



Reflexivity and fragility

Robert Northcott¹ 

Received: 29 November 2021 / Accepted: 12 June 2022 / Published online: 25 June 2022
© Springer Nature B.V. 2022

Abstract

Reflexivity is, roughly, when studying or theorising about a target itself influences that target. Fragility is, roughly, when causal or other relations are hard to predict, holding only intermittently or fleetingly. Which is more important, methodologically? By going systematically through cases that do and do not feature each of them, I conclude that it is fragility that matters, not reflexivity. In this light, I interpret and extend the claims made about reflexivity in a recent paper by Jessica Laimann (2020). I finish by assessing the benefits and costs of a focus on reflexivity.

Keywords Reflexivity · Fragility · Master model · Contextual · Prediction · Laimann

1 Introduction

What influences whether we can predict, explain, and intervene successfully in the human sciences? I compare two potentially relevant phenomena. The first, the topic of this special issue, is *reflexivity* – which is, roughly, when theorising about or studying a target itself influences that target. (More on reflexivity’s definition later.) The second is *fragility* – which is, roughly, when causal and other relations are hard to predict, holding only intermittently or fleetingly. (More on fragility’s definition shortly.) I argue, in a nutshell, that it is fragility that matters, not reflexivity. While reflexivity does sometimes have an impact on optimal methodology, usually fragility has a much bigger one.

In section 2, I introduce the notion of fragility and outline its methodological consequences. In section 3, I explain why, methodologically speaking, fragility matters more than reflexivity, by going systematically through cases that do and do not feature

This article belongs to the Topical Collection: Reactivity in the Human Sciences
Guest Editors: Marion Godman, Caterina Marchionni, Julie Zahle

✉ Robert Northcott
r.northcott@bbk.ac.uk

¹ Department of Philosophy, Birkbeck, University of London, Malet Street, London WC1E 7HX, UK

each of them. In section 4, I use this to build on recent work on reflexivity by Jessica Laimann. Finally, in section 5, I return to reflexivity itself, and in particular to how a focus on it both helps and hinders scientific progress.

2 Fragility

Imagine two worlds. In one, there is underlying order. Causal relations are stable and long-lasting; mechanisms, structures and functional dependencies persist across many cases; laws are unchanging. How best to investigate such a world? By uncovering these cogs and wheels of nature, confident that they will work widely. Knowledge of them, accumulated over many generations, is the route to remarkable power. Technological artefacts can work reliably, by being engineered to combine and exploit cogs and wheels without disruption. To explain an event is a matter of identifying some configuration of these cogs and wheels, and perhaps thereby of seeing how the same cogs and wheels underlie many other, superficially disparate, events too. The science in such a world is one familiar from textbooks and popular image. Beneath the messy imperfection around us stands a Platonic order, and finding this order is science's mission.

Now imagine a second world, this time one in which laws and causal relations are fragile, winking in and out like bubbles in a boiling soup. In this world, things are different. Just because one thing causes another over there, that does not mean it will cause it over here. Working hard to discover underlying cogs and wheels is no longer an efficient use of our energies. When explaining things, we are forced knee-deep into idiosyncratic local detail; no eternal laws, rather each time a new look. Artefacts lose their power because the building blocks on which they rely are fragile. Progress is still possible in this fallen world, and remains vital for our fortunes, but it is piecemeal and patchwork.

Our world is an interlocking mixture of these two worlds. But many target relations are fragile and this needs to be reckoned with, as we will see Northcott (2022).

Speaking less roughly, call a relation *fragile* if, in the salient circumstances, it is not predictable. (Reciprocally, to be *stable* is to be non-fragile.) Predictability is a matter of degree, and therefore so is fragility. Fragility can arise because a relation holds only locally, or intermittently, or with inconsistent strength – and because these variations cannot easily be predicted.

Predictability has both a subjective and objective aspect, and therefore so does fragility. Predictability is relative to our knowledge: in a deterministic world, for example, nothing is hard to predict for Laplace's Demon, even while, of course, many things remain hard to predict for us humans. But it is also true that some relations are harder to predict than others for reasons external to us. The difference between chaotic and non-chaotic systems is one example. Relations may also be fragile because they rarely operate in isolation, or they require advanced levels of knowledge to identify, or they are difficult to observe. Whatever the cause, in many cases it is not realistic to render something predictable just by increasing our knowledge. In those cases, in practice, we must take a relation's fragility to be a given.

Fragility is not the same as *complexity*, although complexity often leads to it. Fragility arises in many non-complex systems, while some relations even in complex systems are not fragile: summers are predictably warmer than winters, for example, even though the weather system is a paradigm of complexity.

2.1 The Master-Model strategy

There are further nuances, but for our purposes the above understanding of fragility will be precise enough. Turn to our main focus, which is what fragility implies for methodology. It is useful to begin with a simplified, benchmark case of a *non-fragile* relation. Suppose we want to predict the motion of a newly discovered moon. To do so, we apply a Newtonian two-body model of gravity, inputting the moon and parent planet's masses, positions and motions. Something like this procedure is a staple of actual space exploration. Why does it work? The answer is stability: the Newtonian model that has been successful elsewhere can be assumed to apply to a new case, because gravity itself can be assumed still to be operating in the same way. Each time, just re-apply the same Newtonian master model. Call this the *Master-Model* strategy.

A major advantage of Master-Model is that it is still effective even in the face of *noise* – understood here as significant effects from disturbing factors that are not captured by our model. For example, the moon's motion may be deflected by gravity from a second moon, by impact with a comet, or (at least for a small moon maybe) by human interference. If so, because of these disturbing factors, the Newtonian two-body master model would no longer predict accurately. Nevertheless, the model would still reliably identify *one* of the factors influencing the moon's motion, namely the gravitational interaction between moon and planet. In this sense, the model would still explain 'partially' (Northcott, 2013). To explain fully, or predict accurately, we would have to add in the effect of unmodeled disturbing factors. This strategy – of developing a master model and then in specific cases adding in disturbing factors as needed – was already advocated by Mill almost two centuries ago (1843). It can work well when we have stability.

In this way, a master model provides some understanding even in the many cases when empirical accuracy is imperfect. Such an achievement, according to this view, is actually *superior* to mere empirical accuracy. Why? Because empirical accuracy in any particular case requires taking account of every local factor, no matter how *sui generis* or transient. But what is of greater interest to science, as a pursuit of systematic knowledge, is those factors that generalize – which is just what a master model captures.

Master-Model relies on stability in two ways. First, stability is essential metaphysically. A master model serves as a reliable base onto which case-specific disturbing factors can be added, only because the relations it describes are stable. (Mill himself was well aware of this: he had in mind economics, where he thought core psychological tendencies such as seeking to increase one's own wealth are indeed stable in the required way.) Second, stability is also essential epistemologically. In easy cases, warrant to apply a master model comes from empirical success here and now: the Newtonian gravity model, for example, is given warrant by successfully predicting the motion of the moon. But often there is no empirical success here and now, because of noise. Then, warrant can come only indirectly, by importing empirical success from elsewhere. For example, even when noise means it predicts badly here and now, still we are justified in thinking the Newtonian model has correctly identified one gravitational force at work. Why? Because of the model's empirical success elsewhere. But such indirect warrant is justified only when there is stability. In this example, it is only because gravity operates in the same way across cases that the Newtonian model's warrant from success elsewhere stays good over here.

In sum, the crucial thing for Master-Model is stability. Noise does not tell against it – but only given stability. Without noise, a master model is empirically accurate across many cases only when the relations it describes are stable. With noise, meanwhile, while a master model is no longer empirically accurate, now we may retreat to Mill's strategy, confident that a master model does at least capture some of the factors present, even if additional disturbing factors are present too.

2.2 The Contextual strategy

But there is an alternative approach – for when we face fragility. To introduce this alternative, imagine now a different moon example. This time, the 'moon' in question is a toy moon on a string, being carried by a child around a toy planet. How might we predict the motion of *this* moon? The best candidate here for a master model is something psychological, perhaps that children continue doing actions that they are enjoying. Call this the Continue model. It predicts that, if the child has carried the toy moon around the planet for two 'orbits' happily, they will continue for at least another two orbits. This prediction will be right sometimes. Other times, it will not be: perhaps the child gets distracted, interrupted, or bored, or perhaps they are following instructions in an online science class (two orbits only), or perhaps they are playing a game with a friend (take it in turns to hold the moon). The underlying problem is fragility: the relation captured by the Continue model does not hold reliably. Using just a single model is no longer effective.

What alternative strategy works better? Many different models are available, some relatively formal, others that we might think of more as loose hypotheses or rules of thumb. Some models cover a child's behaviour in each of the deviant scenarios above: when distracted, bored, in a school class, with a friend, and so on. Others cover plenty of further scenarios: when a child is tired, when they are interacting with a sibling, when they are affected by poverty or divorce or moving to a new house, and so on. The key is which of these many models applies in any particular case. To discover that requires much case-specific work, looking for contextual clues and triggers: the character of this child, the nature of this household, is the child tired late in the day, is the child hungry, is the child – or friend or parent – generally frustrated after a prolonged lockdown, is the weather bright and warm or is it grey and miserable, and so on. In short, exactly the things a parent considers when trying to understand and predict a child's behaviour. Instead of a single master model, we choose from many different models, case by case.

Label this new methodological strategy, *Contextual*. Contextual is not against wide-scope models as such; on the contrary, the larger the available toolbox of such models, the better. Rather, Contextual implies two things. First, a change in balance: relatively more scientific effort must be devoted to local empirical investigation. Because no single model may be assumed always to apply, we must be more sensitive to the details of each case, in order to select wisely from our toolbox. Second, a change in how models are developed. They should not be developed a priori or in the abstract, relying on real-world stability to ensure that if they capture a relation in one place, they will therefore capture it elsewhere too. Instead, models must be constantly empirically refined, in turn by constantly applying them to real-world cases. Such constant refinement is the best way to make models empirically productive, and to learn in what circumstances they are likely to apply (Ylikoski, 2019).

When relations are fragile, even though relevant *models* may be putatively wide-scope, the *explanations* derived from those models are typically narrow-scope. In other words, these explanations cover only one or a few situations rather than, like the Newtonian gravitational model, many. Why? First, because relations at the heart of an explanation will, if fragile, often not hold widely. Second, because, as noted, the warrant of predictive success cannot be imported from elsewhere, which means that accurate prediction is always required here and now. This forces an explanation to include all local causes, no matter how *sui generis*. In turn, this usually requires going beyond just those factors captured by a wide-scope model and instead delving into local details, in the manner of a historian. This localist picture dovetails with contemporary theories of causal explanation, such as Woodward's (2003), which are not framed in terms of general laws but rather in terms of invariance relations that may be of very limited scope. It also dovetails with much recent philosophy of science, which emphasizes the need for local work to know when and how we may apply our models (Cartwright, 2019).

To clarify the difference between Master-Model and Contextual: in both cases, use is made of models that may apply to many situations. And in both cases, contextual work is required to estimate parameter values. The difference lies elsewhere. In the moon example, the moon's position and velocity vary continuously. The Newtonian gravity model tells us not just the details of that variation, but also when to expect it. Master-Model works well. There is no 'surprise' variation in gravity's influence *that requires knowledge from beyond the model* to predict: the inverse-square law itself does not vary unpredictably, and neither is it difficult to predict when gravity will be present. But in fragile cases, we get just such surprise variation. So, we further need to investigate each time whether a model applies – whether its relations have changed their forms, and indeed whether its relations any longer hold at all. That is, we need Contextual.

Sometimes a model applies, even if only partially, in many different places. Discovering a new mechanism can therefore be valuable, because it enlarges our toolbox of available models. In this way, there is still scope for context-general scientific achievement. But if a relation captured by a model is fragile, then we may never just assume that that model applies; that still needs to be established anew each time. We are no longer in Newton's world, so to speak.

With fragility, explanatory warrant requires accurate prediction here and now, as noted. But accurate prediction is now harder: the toy moon's motion is harder to predict than is the real moon's. It is difficult to know which model applies, and noise is ubiquitous. But that, as it were, is nature's fault, not ours. Still, this is not a counsel of despair: we can get a decent grip on the toy moon's motion sometimes, some predictions are more accurate than others, and some explanations are fuller and better warranted. It is up to us to find them.

Summing up: what matters methodologically is fragility. When target relations are stable, it is best to investigate via a single master model (assuming an accurate one can be found), such as a Newtonian model of gravity. This strategy is effective even when empirical accuracy is disrupted by noise. But when target relations are fragile, matters change. A shift of emphasis is required – towards contextual, historian-like sifting. Local investigation is required each time to discover which of many candidate models might apply, and any model selected needs to be empirically accurate here and now.

3 Reflexivity versus fragility

Turn now to reflexivity. Several definitions of reflexivity, and of related (or, as sometimes used, synonymous) notions such as reactivity and performativity, have been offered. Ian Hacking (1995) made famous the notion of unstable kinds. Reflexivity means that human kinds (i.e., kinds concerning humans) are potentially altered by feedback effects, with each alteration of a kind potentially inducing reactions in the target that in turn potentially feed back into a further alteration of the kind, and so on indefinitely. But reflexivity has also been defined, more simply, as unintended effects by theorising on the objects of study, with unstable kinds being merely one possible side-effect of that. Other definitions have been offered too. Below, I articulate reflexivity in terms of unstable kinds, but for our purposes nothing important turns on that.

Reflexivity is clearly distinct from fragility, as the examples below show. As it were, with fragility the bubbles in the boiling soup come and go, whereas with reflexivity the scientist (perhaps unwittingly) dips their spoon into the soup and actively stirs it.

To assess the relative methodological significance of reflexivity and fragility, I will work through the different combinations of the two, using the following 2×2 table of examples. For each example, I report detailed case studies already carried out by others.

Begin in the bottom-left corner, with reflexivity but not fragility. Start with domestic dogs (Khalidi, 2010). The kind ‘dog’ is reflexive: as Khalidi explains, with repeated rounds of breeding over perhaps 15,000 years, this kind has changed dramatically, both morphologically and behaviourally. Traits including tameness, obedience, teachability, shepherding, hunting, and certain physical characteristics, have been selected for, so that what were originally wild wolves became domestic dogs and then later the many different breeds of domestic dog today. Human interaction with the kind changed it. There were many rounds of Hacking-style looping effects as new kinds themselves stimulated new human breeding behaviours, which in turn fed back to change the kinds once more. The key relation underpinning this history is that between breeding and the evolution of dog traits. This relation is stable: we can (for the most part) reliably predict the (rough) impacts of breeding interventions, and this reliability is exploited by dog breeders all the time. Further, we can reliably explain these impacts of breeding by appealing to general evolutionary theory – one of nature’s cogs and wheels. Master-Model works well here for all of predicting, intervening, and explaining. This is despite reflexivity, and because of stability.

The example of gender roles (Laimann, 2020) teaches the same lesson: optimal method tracks fragility (or its lack), not reflexivity. Briefly, the kinds ‘masculine’ and ‘feminine’ are reflexive, being strongly influenced by how people conceive of them. But many relations involving them are stable. (Indeed, Mallon (2016) argues that reflexivity has in this case been a stabilizing force, pushing the kinds back into line, so to speak, if they show signs of changing.) As a result, we can reliably predict the impacts of many interventions, and can reliably explain them, by appeal to wide-scope theories – in this case, sociological theories of gender roles. Because the relevant relations are stable, Master-Model is successful.

Turn next to the upper-left corner of Table 1, which features both reflexivity and fragility.

Schizophrenia and autism are two of Hacking’s own examples. Begin with schizophrenia. According to Hacking (1999, 112–14), because of reflexivity schizophrenia has changed its properties several times. At the start of the twentieth century, when schizophrenia was first named, by Eugen Bleuler, its main symptom was flat affect. Auditory hallucinations (i.e., hearing voices), by contrast, were considered a minor

Table 1 Reflexivity versus fragility

	Reflexive kinds	Non-Reflexive kinds
Fragile relations	Contextual Schizophrenia Autism	Contextual Invasive species World War One truces
Stable relations	Master-Model Domestic dogs Gender roles	Master-Model Newtonian theory Electric toothbrush

issue, not specific to schizophrenia but rather observed in many other psychiatric conditions too. They were not to be worried about, and not something to hide from the doctor. The result was that hallucinations became increasingly widely reported by patients, and by the time a formal list of 12 symptoms of schizophrenia was compiled by Kurt Schneider 30 years later, the kind had changed, with hallucinations being designated the main symptom. But then, after the war, schizophrenia evolved from something viewed indifferently even favourably, to become instead a diagnosis that people wanted to avoid. As a result, patients became less willing to report hallucinations. This led the definition of the kind to be changed again, as hallucinations were gradually de-emphasized once more as a diagnostic criterion (although they are still listed as one of the main symptoms). Schizophrenia is, thus, according to this account, an unstable kind because of reflexivity effects.

How will the schizophrenia kind change next? That is hard to predict, because the social and political relations surrounding schizophrenia are hard to predict. In other words, these relations are fragile. If they were not fragile then, like with the breeding of domestic dogs, Master-Model could get us the answers we want. But there is no such master model available that will predict the future of the schizophrenia kind reliably – no equivalent to evolutionary theory. As a result, in-depth local investigation is required instead, such as the history that Hacking recounts in order to explain the various changes in schizophrenia's definition. Prediction, intervention and explanation are not straightforward in this case. Fragility explains why. Fragility, not reflexivity, is the difference between the schizophrenia and domestic dogs examples.

Similar remarks apply to autism (Hacking, 1995). First named in 1938, this kind has subsequently varied greatly in its definition, as well as in theories of what causes it, and in its degree of stigma. Reflexivity effects are an important part of the story, according to Hacking. The need for detailed local investigation, both to explain the kind's history and to predict its future evolution, are again a result of the fragility of the many social relations swirling in the background.

Turn now to the two right-hand boxes. Being cases of non-reflexive kinds, they would not typically feature in discussions of reflexivity. But optimal methodology differs between them, and this difference is revealing, because only fragility is varying between the cases, not reflexivity.

In the upper-right corner, there is fragility but not reflexivity. Consider invasive species, as reported by Alkistis Elliott-Graves (2016, see also 2018, 2019). The relevant kinds here are things like tree species, soil nutrients, islands, and lakes. In the context of species invasions, none of these kinds is unstable or reflexive. As Elliott-Graves

recounts, despite knowing several mechanisms behind invasions and their harmful effects, we cannot predict reliably the outcome or scope of invasion events. The relevant relations are fragile.

Consider one of Elliott-Graves's examples: plant-soil interaction. There exist both positive (certain fungi, and nitrogen fixers) and negative (pathogenic microbes) potential feedbacks to plants from the soil. Evolutionary interaction tends to favour the negative feedbacks, with the result that plants will tend to be favoured by going to a new area. Does this pattern enable us to predict the fate of plant invasions? Alas, no. How quickly plants accumulate pathogens is critical to an invasion's success, and this in turn varies with several further, local factors, such as the relative abundance of invaders and native plants, and the predation climate. As a result, prediction is difficult. Even when a combination of invader and soil-microbes seems perfect for an invasion to succeed, often an invasion nonetheless fails. The relation between plant-soil set-up and invasion success is fragile. Plant-soil feedback interactions have been modelled extensively, and the relative abundance within a community of all-native plants has been predicted successfully. But when it comes to invasions of new communities, predictions are no longer reliable.

Many rules of thumb explain, or partially explain, invasions sometimes. Examples include: that islands, especially small ones, are more susceptible to invasions than are mainlands; that temperate climates are more susceptible than the tropics; and that within a taxon, smaller animals are more invasive than larger animals. Within plant taxa, the following traits correlate with successful invasions: small seed size; phenotypic plasticity; allelopathy, i.e., producing biochemicals that impact on the success of other organisms; adaptation to fire; and, at different times in different places, small and large size, flowering early and late, and both dormancy and non-dormancy. And the following traits of communities are all correlated with being easy to invade: when humans facilitate the invasion (perhaps inadvertently); when the community is disturbed; lack of biological inertia (i.e., the ecological balance can change relatively easily); particular plant-soil feedbacks, as just discussed; when the supply of resources fluctuates; and, depending on context, both high and low diversity. Similar lists can be compiled for marine ecosystems, insects, vertebrates, and so on.

But none of these many rules of thumb explains invasions reliably, and none predicts them reliably either. They are like the various models in the toy-moon example: they are all fragile. For any given invasion, extensive case-specific work is required to work out which rule of thumb was relevant in the particular case. Master-Model does not work.

Note two other features of the invasive species case. First, epistemic difficulty occurs despite the lack of reflexivity. Second, the case illustrates how fragility is not restricted to human sciences. (More on the scope of fragility in Section 4.)

The World War One truces are another example of fragility without reflexivity, this time taken from the human sciences. Briefly: truces broke out spontaneously in many parts of the Western Front, despite constant pressure against them from senior commanders. What explains this remarkable and moving phenomenon? According to (Northcott & Alexandrova, 2015), the Master-Model strategy, represented in this case by the Prisoner's Dilemma game, is not successful, because the historical details turn out to contradict the Prisoner's Dilemma account in many ways. Indeed, the Prisoner's Dilemma actively directs attention away from the factors that are actually significant,

and that bear on other instances of co-operation. The only way to successfully explain the truces is by contextual work, as exemplified by investigations by historians. The relevant relations are fragile. And the kinds here (war, soldiers, truces) are not reflexive, at least not in the context of this case.

In the lower-right corner, finally, there is neither fragility nor reactivity. Without fragility, Master-Model can again succeed: capturing the cogs and wheels of nature is again the proven route to prediction, intervention, and explanation. A paradigm case is Newtonian theory. This describes relations that are non-fragile – indeed universal – and succeeds dramatically. It is a similar story with technological artefacts, which are deliberately engineered to exploit relations that, in a suitably shielded environment, are stable.

Newtonian theory's target kinds are typically non-reflexive, as are those of technological artefacts. But, as with invasive species and with World War One truces, when it comes to choosing between Master-Model and Contextual, it does not matter whether kinds are reflexive or not. What matters is whether target relations are fragile.

4 Laimann and beyond

Jessica Laimann's (2020) discussion of reflexivity shares many of the above emphases. I discuss now how an analysis in terms of fragility complements and adds to it.

According to Laimann, our concern with reflexivity is ultimately epistemic and methodological. Unstable kinds in themselves are not necessarily a problem; rather, what matters epistemically are the processes and mechanisms behind that instability. How well do we understand those? When we do understand them well, as with domestic dogs and with gender roles, we are able to predict, intervene, and explain satisfactorily.

Laimann states: "Only when we *understand the mechanisms* that support patterns of change and stability among the members of a kind are we in a position to *provide accurate explanations and make inductive inferences* across a variety of contexts." (2020, 1056, italics added) The notion of fragility adds a metaphysical underpinning to the italicised phrases. What matters is whether the background relations are fragile. If they are not fragile, then, by definition, we will 'understand' them well enough to 'provide accurate explanations' and 'make inductive inferences'. Laimann also states: "The problem with human interactive kinds is not merely that the classified objects change, but that they change in ways *unforeseen by our extant theoretical understanding of the world.*" (2020, 1051, italics added) Fragile relations, by definition, lead to changes that are unpredictable without supplementary knowledge, in other words precisely to changes that are 'unforeseen by our extant theoretical understanding of the world'.

A focus on fragility also clarifies why, for many purposes, it is relations that matter, not kinds. This is particularly clear in the case of causation. To explain causally requires us to identify a causal *relation*. Successful interventions, at least according to many theories of causation, also require identifying causal *relations*. And predictability, and thus the ease of successful intervention, tracks the (lack of) fragility of *relations* – by definition of fragility. Master-Model, it is made clear, is an appropriate strategy only for when *relations* are not fragile. That is why it works with domestic dogs and Newtonian theory, but not with invasive species or schizophrenia.

As Laimann points out (against Mallon), the question of how quickly *kinds* change is a red herring. For example, many bacteria change their nature very fast, but they can still be analysed successfully by evolutionary theory. Gender roles, in contrast, do not change at all because (according to Mallon) of stabilizing social effects, but we will nevertheless predict, intervene, and explain wrongly if the underlying social relations are not understood. What matters is not the kinds but the relations.

A central part of Laimann's paper is her argument that human kinds are often hybrid in a particular way. They have a dual nature: they can be understood in terms of the properties that explicitly define the category (the 'base kind'), but also in terms of the social position an individual occupies or social role the individual plays in virtue of being recognised as a member of that category (the 'status kind'). Laimann gives the example of sex as a biological base kind, versus gender as a social status kind. In much everyday and scientific speech, 'man' and 'woman' are hybrid kinds, encompassing both of these aspects. Often, the base kind is stable while the status kind is fragile (although perhaps not in the case of gender, as mentioned earlier).

The fact that many human kinds are hybrid leads to two frequent sources of scientific error, according to Laimann. The first she calls *biased conceptualization*. This is when the status element in a kind is ignored, with the result that, surprised by it, our predictions go wrong. For example, if schizophrenia is treated purely in terms of a specific symptom profile or purely as a neurological condition, then we would miss (according to Hacking) how people diagnosed as schizophrenic are singled out for particular expectations, opportunities, and treatments, and how this in turn leads to a change in the behaviour of schizophrenics and thus, ultimately, to a change in the definition of the kind itself. As Laimann says, if we conceive of a hybrid kind "solely in terms of the base kind, without considering the associated status, causal pathways associated with the status disappear out of sight." (1060) The resulting gap in our causal knowledge renders relations around the hybrid kind fragile. Just knowing schizophrenia's current definition in terms of symptoms or neurological features will leave us unable to predict future changes in the kind reliably. Another example of biased conceptualization concerns the kind 'unemployed'. Here, the base kind is a dry economic definition of being without paid work when available for it, while the status kind is the social stigma of not having a job. According to Laimann, neglect of the latter aspect impedes our understanding of unemployed people's inferior health outcomes.

The second source of error, according to Laimann, is simply *not understanding social status effects*. Social mechanisms are many and complex, and their effects, or whether they are even operating at all, are often difficult to predict. In other words, social status effects are often the products of fragile relations. Not understanding social status effects is not, so to speak, a conceptual error, unlike biased conceptualization. Rather, it is just that even when we recognise the true nature of hybrid kinds, still the social science required to study them effectively can be difficult. Laimann gives the example of the rise of the gay rights movement in the USA after the Stonewall riots in 1969, which greatly and rapidly changed the status kind component of 'homosexual'. This event was the result of a perhaps unique constellation of complex social and political processes. It was hard to predict and remains hard to fully explain.

A key point for Laimann is that neither biased conceptualization nor not understanding social status effects, directly concerns reflexivity. Reflexivity is not the key feature. Rather, each shortcoming is at root a deficit in our knowledge of causal

relations surrounding the social status aspect of hybrid kinds. If we had had this knowledge then, reflexivity notwithstanding, we could have predicted and explained successfully. I agree with Laimann here. Thinking in terms of fragility allows us to pinpoint exactly what this deficit in our knowledge of causal relations is.

Laimann convincingly and usefully shows one route – hybrid kinds – by which human sciences fall prey to fragility. But there exist other routes too, which have nothing to do with hybrid kinds (World War One truces). And fragility is not unique to human sciences (invasive species). Other instances of fragility in natural sciences arguably include many cases from ecology generally (Sagoff, 2016), and indeed from field biology generally (Dupré, 2012). Fragility is also ubiquitous in many medical treatments, in data science (Pietsch, 2016), and in complex systems generally. (Reflexivity itself is also not unique to human sciences; in addition to domestic dogs, there are arguably other cases too from biology (Cooper, 2004).)

5 Reflexivity revisited

Does a focus on reflexivity help or hinder? On the positive side, reflexivity can be a useful indicator of fragility. Laimann's mechanisms of biased conceptualization and of not understanding social status effects are two ways by which it can. As an indicator, reflexivity is not infallible though. Sometimes, as with domestic dogs and gender roles, reflexivity comes without fragility; other times, as with invasive species and World War One truces, fragility comes without reflexivity.

Being aware of reflexivity also has a second positive effect. It can alert us to specific social mechanisms, knowledge of which helps us *reduce* fragility. A familiar case illustrates: the self-fulfilling prophecies behind bank runs. Banks' cash reserves typically cover only a fraction of their depositors' credit, so if all depositors demand their money at the same time, the bank faces a liquidity crisis and can go bust. In normal times, this does not happen. But rumours or reports that a bank is in trouble can spur all depositors, made worried about the bank's solvency at the same time, indeed to try to withdraw their money at the same time. In this way, mere rumours of trouble can cause actual trouble – even when they are initially false. This is reflexivity in action. The analysis of a target – in this case, the rumours regarding the bank's solvency – itself influences that target. The positive thing for science, is that knowing this mechanism of the self-fulfilling prophecy allows us both to predict and to explain the run on the bank. Many historical bank runs have been explained in this way, at least in part, such as the collapse of the Dutch Tulip mania, the British South Sea bubble, many American banks in the Great Depression and, more recently, the collapses in 2007 of Northern Rock bank in the UK and IndyMac bank in the USA. And knowing the mechanism of the self-fulfilling prophecy does more. It also guides us towards what interventions can stymie this mechanism and thereby keep banks stable, such as granting regulators the power to prevent deposit withdrawals, or guaranteeing all deposits up to a certain value. Such measures are designed to allay depositors' fears of a run on the bank losing them their money, thereby preventing the run in the first place, and so preserving the bank's solvency. Preventing bank insolvency in this way means that relations such as that between depositing money in a bank and being confident of having access to that

money at a later date, are rendered reliable rather than fragile. That is, we are using knowledge of reflexivity to ensure stability.

A similar story is true of many other rational-expectations economic models. Knowledge of reflexivity enables us to make stable some relations that were previously fragile.

Is reflexivity an effect of fragility, or a cause of it? It can be either. Some fragile relations are causally upstream of reflexivity, others causally downstream. For example, in a bank run, will authorities intervene effectively? Suppose that this is hard to predict. Then the following relation is fragile: that a bank being rumoured to be in trouble causes the authorities to intervene effectively. In turn, because of this fragility, depositors lack reassurance, and so rumours of trouble can become self-fulfilling prophecies. That is, here fragility causes reflexivity. But matters do not stop there. For the occurrence of bank runs then causes a new relation to become fragile, as noted above, namely the relation between depositing money and being confident of having access to that money at a later date. Having itself been caused by one case of fragility, reflexivity now causes a new case of fragility.

That is the positive side, so to speak, of a focus on reflexivity. It can be an indicator of fragility, either as a cause or effect of it, and it can be a source of insight into social mechanisms. Turn now to the negative side.

The main danger is simply misdirection: we should be focusing on fragility instead, because it matters more. Reflexivity is a distraction.

But there is, in addition, a second danger: a mistaken scepticism about social prediction. For in its more radical forms, an emphasis on reflexivity denies that systematic predictive success in human sciences is possible at all. This scepticism is *a priori*. Given free will, the argument runs, humans are always free to react to any prediction about themselves in such a way as to falsify it. Suppose, for example, that an unpopular candidate is predicted to win an election because of low voter turnout for their opponent. This very prediction may then inspire the previously apathetic supporters of the opponent to come out and vote, thereby preventing the unpopular candidate from winning – and so falsify itself. Because prediction about humans is always vulnerable to reflexivity effects in this way, it is inevitably unreliable. Therefore, the argument concludes, in human sciences we cannot use prediction to test scientific theories, or as a basis for action, in the same way as we can in natural sciences.

This scepticism has had distinguished proponents. They include Hayek, Popper, MacIntyre, many interpretivists, and various intellectuals outside academia, such as George Soros, Michael Frayn, and Jonathan Miller.

Some responses to this scepticism, such as Mill's, have claimed that free will is compatible with determinism. But whether it is or not, a better reaction, I think, is to point out the obvious fact that social predictions often are successful. The interesting thing is when and why they are.

No doubt, social prediction is challenging. Prediction generally is challenging. But from a methodological point of view, reflexivity is merely one source of fragility. It no more makes third-person causal investigation impossible than does any other source of fragility. As Laimann emphasizes, what matters is not that reflexivity moves the goalposts, but rather whether we understand the mechanisms behind the moving of the goalposts. If we do, then we can still predict and explain perfectly well, as we do

with domestic dogs and gender roles. Reflexivity does not somehow magically negate this. A priori arguments that it does are falsified by ample experience.

But it is not just that scepticism about social prediction is misguided. It is also pernicious. Why? Because, in effect, it seeks to deny to human sciences the possibility of empirical testing, which is the key to advancing knowledge. Here is one example. (There are many others.) In a UK government press conference in April 2020, Health Secretary Matt Hancock was asked whether total UK Covid-19 deaths could still be kept below 20,000. (The official figure had just reached 10,000.) He replied: “The future path of this pandemic in this country is determined by how people act. That’s why it’s so important that people follow the social-distancing guidelines. *Predictions are not possible, precisely because they depend on the behaviour of the British people.*” (BBC, 2020, italics added) Here, Hancock explicitly endorses the a priori scepticism about social prediction. At one level, his statement can be read simply as a statement of epistemic humility, correctly noting that part of the causal chain determining case numbers would be the public’s behaviour. But the statement also makes it impossible to hold policy to account. It is deemed that we cannot fairly assess whether a prediction of a particular policy’s effect is rational and, thus, whether the policy is worthy of praise or blame. No prediction can be deemed better than any other. Responsibility for the outcome is conveniently evaded – and, in this case, put on the public instead.

Of course, the context here was that Hancock wanted (understandably) to maximize popular following of restrictions, and so had reason to emphasize the importance of that rather than of any particular prediction. Perhaps telling the public that the outcome depended on its own behaviour was merely in the service of this urgent practical imperative, and Hancock himself did not really believe that social predictions cannot be fairly evaluated. But what Hancock himself did or did not believe is beside the point here. When it gives cover to such evasions of responsibility, the a priori scepticism is pernicious, both scientifically and morally.

Denying rational social prediction is, in effect, philistine. It denies that much actual, successful scientific inquiry is even possible. From where does this philistinism arise? A full answer is beyond this paper, but here is a speculation. The root cause of the error is a mistaken focus on reflexivity rather than fragility and, in turn, this mistaken focus reflects an agenda drawn primarily not from philosophy of science or from the science itself, but instead from wider philosophy. This agenda is ultimately external. It can be seen in many traditional handbooks, anthologies and introductions to philosophy of social science, or in, to take one example, how David Papineau remembers and presents the agenda of philosophy of social science in (Papineau, 2008). Typical questions are: ‘What is intentional explanation?’ and ‘How can we causally explain human action?’, which reflect the agendas of philosophy of action and philosophy of mind; ‘Is there collective agency?’, which reflects the agendas of metaphysics, ethics, and philosophy of action; and ‘Do special sciences reduce to physics?’, which reflects the agendas of metaphysics and philosophy of mind. The same questions drive much of the attention given to reflexivity. But none of them is primarily motivated by wondering what methods make human sciences successful or unsuccessful, where this is measured in the currency of predictions, explanations, and interventions. For that, we should focus instead on fragility.

Acknowledgements I am grateful for useful feedback, first, from the audience at the Reactivity, Prediction and Intervention in the Human Sciences conference in Helsinki in August 2021, and second, from two anonymous referees for this journal.

Declarations

Ethics approval N/A

Informed consent N/A

Conflict of interest None.

References

- BBC. (2020, April 12). Coronavirus: ‘Sombre day’ as UK deaths hit 10,000. *BBC News Online*. <https://www.bbc.com/news/uk-52264145>
- Cartwright, N. (2019). *Nature, the artful modeler*. (Open Court).
- Cooper, R. (2004). Why Hacking is wrong about human kinds. *British Journal for the Philosophy of Science*, 55, 73–85.
- Dupré, J. (2012). *Processes of life*. Oxford University Press.
- Elliott-Graves, A. (2016). The problem of prediction in invasion biology. *Biology and Philosophy*, 31, 373–393.
- Elliott-Graves, A. (2018). Generality and causal interdependence in ecology. *Philosophy of Science*, 85, 1102–1114.
- Elliott-Graves, A. (2019). The future of predictive ecology. *Philosophical Topics*, 47, 65–82.
- Hacking, I. (1995). The looping effects of human kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 351–394). Clarendon Press.
- Hacking, I. (1999). *The social construction of what?* Harvard University Press.
- Khalidi, M. (2010). Interactive kinds. *British Journal for the Philosophy of Science*, 61, 335–360.
- Laimann, J. (2020). Capricious kinds. *British Journal for the Philosophy of Science*, 71, 1043–1068.
- Mallon, R. (2016). *The construction of human kinds*. Oxford University Press.
- Mill, J. S. (1843). *A system of logic*. Parker.
- Northcott, R. (2013). Degree of explanation. *Synthese*, 190(15), 3087–3105.
- Northcott, R. (2022). *Science for a fragile world*. Oxford University Press (Forthcoming).
- Northcott, R., & Alexandrova, A. (2015). Prisoner’s Dilemma doesn’t explain much. In M. Peterson (Ed.), *The Prisoner’s Dilemma* (pp. 64–84). Cambridge University Press.
- Papineau, D. (2008). Five philosophy of social science answers. In D. Rios & C. Schmidt-Petri (Eds.), *Philosophy of the social sciences: Five questions* (pp. 99–110). Automatic Press.
- Pietsch, W. (2016). The causal nature of modeling with big data. *Philosophy and Technology*, 29, 137–171.
- Sagoff, M. (2016). Are there general causal forces in ecology? *Synthese*, 193, 3003–3024.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University Press.
- Ylikoski, P. (2019). Mechanism-based theorizing and generalization from case studies. *Studies in History and Philosophy of Science*, 78, 14–22.

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.