions, intentional action and Aristotelian practical syllogisms"* Erkenntnis 18.2. 239-260

Castaneda, H-N. (1992) *"Indexical reference and bodily causal diagrams in intentional action"* Studia Logica 51.3. 439-462

Chisholm, R.M. (1970) *"The structure of intention"* Journal of Philosophy 67.19. 633-647

Collier, D.; Hidalgo, F. D.; Maciuceanu, A. O. (2006) *"Essentially contested concepts: Debates and applications"* Journal of Political Ideologies. 11.3. 211–246

Collingwood, R.. G. (1994) '*Human Nature and Human History'* in 'Readings in the Philosophy of Social Science', edited by Michael Martin and Lee McIntyre. MIT. First Edition. 163-172

Collingwood, R. G. (1946/1994) *The Idea of History: With Lectures 1926-1928* edited by J. Dussen. Oxford University Press. Revised Edition

Cooper, B.; Glaesser, J.; Gomm, R.; Hammersley, M. (2012) *Challenging the Qualitative- Quantitative Divide.* Continuum International Publishing

Crasnow, S. (2015) "*The Measure of Democracy: Coding in Political Science"* Forthcoming in 'Standardization in Measurement: Philosophical, Historical and Sociological Issues', edited by Lara Huber and Oliver Schlaudt. Routledge. Taylor & Francis. 149-160 In text page references refer to the pre-publication draft accessed on www.academia.edu

Davidson, D. (1963) *"Actions, Reasons, and Causes"* The Journal of Philosophy. 60.23. 685-700

Davidson, D. (1969) *"The individuation of events"*. In 'Essays in honor of Carl G. Hempel' Springer Netherlands 216-234

Davidson, D (2002) *Essays on Actions and Events* Oxford: Clarendon Press. Second Edition

Davidson, D (1963/ 2002) *"Actions, reasons and causes"* in D. Davidson 'Essays on Actions and Events' Oxford: Clarendon Press. Second Edition. 3-20

Davidson, D (1974/ 2002) *"Psychology as philosophy"* in D. Davidson 'Essays on Actions and Events' Oxford: Clarendon Press. Second Edition. 229-244

Davidson, D (1978/ 2002) *"Intending"* in D. Davidson 'Essays on Actions and Events' Oxford: Clarendon Press. Second Edition 83-102

Davis, W. (1984) *"A causal theory of intending"* American Philosophical Quarterly. 21.1. 43-54

Day, M. (2012) *The Philosophy of History*. Continuum International Publishing Group. London. Second Edition

Decock, L., Douven, I. (2014) "*What Is Graded Membership?*" Noûs. 48.4. 653-682

Della Sala, S.; Marchetti, C.; Spinnler, H. (1994) *"The anarchic hand: a fronto-mesial sign"*. In 'Handbook of Neuropsychology' edited by F. Boller and J. Grafman. Elsevir, 9. 233-255

Demetriou, C. (2009) *"The realist approach to explanatory mechanisms in social science"* Philosophy of the Social Sciences 39.3. 440-462

Douven, I (2016) "*Vagueness, graded membership and conceptual spaces"* Cognition 151. 80-95

Dray, W. (1957) *Laws and Explanation in History.* Oxford University Press

Dretske, F. I. (1998) '*Laws of Nature'* in 'Philosophy of Science- The Central Issues' edited by Martin Curd & J. A. Cover. Norton & Company. First Edition. 826-845

Droz, S.; Oswiecimkaa, P.; Kuliga, A,; Kwapien, J.; Bazarnikb, K.; Grabska-Gradzin, I.; Rybickib, J,; Stanuszekd, M. (2015) *"Quantifying origin and character of long-range correlations in narrative texts"*, Submitted to the Journal of Information Sciences

Duives, D. C.; Daamen, W.; Hoogendoorn, S. P. (2013) "*State-of-the-art crowd motion simulation models"*. Transportation Research Part C 37. 193-209

Earman, J.; Roberts, J. (1999) *"Ceteris paribus, there is no problem of provisos"* Synthese 118.3. 439-478

Education Endowment Foundation (2017) *EEF Projects: Magic Breakfast Executive Summary*. Accessed at https://educationendowmentfoundation.org.uk/our-work/projects/magic-breakfast

Education Endowment Foundation (2016) *EEF Projects: Magic Breakeakfast Evaluation Report and Executive Summary*. Accessed at https://educationendowmentfoundation.org.uk/our-work/projects/magic-breakfast

Ehrenberg, K. M. (2011) *"Law is not (best considered) an essentially contested concept"* International Journal of Law in Context 7.2. 209-232

Eklund, M. (2011) *"Recent work on vagueness"* Analysis Review. Oxford Studies in Metaphysics 71.2. 352-363

Ellis, B. *'Essentialism and Natural Kinds'*. (2008) 'The Routledge Companion to Philosophy of Science', edited by Stathis Psillos and Martin Curd. Taylor & Francis (Routledge). 139-148

Ereshefsky, M. (2004) "*Bridging the Gap between Human Kinds and Biological Kinds"*. Philosophy of Science 71.5. 912-921

Fay, B. (1994) *'General Laws and Explaining Human Behaviour'* in 'Readings in the Philosophy of Social Science' edited by Michael Martin and Lee McIntyre. MIT. 91-111.

Fearon, J. D. (1996*) 'Causes and Counterfactuals in Social Science: Exploring an analogy between cellular automata and historical processes'* in 'Counterfactual Thought Experiments in World Politics: Logical, Methodological and psychological Perspectives.' Edited by Philip E. Tetlock and Aaron Belkin. Princeton University Press 39-67

Ferguson, N. (1999) *Virtual History*. Picador. First Edition

Ferguson, N. (2007) *"Capturing human behaviour"* Nature 446. 733.

Frankfurt, H.G. (1978) *"The problem of action"* American Philosophical Quarterly 15. 2. 157-162

Freeden, M. (1996) *Ideologies and Political Theory: A Conceptual Approach.* Clarendon Press. 2006 edition

Gallie, W.B (1957) *"Essentially Contested Concepts"* Proceedings of the Aristotelian Society, New Series. 56. 167-198

Garance, F-R. (2013) "*All the Previous Declarations of War"*. The Atlantic. August. www.theatlantic.com

Garcia, J. L. A. (1990) *"The intentional and the intended"* Erkenntnis 33.2. 191-209

Giere, R. (1999) *Science Without Laws* University of Chicago Press

Goertz, G. (2006) "*Introduction to the Special Issue 'Causal Complexity and Qualitative Methods'*". Political Analysis. 14.3. 223-226

Goertz, G. (2006i) *Social Science Concepts: A User's Guide*. Princeton University Press

Goertz, G. (2008) "*Concepts, Theories, and Numbers: A checklist for constructing, evaluating, and using concepts or quantitative measures"* in 'The Oxford Handbook of Political Methodology' edited by Janet M. Box-Steffensmeier, Henry E. Brady and David Collier. Oxford University Press. 97-119

Gram, M. (2010) *"Self-reporting vs. observation: Some cautionary examples from parent/child food shopping behaviour"*. International Journal of Consumer Studies 34.4. 394-399

Grunbaum, T. (2007) *"The body in action"* Phenomenology and the Cognitive Sciences 7.2. 243-261

Grunbaum, T. (2010) *"Action and Agency"* in 'Handbook of Phenomenology and Cognitive Science' edited by S. Gallagher and D. Schmickinh. Springer Science. 337-354

Hacking, I. (1991) "*A Tradition of Natural Kinds"*. Philosophical Studies 61.1. 109-126

Hacking, I. (1995) "*Chapter 12 - The Looping Effects of Human Kinds"* in 'Casual Cognition: A Multidisciplinery Debate' edited by Dan Sperber, David Premack, and Ann J. Premack. Oxford Clarendon Press. 351-394

Hacking, I. (1998) *Rewriting the Soul: Multiple Personality and the Sciences of Memory*. Princeton University Press

Hacking I. (1999) *The Social Construction of What?* Cambridge, Harvard University Press

Hacking, I. (2004) *Historical Ontology*. Harvard University Press

Hacking, I. (2009) "*Humans, Aliens and Autism"*. Daedalus 138.3. 44-59

Halvey, M.; Keane, M. T.; Smyth, B. (2006) *"Mobile web surfing is the same as web surfing"* Communications of the ACM 49.3. 76-82

Hammersley, M. (2011) "*On Becker's Studies of Marijuana Use as an Example of Analytic Induction*" Philosophy of the Social Sciences 41.4. 535

Hausmann, D. M. (2008) *The Philosophy of Economics*, Cambridge University Press. Third Edition

Hawley, K.; Bird, A. (2011) "*What Are Natural Kinds?"* Philosophical Perspectives 25.1. 205-221

Hedström, P.; Ylikoski, P. (2010) *"Causal mechanisms in the social sciences*" Annual Review of Sociology 36. 49-67

Heilbron, J. (2008) *"Social Thought and Natural Science"*. In 'The Cambridge History of Science, Volume 7, The Modern Social Sciences' edited by Theodore Porter and Dorothy Ross. Cambridge University Press. Second Edition. 40-56

Helbing, D.; Farkas, I.; Vicsek, T. (2000) *"Simulating Dynamical Features of Escape Panic"*. Nature 407. 487–490

Hempel, C. G. (1942) *"The function of general laws in history"* The Journal of Philosophy 39.2. 35-48

Hempel, C. G. (1994) "*The Function of General Laws in History"* in 'Readings in the Philosophy of Social Science' edited by Michael Martin and Lee McIntyre. MIT. 43-53

Hempel, C. G. (1962/1998) *"Two basic types of scientific explanation"* in 'Philosophy of Science- The Central Issues' edited by Martin Curd & J. A. Cover. Norton & Company. First Edition. 685-694

Hoogendoorn, M. (2013) *"Predicting human behaviour in crowds: Cognitive modelling versus neural networks"* in 'Recent Trends in Applied Artificial Intelligence' edited by M. Ali *et al* Springer-Verlag Berlin. 73- 82

Howard, M. (2009) *War in European History* Oxford University Press. Second Edition

Huberman, B. A.; Pirolli, P. L. T.; Pitkow, J. E. Lukose, R. M.; (1998) *"Strong Regularities in World Wide Web Surfing"* Science. New Series. 280(5360) 95-97

IFS (2017) *Magic Breakfast: Evaluating school breakfast provision.* Accessed at https://educationendowmentfoundation.org.uk/our-work/projects/magic-breakfast

Jackson, P. T. (2011) *The Conduct of Enquiry in International Relations.* Routledge

Jones, A. (2001) *The Art of War in the Western World* University of Illinois Press

Katz, J. (2001) *"Analytic induction"* in 'International Encyclopedia of the Social & Behavioral Sciences' edited by Neil J. Smelser and Paul B. Bates. Elsevier, Amsterdam. 480-484

Khong, Y. F. (1996) *"Confronting Hitler and its consequences"* in 'Counterfactual Thought Experiments in World Politics: Logical, Methodological and psychological Perspectives.' Edited by Philip E. Tetlock and Aaron Belkin. Princeton University Press 95-118

Kieran, S. (2015) *"Intention"* in *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition) edited by Edward N. Zalta.

http://plato.stanford.edu/archives/sum2015/entries/intention/

Kim, J. (2009) *"Causation"* in 'The Cambridge Dictionary of Philosophy' edited by Robert Audi. Cambridge University Press. Second edition. 125-127

Kincaid, H. (1994) *"Defending laws in the social sciences"* in 'Readings in the Philosophy of Social Science' edited by Michael Martin and Lee McIntyre. MIT. 111-130

Kincaid, H. (2004) "*There are Laws in the Social Sciences*", in 'Contemporary Debates in the Philosophy of Science' edited by Christopher Hitchcock. Blackwell Publishing. First Edition. 168-186

King, G.; Keohane, R.; Verba, S. (1994) *Designing Social Enquiry.* Princeton University Press

Kurki, M. (2008) *Causation in International Relations.* Cambridge University Press

Lange, M. (2000) *Natural Laws in Scientific Practice.* Oxford University Press

Lange, M. (2005) *"Reply to Ellis and to Handfield on Essentialism, Laws and Counterfactuals".* Australasian Journal of Philosophy. 83.4. 581-588

Lange, M. (2009) *Laws and Lawmakers.* Oxford University Press

Lebow, R. N.; Stein, J. G. (1996) *"Back to the past: Counterfactuals and the Cuban Missile Crisis"* in 'Counterfactual Thought Experiments in World Politics' edited by Philip Tetlock and Aaron Belkin. Princeton University Press. 119-148

Lewis, D. (1973) *"Causation."* The Journal of Philosophy 70.17. 556-567

Lewis, D. (1979) *"Counterfactual Dependence and Time's Arrow"* Noûs 13.4. 455-476

Lewis, D. (2005) "*Counterfactuals"* Blackwell Publishing

Lewis, D. (2009) *On the Plurality of Worlds* Blackwell Publishing. Third Edition

Li, R. M. (2011) *The Importance of Common Metrics for Advancing Social Science Theory and Research: A Workshop Summary.* National Research Council. National Academies Press. Washington DC

Little, D. (1991) *Varieties of Social Explanation: An Introduction to the Philosophy of Social Science* Westview Press

Little, D. (1993) *"On the Scope and Limits of Generalisations in the Social Sciences"* Synthese 97.2. 183-207

Little, D. (2015) *"Mechanisms and Method"* Philosophy of the Social Sciences 45.4-5. 462-480

Lynch, K. (2012) "*A Multiple Realization Thesis for Natural Kinds."* European Journal of Philosophy 20.3. 389-406

Machlup, F. (1994) "*Are the Social Sciences Really Inferior?"* in 'Readings in the Philosophy of Social Science' edited by Michael Martin and Lee McIntyre. MIT. 5-19

Mäki, Uksali (2012) *"Realism and Antirealism about Economics"* in 'Handbook of the Philosophy of Science: Volume 13 Philosophy of Economics', edited by Dov M Gabbay, Paul Thagard and John Woods. Elsevier. First Edition. 3-24

Mallon, R. (2007) "*Human Categories Beyond Non-Essentialism*" The Journal of Political Philosophy 15.2. 146-168

Mason, K.; Easton, G.; Lenney, P. (2013) *"Causal social mechanisms; from the what to the why"* Industrial Marketing Management 42.3. 347-355

McCauley, J. L. (2004) *Dynamics of Markets- Econophysics and Finance*. Cambridge University Press

McGreer, V. (2009) "*The Thought and Talk of Individuals with Autism: Reflections on Ian Hacking".* Metaphilosophy 40.3-4. 517-530

McIntyre, L. C. (1996) *Laws and Explanation in the Social Sciences*. Westview Press

McIntyre, L. C, (1999) *"Davidson and social scientific laws"* Synthese 120.3. 375-394

McIntyre, L. C. (2000) *"Reduction, supervenience, and the autonomy of social scientific laws"* Theory and Decision 48.2. 101-122

McKnight, C. (2003) "Medicine as an essentially contested concept" Journal of Medical Ethics 29.4. 261-262

Mellor D.H. (1977) "*Natural Kinds"* The British Journal for the Philosophy of Science 28.4. 299-312

Merricks, T. (2001) "Varieties of vagueness" Philosophy of Phenomenological Research LXII.1. 145-157

De Mesquita B. B. (1996) *"Counterfactuals and international affairs: Some insights from game theory"* in 'Counterfactual Thought Experiments in World Politics: Logical, Methodological and psychological Perspectives.' Edited by Philip E. Tetlock and Aaron Belkin. Princeton University Press 211-229

Millikan, R.G. (1999) "*Historical Kinds and the "Special Sciences""* Philosophical Studies 95.1-2. 45- 65

Moaz, Z.; Russett, B. (1993). *"Normative and structural causes of democratic peace 1946-1986"* The American Political Science Review 87.3. 624-638

Norkus, Z. (2005) *"Mechanisms as miracle makers? The rise and inconsistencies of the "Mechanismic Approach" in social science and history"* History and Theory 44.3. 348-372

Northcott, R. (2008) *"Causation and contrast classes"* Philosophical Studies 139.1. 111-123

Odenbaugh, J.; Alexandrova, A. (2011) "*Buyer beware: robustness analyses in economics and biology"* Biology and Philosophy 26.5. 757–771

O'Shaughnessy (1980) *The Will, a Dual Aspect Theory* Volume 2. Cambridge University Press

Parsons, K.P. (1973) "*Three Concepts of Clusters"* Philosophy and Phenomenological Research 33.4. 514-523

Persky, J. (1990) *"Retrospectives: Ceteris paribus"* Journal of Economic Perspectives 4.2. 187-193

Pigliucci, M. (2005) "*Wittgenstein Solves (Posthumously) the Species Problem"* Philosophy Now. March/April. 50

Pitcher, G. (1970) *"In Intending' and side effects"* Journal of Philosophy 67.19 659-668

Popper, K. (1961) *The Poverty of Historicism.* New York, Harper Torchbooks

Porter, T. (2008) "*Genres and Objects of Social Enquiry, from the Enlightenment to 1890".* In 'The Cambridge History of Science, Volume 7, The Modern Social Sciences' edited by Theodore Porter & Dorothy Ross. Cambridge University Press. 13-39

Ragin, C.C. (2008) *Redesigning Social Enquiry: Fuzzy Sets and Beyond.* University of Chicago Press

Reiss, J. (2009) *"Counterfactuals, Thought Experiments, and Singular Causal Analysis in History."* Philosophy of Science 76.5. 712-723

Richardson, L.; Le Grand, J. (2002) *"Outsider and Insider Expertise: The response of residents of deprived neighbourhoods to an academic definition of social exclusion"* Social Policy & Administration 36.5 496-515

Roberts, J., T. (2004) *"There are no laws of the social sciences"* in 'Contemporary Debates in the Philosophy of Science' edited by Christopher Hitchcock. Blackwell Publishing. 151- 167

Robinson, G. (2007) "*Nature and Natural Kinds"* Philosophy 82.4. 605-623

Rodway, P.; Schepman, A..; Thoma, V. (2016) *"Reachability does not explain the middle preference: A comment of Bar-Hillel"* i-Perception. 7.2 1-5

Rosenberg, A. R. (1995) *Philosophy of Social Science.* 2nd Edition. Westview Press

Rosenberg, A. R. (2001) "*How is biological explanation possible?"* British Journal of Philosophy of Science 52.4. 735-760

Rosenberg, A. R. (2012) *"Why do spatiotemporally restricted regularities explain in the social sciences"* British Journal for the Philosophy of Science 63.1. 1-26

Salmon, M. H. *et al* (1999) *Introduction to the Philosophy of Science.* Hacking Publishing Company, Indianapolis/ Cambridge

Sambanis, N. (2004) "*What is Civil War?*" Journal of Conflict Resolution 48.6. 814-858

Scriven, M. (1964) *"Views of Human Nature"* in 'Behaviourism and Phenomenology: Contrasting Bases for Modern Psychology' edited by T. Wann. University of Chicago Press

Scriven, M. (1994) *"A Possible Distinction Between Traditional Scientific Disciplines and the Study of Human Behaviour"* in 'Readings in the Philosophy of Social Science' edited by Michael Martin and Lee McIntyre. MIT. 71-77.

Sen, A. (2004) *"Capability and well-being",* in Nussbaum, M. 'The Quality of Life', New York: Routledge. 30-53

Setiya, K. (2015) *"Intention",* The Stanford Encyclopedia of Philosophy (Summer 2015 Edition) edited by Edward N. Zalta

http://plato.stanford.edu/archives/sum2015/entries/intention/

Song, C.; Qu, Z.; Blumm, N.; Barabási A-L. (2010) *"Limits of Predictability in Human Mobility"* Science, New Series 327.5968. 1018-1021

Staley, K. W. (2014) *An introduction to the Philosophy of Science* Cambridge University Press. First Edition.

Steel, D.P (2008) A*cross the Boundaries: Extrapolation in Biology and Social Science*. Oxford University Press

Steuer, M. D. (2003) *The Scientific Study of Society*. Kluwer Academic

Taylor, C. (1971/ 1994) *"Interpretation and the Sciences of Man."* In 'Readings in the Philosophy of Social Science', editd by Michael Martin and Lee McIntyre. Massachusetts: MIT. 181-212

Tetlock, P. E.; Belkin, A. (1996) *Counterfactual Thought Experiments in World Politics: Logical, Methodological and psychological Perspectives.* Princeton University Press

Tsou, J.Y. (2007) "*Hacking on the Looping Effects of Psychiatric Classifications: What is an Interactive and Indifferent Kind?*" International Studies in the Philosophy of Science 21.3. 329-344

Tucker, A. (2009) *Our Knowledge of the Past: A Philosophy of Historiography*. Cambridge University Press. Second Edition. Originally published 2004

UK Government (2014) Child Poverty: a draft strategy. https://www.gov.uk/government/consultations/child-poverty-a-draft-strategy

van Fraassen, B. C. (1980) *The Scientific Image* Oxford. The Clarendon Press

Weber, Max. (1978) *Economy and Society: An Outline of Interpretive Sociology*. University of California Press.

Williams, M. (1980) *"Similarities and differences between evolutionary theory and the theories of physics"* Proceeding of the Biennial Meeting of the Philosophy of Science Association. 1980.2. 385-396

Wittgenstein, L. (2009) *Philosophical Investigations* Wiley Blackwell. Fourth Edition

Woodward, J. (2003) *Making Things Happen.* Oxford University Press

Woodward, J. (2011) *"Mechanisms revisited"* Synthese. Neuroscience and its Philosophy 183.3. 409-427

Woodward, J. (2015) *"Methodology, ontology, and interventionism*" Synthese. 192. 3577-3599

Woodward, J. (2016) *"The problem of variable choice"* Synthese 193. 1047-1072

Wright, F. (2011) *"The more things change, the more they stay the same: Criminal law, Downs Syndrome, and a life worth living"* Law, Crime and History 1. 62-85