

# Ex Ante Facts, Causal Structures, and Prediction

By Chad Harris\*

## Abstract

In his *A Realist Philosophy of Economics* Karl Mittermaier (2023) commends the importance of distinguishing between two order of facts which he designates as *ex post* and *ex ante* facts. He argues that making this distinction allows us to identify and exploit useful *ex ante* facts such as structural and causal factors, which in turn can be used as an epistemic warrant for making predictions. This article scrutinises work in philosophy, economics and the social sciences where attempts are made to follow through on this strategy and make predictions underpinned by structural and causal facts. It explores points of comparison and contrasts between this work and the exposition of *ex ante* facts provided by Mittermaier.

*JEL Code: B40*

*Keywords: Ex ante Facts, Ex post Facts, Structural Facts, External Validity, Causal Structures, Prediction*

## 1. Introduction

In *A Realist Philosophy of Economics* Karl Mittermaier draws our attention to the benefits of distinguishing between two orders of facts in economics, namely *ex post* (loosely “after the event”) and *ex ante* (loosely “before the event”) facts. Mittermaier argues that *ex ante* facts are particularly useful in disciplines such as economics because they provide epistemic warrant for predictions. The intriguing story of the genesis and subsequent neglect of his work in his dissertation meant that efforts towards further explicating the distinction and executing the work Mittermaier envisioned for *ex ante* facts were never fully followed up during his lifetime.<sup>1</sup> However, over time, scholars in other fields and disciplines developed theories and arguments that proceeded according to the strategy hinted at by Mittermaier, namely developing accounts of prediction based on the discovery of *ex ante* structures. In this article, I examine some of this work with the aim of comparing these accounts of prediction with Mittermaier’s characterisation of *ex ante* facts. I start by describing the *ex ante/ex post* fact distinction and explaining how Mittermaier envisioned *ex ante* facts could be put to use for predictive work. Thereafter, I outline and critically evaluate the theory of institutional facts put forward by John Searle and discuss how they compare against *ex ante* facts. Next, I describe the problem of external validity and compare some

---

\* Department of Philosophy, University of Johannesburg, Kingsway Campus Auckland Park, 2092 Johannesburg, South Africa. The author can be reached at [cmhharris@uj.ac.za](mailto:cmhharris@uj.ac.za).

<sup>1</sup> For a detailed explanation, see the extended preface and prologue.

of the prominent purported solutions to this problem with the strategy of basing predictions on *ex ante* facts, before concluding.

## 2. Mittermaier's Distinction

In his book *A Realist Philosophy of Economics*, Karl Mittermaier makes the case for distinguishing between two orders of facts, which he terms *ex post* and *ex ante* facts. He describes the former as “those facts which owe the meaning they have for us to their position in a possibly unique and fortuitous or stochastic course of events” (Mittermaier 2023, 18). On the other hand, he defines another, *ex ante* order, or the order of “facts, which relates to structures, that is, to not altogether transient patterns, and which may, therefore serve as a guide to our purposeful action” (*ibid.*). The advantages of drawing this distinction are suggested by the following considerations:

Let us consider what rising price indices tell us about inflation. Do they tell us *why* prices are rising or do they tell us *that* prices *have* risen? Do they convey the sort of information we should need to devise a policy to counter inflation, i. e. do they provide us with a guide to action in the future? Common sense would tell us to distinguish between facts that merely record past events and facts that somehow relate to the structure of the economy. If we want to recommend an anti-inflationary policy we may well look at price indices to see what actually has taken place, but we would know that this in itself would not be enough. We also would need knowledge of certain structural features of the economy (*ibid.*, 17).

Mittermaier is evidently seized with the following problem: merely having a list of facts that have already occurred is fine for certain purposes, but this is not particularly useful for many other important purposes in economics and other sciences. If we really want to gain traction when it comes to our predictions about economic phenomena and our explanations for why they occur, we need to do more than simply list a series of past events. We must also find out about the underlying structures in the world that are responsible for the occurrence of those facts.

He illustrates this principle with a mechanical analogy. Consider a simple physical structure such as a seesaw. If we placed an object on one side of the seesaw, we can easily predict that this object will rise if we push the other end down. We can make this prediction because of our knowledge of the underlying structure consisting of the lever and fulcrum. The existence of this structure, for Mittermaier, is an *ex ante* fact, and knowing this fact allows us to explain and predict a selection of related *ex post* facts that come about, or could potentially come about, as a result of this structure. As he explains:

I have now used the structure as a guide to action. By pressing down one end of the lever, I cause the other end to go up – there is a causal relation between the going down of one end and the going up of the other. Moreover, I can predict the consequences of my action without actually engaging in it. It may seem trivial when put this way, but I believe that this limited sense of the words cause and predict is the most useful sense, and that it is in this sense that prediction and causation are most usually thought of in the applied physical sciences (*ibid.*, 36).

This strategy clearly makes sense when it comes to prediction of physical phenomena. The question remains whether a similar strategy can be made to work in the context of the social world. In fields such as economics and other social sciences, identi-

fying *ex ante* facts can be just as useful as enduring physical structures. As Mittermaier's original example in the quotation suggests, understanding the underlying structures in economic and social phenomena would allow us to effectively predict and explain events and phenomena relevant to our knowledge and theories in these fields, such as inflation. What the quotation suggests is that tracking the past movements of something like the inflation rate is of limited use in the absence of facts about the underlying structures responsible for inflation. To truly understand, explain and predict inflation rates and their movements, we must also have knowledge of *ex ante* facts, those facts related to the structure of the economy.

Doing so means going beyond reliance on the mechanical analogy above and providing a more formal characterisation of *ex ante* facts in economics. Mittermaier explains that *ex ante* can be considered an umbrella term of sorts, covering two major types of facts, namely structural and causal facts (*ibid.*, 39). He defines these as follows:

I define *structural knowledge* or a *structural fact* as the belief that if in a particular situation one observes the set of elements A identified by defining procedures B, then by following the relating procedure C, one may observe in the same situation or a later situation specified by the relating procedure C, a set of elements D that can be identified by defining procedures E. [...] I define *knowledge of causal relations* or *causal facts* as the belief that if one performs the operation X in a situation in which one observes the set of elements A identified by defining procedures B, then by following the relating procedure C, one may observe in a later situation specified by the relating procedure C, a set of elements D that can be identified by defining procedures E (*ibid.*, 36–39).

After introducing and defining these terms, Mittermaier expends most of his efforts in the rest of the book pointing out the ways in which *ex ante* and *ex post* facts are conflated in orthodox economics. This, combined with the fact that this material has not been circulated until now, means that there has not been any deep scrutiny of the strategy of discovering and exploiting *ex ante* facts for the purposes envisioned by Mittermaier. However, in the period since this distinction was drawn, other work in economics, social science and philosophy has proceeded along lines that resonate with the definitions of causal and structural facts given above. In what follows I critically examine the work of philosopher John Searle on institutional facts and the work of some theorists in economics and the social sciences addressing the problem of external validity. This should not be interpreted as me claiming that *ex ante* facts are identical to institutional facts, or that the problem of external validity is tantamount to Mittermaier's concern with using *ex ante* facts to make predictions. I choose these specific strands of enquiry because they provide interesting points of comparison and contrast with Mittermaier's distinction. The unifying element in all of these approaches is that they are premised on using structural and causal facts as an epistemic warrant for predictions or extrapolations.

### 3. Institutional Facts

The impetus for Searle's work on institutional facts is similar to the concern that motivated Mittermaier's work on *ex ante* facts, namely the insight that we can distinguish

different orders of fact in the world. Searle pointed out that certain phenomena can be considered facts, even if they do not seem to work on the same model as brute physical facts. There is an order of facts, emanating from the social world, that are a function of human behaviour and appear to depend on shared beliefs and conventions. In his article “What is an Institution?” Searle explains the difference between these orders of facts as follows:

In some intuitively natural sense, the fact that I am an American citizen, the fact that the piece of paper in my hand is a 20 dollar bill, and the fact that I own stock in AT&T, are all institutional facts. They are institutional facts in the sense that they can only exist given certain human institutions. Such facts differ from the fact, for example, that at sea level I weigh 160 pounds, or that the Earth is 93 million miles from the sun, or that hydrogen atoms have one electron (2005, 2–3).

As is the case with *ex ante* facts, having knowledge of institutional facts increases one’s predictive abilities. Consider, for example, observing the behaviour of groups of students in a university building during the course of a morning. For most of the time, the distribution of students might be evenly spread across different venues such as passageways, lecture venues, seminar and reading rooms. However, at a specific time, all of these venues empty out and the vast majority of the traffic gets directed towards one specific location. If we are apprised with knowledge of a specific institutional fact, namely the fact of the university’s timetable showing the 11 am break which is collectively observed as the morning coffee break, we would then be in a good position to predict this sudden surge of students towards one particular venue, namely the coffee shop, at that specific time. Considered this way, the institutional fact of the observance of the 11 am break functions in much the same way as Mittermaier’s seesaw. It is a persisting structure, albeit a social structure in this case, that allows us to make sense of a flux of events that would have no explanation in the absence of the structure.

This serves us nicely as an informal illustration of what institutional facts are and how they can assist our predictive and explanatory abilities. This still falls short, though, of being a theory of institutional facts. For the purposes of this article I cannot provide an exhaustive exposition of Searle’s theory, but I will rely on Moural’s (2002) summary of his framework. According to Moural, Searle’s core theory is comprised of two main elements. These are, firstly, an analysis of institutional facts in terms of the “counting as” form: “(I) X counts as Y in context C”; secondly, an account of the role of power in the functioning of institutional facts (Moural 2002, 273–74).

To apply these two elements to our simple example, we can say that the variable X stands for the period of time starting from 11 am, the variable Y stands for the coffee break, and the context C refers to the university. In terms of power, we can say that in this context, the institutional fact of the 11 am observance of the coffee break manifests itself in the legitimate practice of certain powers and practices. A student would not normally have the power to get up in the middle of a lecture and leave the venue, but her leaving the venue would be considered normal during the 11 am coffee break. This also makes it clear how the institutional facts of the coffee break are related to other institutional facts, such as the institutional fact of the lecture and the behaviour that it engenders.

Of course, it does not imply that the observance of the break, or the lecture, is a practice that is conformed to without exception. To say that a student cannot act at will should not be interpreted as equating social conventions with laws of nature. The student wanting to get up and leave is not constrained in the same way that a light particle is constrained to travelling at a certain speed. While there may be negative consequences to breaking the rule about staying seated during the lecturer or of only going to the coffee shop during the break, it is of course possible for there to be exceptions to these behaviours. However, the force of the social convention means that it is possible to predict the behaviour of the students with an impressive degree of certainty, albeit short of full certainty. What makes the practice of this convention possible, and hence another crucial element in any analysis of institutional facts, is the idea of collective recognition. For this reason, Dörge and Holweger build the idea of collective recognition into the very definition of an institutional fact: “That the (putative) fact that *p* is an ‘institutional fact’ means that *p* is a fact that exists by virtue of the collective recognition of its existence” (2021, 4955).

Now that we have an idea of what institutional facts are, it becomes clear why they are considered useful when it comes to predicting events in the social realm. However, there is still the underlying and unresolved question of weighing up the evidence for whether they exist. To the uninitiated, this might seem strange. After all, we have been speaking about the differences between institutional and other facts, and we have a story to tell about how institutional facts power our predictions and explanations. Surely, one seems justified in arguing, this means that they must exist and to further belabour this point is merely of philosophical interest. However, the question of the precise ontological status of institutional facts, that is, in simple terms, whether they actually exist, is not as simple as this suggests. To see why, consider a similar analysis of the notion of unicorns. We can stipulate with confidence how they are defined, and we can make up a story about the benefit of having them around, but this does not amount to a vindication of their existence.

In order to gain traction towards such a vindication, we need to state with precision exactly in what sense we are considering the existence of institutional facts. We can take as our starting point Moural’s (2002, 272) insight that Searle himself settles for the resolution of this issue by stipulating that institutional facts, given that they depend on our collective recognition, should be considered ontologically subjective but epistemically objective. However, settling for this outcome seems risky if we consider the *ex ante* facts we have been relying on. Think back to Mittermaier’s seesaw, our initial example of how *ex ante* facts are meant to work. The point of this mechanical analogy is that in the same way that there are actually existing structures that allow us to make predictions about physical objects, we need similar structures in the context of economic phenomena in order to make predictions about these events. In other words, the fundamental question we are dealing with is *not* whether such facts are subjective or objective, but whether they exist or not. To speak of these facts as “ontologically subjective” conflates the epistemic and the metaphysical aspects of this question. If Mittermaier’s mechanical analogy is to be taken seriously, then we must ask whether these facts truly exist in the metaphysical sense, epistemic status notwithstanding.

One approach to answering this question for institutional facts is found in the work of Dörge and Holweger (2021), who argue that there are good reasons to think that institutional facts do not really exist. To demonstrate why, they appeal to six existence criteria, or factors that we typically use in deciding whether anything exists or not. In their article, they argue that institutional facts fail to meet the following existence criteria: representation independence (4960), causal integration (4961), spatiotemporality (4962), absoluteness (4963), consistency (4965) and non-graduality (4966).

Space constraints prevent me from going into a detailed evaluation of their arguments for why institutional facts fail each of these existence criteria, but some of the key points include that institutional facts are not representation-independent because they require our collective recognition. If you consider our coffee break example, the institutional fact of the observance of the coffee break only exists for as long as those who are members of the university observe it. If the university ceases to exist, or if the governing body decides that the timetable will be changed, then the institutional fact of the 11 am coffee break will cease to exist. This is in contrast to a fact that exists unequivocally, such as the fact that the Earth came into existence billions of years ago. This fact exists and our collective recognition of it has no impact on its coming into and going out of existence.

Institutional facts are also not spatiotemporal except in a derived sense. In other words, an institutional fact is only spatiotemporal to the extent that it is made up of other basic facts that are themselves spatiotemporal. For instance, the institutional fact of the coffee break ultimately boils down to the location of a certain number of bodies in a certain space at a certain time. These bodies themselves have a spatiotemporal profile, but this is not the same as saying that the coffee break has a spatiotemporal existence.

Defenders of institutional facts could legitimately take issue with the list of existence criteria that Dörge and Holweger advance. It could be argued that this list does not constitute all the necessary and sufficient factors that are relevant for determining whether something truly exists. It is also possible to disagree with Dörge and Holweger about their interpretation of why institutional facts fail to meet certain existence criteria. However, regardless of the specific list of existence criteria used, the real problem with making the case for institutional facts being real, according to Dörge and Holweger, is linked to two phenomena they describe as ontological projection and the Thomas Theorem.

Ontological projection refers to “the generation of an entity or its maintenance in existence by virtue of collectively representing it as existing” (2021, 4598). They argue that proponents of the reality of institutional facts would need to justify the possibility of ontological projection, but the authors deny that ontological projection is feasible. Without an explanation for how belief in ontological projection is warranted, the argument for the existence of institutional facts falls flat. The Thomas Theorem (2021, 4970) describes a psychological phenomenon whereby groups of people, because of a collective recognition of something, can behave as if something existed even when that thing does not exist in reality. As a simple example, think of the phenomenon of the bogeyman, and how this is used to regulate the behaviour of young children. As another example, imagine what would have happened if, for some reason,

the death of Queen Elizabeth II had been withheld from the public for a few weeks after it occurred. Due to the collective recognition of the status of the monarch, most of British society would have proceeded as if the Queen still existed, even though she no longer did.

The impossibility of ontological projection and the lessons of the Thomas Theorem do not augur well for those committed to establishing the ontological existence of institutional facts. This is because institutional facts, by definition, are entities that come about through collective recognition and, if ontological projection is indeed impossible, there does not seem to be any good reason for assuming that they exist in the same way that mind-independent entities exist. However, there is a way of saving face and rescuing some form of existence claim for the institutional facts. In similar vein to the comments made about spatiotemporality above, it is possible for the proponent of institutional facts to claim that they exist to the extent that they are underpinned by other, more basic facts, which have a far more convincing claim to satisfying any existence criteria. This is clearly the line taken by Searle himself in the following passage:

...all sorts of things can be money, but there has to be some physical realization, some brute fact – even if it is only a bit of paper or a blip on a computer disk – on which we can impose our institutional form of status function. Thus there are no institutional facts without brute facts (1995, 56).

This strategy suggests a similar avenue for the justification of the existence of *ex ante* facts in the context of economics. Here, the brute facts grounding the *ex ante* facts would be facts related to the psychological make-up of humans, and/or the social dynamics governing institutional decisions and group behaviour.<sup>2</sup> Provided these factors are accepted as uncontroversially real, they can legitimately be taken as providing the requisite support to the existence of the type of *ex ante* facts that provide epistemic warrant to predictions in economics. As the discussion above reveals, we can legitimately interpret Mittermaier's work on *ex ante* and *ex post* facts as consistent with, and related to, the Searlian conception of institutional facts. Perhaps the most salient difference is that Searle's work on institutional facts was not formulated so as to be directly relevant to the type of prediction required in economics. Next, I turn to a problem that is closer to home in the sense that it is a contemporary concern in economics and other social sciences.

#### 4. External Validity and Causal Structures

The distinction between internal and external validity traces its origin to the work of psychologist Donald Campbell, who drew attention to the importance of differentiating between two types of criteria in judging the success of an experimental intervention. On the one hand, we need to ask “did in fact the experimental stimulus make some significant difference in this specific instance?” (1957, 297), in which case we are asking about the internal validity of the intervention. On the other hand, we might want to know “to what populations, settings, and variables can this effect be

---

<sup>2</sup> I am indebted to Professor Rod O'Donell, University of Sydney, for suggesting this possibility in response to an earlier draft of this paper.



generalized” (*ibid.*), in which case we are curious about the external validity of the intervention. More formally, we can adopt the definition given by philosopher of science Francesco Guala:

Internal validity is achieved when some particular aspect of a laboratory system (a cause-effect relation, the way in which certain factors interact, or the phenomena they bring about) has been properly understood by the experimenter. For example: the result of an experiment E is internally valid if the experimenter attributes the production of an effect Y to a factor (or a set of factors) X, and X really is a cause of Y in E. Furthermore, it is externally valid if X causes Y not only in E, but also in a set of other circumstances of interest F, G, H, ... (Guala 2005, 142).

This definition demonstrates why this distinction is relevant to our discussion. Internal validity can be seen as the attempt to excavate or discover actual causal structures. External validity is about using these causal structures to make predictions about whether the same effects can be replicated in a different contexts. It is also about using these causal structures as guides to action in the way Mittermaier hoped *ex ante* facts could be used as guides to action. For example, even if we know that investing in infrastructure has led to economic development in one country, we cannot necessarily infer that the same programme of infrastructure build will have the same effect in another country. It may, or it may not. Understanding why this is the case and developing an account of when the same intervention will lead to the same results require an account of external validity, not just a method for detecting causal structures. Or it may lead to the more general proposition that programmes of infrastructure development benefit countries, with different programmes suitable for different countries in different contexts.

When we look at some of the prominent attempted theoretical solutions to the problem of external validity, we notice more commonalities with Mittermaier’s appeal to *ex ante* facts. Guala (2005, 2012), for example, puts forward an account informed by the actual techniques used by economists who conduct experimental research. He explains that he is following in the tradition of the founders of experimental economics, specifically Vernon Smith and Chris Starmer, who see external validity as an empirical problem that should be resolved empirically. This means that the best way to approach external validity is by combining any available evidence from the original study or trial, with other information regarding the environment of interest. Specifically, we are advised to look for evidence pointing to any reasons for suspecting that the causal structure identified in the experimental context may be missing, or significantly different, in the target environment (Guala 2012, 628).

As an example of using his causal-analogical approach to solve external validity problems, Guala explains the strategy of “exporting the lab.” This strategy works by “modifying the experiment to include the features of the target that could be responsible for the alleged external validity failure, and see whether they in fact make a difference or not” (*ibid.*, 628). This method is a form of falsificationism as applied to experimental inferences, but it uses potential threats to external validity as a way of updating and improving the experimental method:

The idea, roughly, is that we need to create (or select) circumstances in which it is really unlikely to observe certain data, unless the external validity hypothesis is true. In this case, the data is the correspondence between observed features of the target and observed features of



the experimental system; the external validity hypothesis is that the relata belong to similar causal mechanisms. Now, the probability of observing such a correspondence (were the hypothesis false) is low if we have eliminated alternative reasons why such a correspondence might occur, other than the causal similarity between the two systems (Guala 2003, 1202).

This shows the affinities between Guala's resolution of the external validity problem and Mittermaier's appeal to *ex ante* facts, including facts pertaining to causal structures, as support for our predictions. In essence, Guala is pointing to the usefulness of employing already-established knowledge about causal structures to make better inferences and predictions.

Another theorist that has made contributions to the external validity debate consistent with Mittermaier's distinction is Nancy Cartwright (2007). While she does not employ the term *ex ante fact*, she stresses the importance of identifying stable tendencies as a necessary first step in making predictions or extrapolations about physical phenomena. In a different work, she makes use of the notion of a nomological (or law-based) machine as fulfilling a similar function. A nomological machine is described as "...a fixed (enough) arrangement of components, or factors, with stable (enough) capacities that in the right sort of stable (enough) environment will, with repeated operation, give rise to the kind of regular behaviour that we represent in our scientific laws" (Cartwright 1999, 50). These nomological machines, if they are properly identified, can be expected to hold outside of laboratory context, and this holds the key to how scientific knowledge gained through experimental methods allows us to deal with prediction, extrapolation and the external validity problem.

When it comes to economics and the social sciences Cartwright's approach is similar. The stable tendencies identified in the types of experiments conducted in economics are couched in terms of causal roles or causal principles, but the basic idea remains the same. In Cartwright and Hardie (2012), for example, the authors discuss approaches that interpret external validity as a problem about exploiting knowledge about similarities and differences of causal structure between experimental and target environments. A nice illustration of this is the case study, commonly used as an example of external validity failure, about the failure of the UN's Bangladesh Integrated Nutrition Programme, despite the success of an identical intervention in neighbouring India. They point out that the two environments, despite their similarities, displayed small cultural differences that ultimately meant there were significant differences in the causal structure across the two environments. This difference meant that the intervention yielded very different results in the new environment, and hence resulted in the failure of the prediction that the programme would be successful in Bangladesh.

Another approach that develops a model for extrapolation based on the discovery of underlying structures is Daniel Steel's model process-tracing approach, outlined in his book *Across the Boundaries: Extrapolation in Biology and Social Science*. The foundation of the process-tracing approach is to identify the relevant mechanisms or structures underpinning the causal relationships we want to use for the purposes of prediction. Steel characterises these causal mechanisms as being composed of "...interacting components that generate a causal regularity between some specified beginning and end points" (Steel 2008, 40). He gives different examples of mechanisms from both the physical and social sciences. As an example of a simple physical mech-

anism he lists the working of a car engine, and for a more sophisticated example from the biological sciences he describes the process of protein synthesis in cells:

First, a strand of DNA unwinds and the adjoining nucleotide bases separate. The next step is the transcription of the unwound DNA by messenger RNA (mRNA), the order of the bases of the mRNA being determined by the order of the complementary nucleotide bases in the DNA strand. Finally, the strand of mRNA serves as a template for transfer RNA (tRNA), which assembles a string of amino acids into a protein (Steel 2008, 40).

The functioning of this mechanism is what allows biologists to make predictions about the manifestation of certain traits in organisms, based on knowledge of their genes. This relationship between genes and trait, although not a causal relationship in the strictest sense, allows biologists to be reasonably certain when the regularity between certain genes and certain traits can be expected to hold.

This feature of mechanisms illustrates how, with the appropriate knowledge about the functioning of the relevant mechanisms, we are able to make the sort of inferences that allow us to determine whether the causal relationships discovered in our models or experiments will transport to reality. In other words, mechanisms provide a potential method for making predictions, extrapolations or external validity inferences:

First, learn the mechanism in the model organism... Second, compare stages of the mechanism in the model organism with that of the target organism in which the two are most likely to differ significantly (*ibid.*, 89).

One of his most insightful examples of how this method aids extrapolation is the case of the metabolism of aflatoxin (*ibid.*, 91). Scientists were interested in investigating the potential carcinogenic effects of aflatoxin, a naturally occurring toxin in certain foodstuff, for humans. It was established, through laboratory studies, that exposure to aflatoxin led to cancer in rats, but it was not clear that this was good evidence for it being carcinogenic in humans because different species have different metabolic mechanisms. Experimenters could also establish that mice were immune from the carcinogenic effects of the toxin. The crucial piece of evidence for this extrapolation was the discovery that when it came to the metabolism of a specific enzyme in the liver, humans were closer to rats than to mice. It was this difference in metabolic mechanisms that allowed scientists to predict that the toxin would be carcinogenic in humans, based on an extrapolation regarding our similarities to other animals.

For Steel, this shows the potential of process-tracing as a device to aid extrapolation. It also demonstrates why process-tracing is an improvement on simple induction. According to Steel the salient difference pertains to the sophistication of the generalisations we are able to engender using process-tracing as opposed to the generalisations emanating from simple inductions:

Simple induction depends upon generalizations of the form “What is carcinogenic for rats is probably carcinogenic for humans, too.” In contrast, comparative process tracing depends upon generalizations like “Features A, B, and C of carcinogenic mechanisms in rodents usually resemble those in humans, while features X, Y, and Z often differ significantly” (*ibid.*, 89).

The process-tracing method works by marshalling our knowledge of mechanisms to bolster our extrapolation. The type of inference illustrated by the aflatoxin example is

reliable to the extent that it identifies mechanisms that are significant in that they determine the success or the failure of the causal relationship of interest. This allows the extrapolator, when she has knowledge of potential differences in these mechanisms between the experimental and target organisms, to predict the likely success of the causal extrapolation.

When it comes to extrapolation in the social sciences, Steel admits that the prospects for mechanism-based extrapolation are less promising than they are in biology. This is because of two difficulties related to social mechanisms. First, in the social sciences it is not always easy to identify the relevant mechanisms for understanding the causal relationships in a social system. The second problem is that intervening on social systems may alter the structure of the mechanisms responsible for the regularities we hope to achieve. This issue is much debated in the economics literature as illustrated by the Lucas Critique and the Hawthorne effect. In short, the argument is that mechanism-based extrapolation in the social sciences is susceptible to being scuppered because of an absence of mechanisms that are sufficiently stable under intervention.

Nevertheless, Steel makes the case that a process-tracing perspective can be adapted to improve predictions based on causal structures in the social sciences. As an example of process-tracing in social science, Steel adapts an example from anthropology regarding the practices of gift-giving and marital custom in Trobriand society (*ibid.*, 189–190). Steel explains how these entrenched customs and protocols create social processes and practices that link together to form social mechanisms. For example, the Trobrianders have a custom whereby brothers give gifts of yams to the husbands of their married sisters. The size of the gift is determined by the social status of the husband, and so chiefs get more sizeable gifts from their in-laws. This practice, combined with the practice of Trobriand chiefs financing their projects with yams, helps explain the link between polygamy and wealth in Trobriand society (*ibid.*, 189).

Steel thinks this example demonstrates the potential of process-tracing in social science because it works by considering the social actors and roles as components of the social system. It then assumes, based on judgement about the behaviours of and inter-relationships between these components, the existence of certain causal regularities between the components. In this example Steel claims the causal generalisations would be psychological in nature, related to “aspirations for wealth and social status” (*ibid.*). These generalisations can then be used to infer social mechanisms, in this case taking the form of the social customs in Trobriand society as a whole. These mechanisms can then be legitimately employed as evidence, even in the absence of statistical data about the components of the social system, for the existence of causal relationships such as the one between wealth and polygamy. Therefore, process-tracing allows us to predict the existence of the link between wealth and polygamy in Trobriand society even in the absence of data about wealth and number of wives among Trobriand chiefs.

The final approach I will discuss is unique in that it is not only premised on exploiting causal structures for prediction, but also develops a formal notation for representing those causal structures. In “External Validity: From Do-Calculus to Transportability Across Populations” (2015), Judea Pearl and Elias Barenboim, who approach the problem from the perspective of computer science, outline a purported solution to the

problem of external validity based on the tools developed in Pearl's (2009) account of causal inference. This approach works by modelling systems of causal relationships using directed acyclical graphs (DAGs) in conjunction with structural equations (SEMs). This approach to causal inference, which is part of a broader outlook sometimes called the potential outcomes approach (POA), has become increasingly influential in many fields including statistics, economics and econometrics (Rubin 2005; Heckman and Pinto 2015), as well as epidemiology and other fields where discovering causal relationships through experimentation is critical. The Pearlian method is praised as a step forward for causal inference because it provides a formal and rigorous way of dealing with causal information. For its proponents, this step means that causal inference and related problems are now amenable to genuine scientific treatment, and can be rescued from the realms of philosophical speculation. When it comes to the specific problem of extrapolation or external validity, what Pearl refers to as "causal transportation," the Pearlian approach involves constructing inferential queries that can be handled by a calculus of interventions developed by Pearl, called do-calculus. In the discussion that follows I will base my characterisation of this approach on Pearl and Barenboim's (2015) article, which is representative of the series of articles (2011, 2013 and 2015) by these authors in which the causal transportation problem is broached. I refer to this as the "P&B" approach from here on.

To see how the P&B approach would handle causal transportation, we are given illustrative examples of scenarios where we might have to make decisions regarding causal transportation. We are asked to consider three DAGs based on hypothetical scenarios where the results of a randomized controlled trial (RCT), conducted in Los Angeles, detects a causal relationship between variable  $X$  and  $Y$  such that  $X$  causes  $Y$ .<sup>3</sup> The first DAG represents the information from the RCT that tells us the effect of  $X$  on  $Y$  for every age group ( $Z$ ). P&B ask us to consider how we would go about making sense of the information we would need in order to make an external validity inference about what we can expect about the effect of  $X$  on  $Y$  in a different environment. In their example the new environment is New York, and in addition to the results of the RCT we are to assume that we also have knowledge that the average age in New York is significantly higher than in Los Angeles. This means that the value given to variable  $Z$  will be different in the two environments. However, P&B maintain that for this experiment it would be relatively easy to combine the information about the two environments and estimate, using what they call a transport formula, the average effect size for each age group in New York, based on the data about effect size in the different age groups from the Los Angeles RCT.

In the second DAG we do not have information about the average age in the environments, but we do have data about language proficiency, which is correlated with age. In the third DAG,  $Z$  is meant to represent something like a bio-marker that is dependent on  $X$  but that could have an effect on  $Y$ . The crucial point about these two scenarios, according to P&B, is that any attempt at transporting the causal relationship is going to be a lot more complicated, if it is even possible, than it is in the first case. For example, the difference in the values of  $Z$  between the two populations could be attributed either to genuine age differences between the two environments, or differ-

<sup>3</sup> The full DAGs can be seen in Pearl and Barenboim 2015, 584.

ences in the way language proficiency correlates with age between the environments. These two possibilities call for very different expectations regarding transportability of the causal relationship, very different transport formulae appropriate for estimating the size of the effect in New York, and consequently different approaches to the extrapolation. The potential for *Z* to confound the relationship of interest in the third case means, for example, the most appropriate approach to external validity may be to ignore the data from the original trial, and it may include finding other information that could help us estimate the effect in the new environment.

The upshot of all of this, according to P&B, is that licensing transportability requires “knowledge of the mechanisms, or processes, through which population differences come about; different localization of these mechanisms yield different transport formulae” (*ibid.*, 586). Where such knowledge can be gathered, P&B have developed a formal representation to encode the differences between environments that come about as a result of these mechanisms. Their solution is thus to augment that DAGs discussed in connection with the example above by attaching new variables, which they call selection or S-variables, to all the variables where population differences are suspected to occur. These S-variables are meant to represent all the factors could possibly threaten the transportation of the causal relationship between study and target populations. According to P&B, the presence of an S-variable on an augmented DAG implies the existence and location of a mechanism responsible for differences between populations and the absence of S-variables implies the absence of causally relevant differences.

The augmented DAGs for the examples we have discussed in this section would represent the information that the New York and Los Angeles population differ in terms of average age would be represented by the S-variable attached to the *Z* variable.<sup>4</sup> There are similar S-variables for the population differences in the cases involving language proficiency and the bio-marker as the mechanisms responsible for the population differences. These augmented DAGs can be used as a guide to make predictions or external validity inferences. In other words, the DAG is a representation of a causal structure, or ex ante fact in Mittermaier’s terms. The augmented DAG then provides guidance about how that ex ante fact can warrant a prediction about how the causal relationship will play out in the new environment.

## 5. Conclusion

It is unclear exactly how matters would have unfolded in the counterfactual scenario in which Mittermaier’s dissertation was published shortly after he wrote it. The ideas he developed at the time, such as his distinction between *ex ante* and *ex post* facts, clearly had the potential to influence the direction of subsequent work in economics and related fields. As we now know, this possibility was cut off due to the vagaries of academic life and the circumstances surrounding Mittermaier’s academic career. Nevertheless, what the discussion in this article has demonstrated is that there are areas of enquiry in philosophy, economics and social science that are consistent with his

---

<sup>4</sup> The full augmented DAGs can be seen at Pearl & Barenboim 2015, 587.

work and that would likely have been enriched had they been able to exploit Mittermaier's distinction and his definitions of causal and structural facts as an additional resource. How this would have changed or affected these fields is anybody's guess. What is clear is that far from being an antediluvian relic of a bygone era of theorising, Mittermaier's distinction between orders of fact and his advice about using *ex ante* facts for prediction would have fitted comfortably into contemporary discussion on these matters. In this, as in other respects, the theoretical ground covered in *A Realist Philosophy of Economics* is still relevant and definitely worth our considered attention today.

## References

- Campbell, D. 1957. "Factors Relevant to the Validity of Experiments in Social Settings." *Psychological Bulletin* 54 (4): 297–312.
- Cartwright, N. 1999. *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Cartwright, N. 2007. *Hunting Causes and Using Them*. Cambridge: Cambridge University Press.
- Dörge, F. C. and M. Holweger. 2021. "The Non-Existence of Institutional Facts." *Synthese* 199 (3): 4953–4974.
- Guala, F. 2003. "Experimental Localism and External Validity." *Philosophy of Science* 70 (5): 1195–1205.
- Guala, F. 2005. *The Methodology of Experimental Economics*. New York: Cambridge University Press.
- Guala, F. 2012. "Experimentation in Economics." In: *Handbook of the Philosophy of Science*, edited by U. Mäki, 597–640. Amsterdam: Elsevier.
- Heckman, J. and R. Pinto. 2015. "Causal Analysis after Haavelmo." *Econometric Theory* 31 (1): 115–151.
- Mittermaier, K. 2023. *A Realist Philosophy of Economics*. Bristol, UK: Bristol University Press.
- Moural, J. 2002. "Searle's Theory of Institutional Facts: A Program of Critical Revision." In: *Speech Acts, Mind, and Social Reality*, edited by G. Grewendorf and G. Meggle, 271–286. Dordrecht: Springer.
- Pearl, J. (2000) 2009. *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press.
- Pearl, J. and E. Bareinboim. 2011. "Transportability of Causal and Statistical Relations: A Formal Approach." *Proceedings of the Twenty-Fifth National Conference on Artificial Intelligence (AAAI 2011)*, 247–54. Accessed 4 Feb 2025. <https://cdn.aaai.org/ojs/7861/7861-13-11389-1-2-20201228.pdf>.
- Pearl, J. and E. Bareinboim. 2013. "A General Algorithm for Deciding Transportability of Experimental Results." *Journal of Causal Inference* 1 (1): 107–134.
- Pearl, J. and E. Bareinboim. 2015. "External Validity: From Do-Calculus to Transportability Across Populations." *Statistical Science* 29 (4): 579–595.

- Rubin, D. 2005. "Causal Inference Using Potential Outcomes." *Journal of the American Statistical Association* 100 (469): 322–331.
- Searle, J. 1995. *The Construction of Social Reality*. New York: Free Press.
- Searle, J. 2005. "What Is an Institution?" *Journal of Institutional Economics* 1 (1): 1–22.
- Steel, D. 2008. *Across the Boundaries: Extrapolation in Biology and Social Science*. Oxford: Oxford University Press.