Norwegian University of Life Sciences
School of Economics and Business

Causal Insights from Failure

Post-Marketing Risk Assessment of Drugs as a Way to Uncover Causal Mechanisms

Elena Rocca, Rani Lill Anjum and Stephen Mumford

Working Papers No. 6/ 2018

ISSN: 2464-1561

This manuscript is a forthcoming publication.

Please cite as: Rocca E, Anjum RL, Mumford S. Post-marketing risk assessment of drugs as a way to uncover causal mechanisms. In: La Caze A, Osimani B, editors. Uncertainty in Pharmacology: Epistemology, Methods and Decisions. Boston Series for the History and Philosophy of Science: Springer; forthcoming: 2018.

Causal Insights from Failure

Post-Marketing Risk Assessment of Drugs as a Way to Uncover Causal Mechanisms

Elena Rocca, Rani Lill Anjum and Stephen Mumford

Abstract

Pharmacology is as much about establishing causation as it is about understanding it. To establish that a certain treatment causes a positive outcome is only half the task. One needs to be able to make predictions about potential harms and benefits for the single patient. For this, deep causal knowledge is crucial. That is, one needs to understand the causal mechanism underlying a certain outcome. Here we argue that, while the repeated corroboration of the same causal hypotheses through experimentation is useful to establish causation, it is little valuable to advance its mechanistic understanding. We base our argument on a particular philosophical framework, causal dispositionalism, which urges that causes should be understood as complex, intrinsic, tendential and context-sensitive. Applied to pharmacology, this suggests that it is by digging into contextual influence, and not by eliminating it, that we can learn more about the how and why the intervention does its causal work. When a treatment fails to give the expected outcome, therefore, it offers an opportunity to investigate the local context of failure and identify possible interferers. Potentially, this helps uncovering more of the causal nexus by which the outcome is produced. Both pre-clinical research and clinical experimentation alone are poorly fit for uncovering causal mechanisms through failure, since they are based on screening off, or disregard, of interferers. This leaves the post-marketing drug monitoring as the best scenario for systematically generating mechanistic hypotheses about a treatment through studying instances of causal failure of treatment in individuals. We suggest an integrated framework in which post market monitoring, through the study of treatment failure, feeds pre-clinical and clinical research with mechanistic hypothesis.

Keywords: pharmacovigilance, causal mechanism, causal dispositionalism, treatment failure, clinical predictions.

1. Deep causal knowledge for pharmacology

Pharmacology is as much about understanding causation as it is about establishing it. Accordingly, the discovery that a certain drug causes a positive outcome is no more than half the task. One needs also to understand the potential harms that might accompany the benefits. We will be considerably aided in this aim if we have, what we will call, deep causal knowledge. Deep causal knowledge consists in more than simply knowing *that* one thing caused another: it involves also knowledge of *how*.

To justify our stance, we use the ontological framework of causal dispositionalism (developed in detail in Mumford and Anjum 2011) to motivate an explorative approach to causation: one that develops mechanistic knowledge concerning how a pharmacological intervention works. The framework explains why cases of causal failure, if properly acknowledged, provide opportunities for new insights, paving the way for deep causal knowledge. The features of causation emphasised in causal dispositionalism also highlight the importance of post-marketing risk assessment of drugs. This is in part because post-marketing research can contribute to evidence of mechanisms, but also because of the nature of causation itself. A trial might show us that A causes B, but in a context which is often artificial. The context-sensitive, intrinsic, tendential and complex nature of causation explains why we should expect that new information comes to light in the post-marketing phase, especially about off-target effects. Accordingly, we argue in support of a better integration of post-marketing monitoring with the other phases of pharmacological research.

2. What is this thing called causal science?

It will help if we set out some of our basic assumptions and define our main terms, starting with the idea of causal science (with an acknowledgment to Chalmers 1999). We call it causal science because it is the part of science that is concerned with causes: for instance, the discovery of what causes what, or the application of causal knowledge in deciding a medical intervention or developing a new technology. Given that there are multiple ways in which we can discover causal connections, and multiple uses to which we can put that knowledge, it might be better to speak of the causal sciences, plural. Randomised controlled trials (RCTs), for instance, would be a part of causal science, if one interprets this method's purpose as the discovery of effects of a trial intervention. It also seems plausible that causal science is the largest part of science. It might not be the whole of it, but science is at least partly a tool that enables us not just to understand but also to manipulate the world. Central to any such endeavour is the acquisition of causal knowledge. Once we know what tends to follow from what, we are in a position to decide courses of action.

Pharmacology is primarily a causal science. What explains its utility is that it can tell us the effects of drugs, with the ultimate aim of using them to produce benefit and reduce suffering, while minimising harms. We can also, then, use this science to illustrate another important distinction: between deep causal knowledge and surface causal knowledge. It is one thing to know *that* A causes B, for instance. Evidence-based medicine professes to have made significant progress on the correct

methods for finding causal knowledge of this kind (Howick 2011). But, on our classification, to know that an intervention is causally linked to a specific outcome, without knowing anything about how it does so, would count as surface causal knowledge only. Deep causal knowledge, however, requires both knowing-that and knowing-how A causes B, in our terminology. Knowledge-that could for instance be indicated statistically, through covariance or increased relative risk, without the backing of a causal theory. Knowledge-how, on the other hand, would require a theoretical approach. Such an approach would include a range of knowing-thats, but it would not end there.

Clearly knowing *how* A causes B entails knowing *that* A causes B. But the entailment does not automatically run the opposite way. Knowing-that would be a necessary, but not sufficient condition for knowing-how. This shows that deep causal knowledge is more than surface causal knowledge. Now, of course, knowing-how an effect is produced may involve some further knowledge-that; for example, that there is such-and-such a mechanism at work. But this knowledge is deeper than our original knowledge simply that A causes B: the knowing-how consists in knowing all sorts of relevant causal facts and how they relate in order to produce the effect. So such knowledge is deeper knowledge in the sense that to know how A causes B is to know something about the mechanisms that take the causal process from A to B. In the case of pharmacology, this will require an understanding, among other things, of the molecular interactions involved. Typically, this will concern the interaction between the drug and the individuals in the target group, including their genetics and biochemical milieu.

Knowing in one case how the cause led to the effect amounts to knowledge of the mechanism involved, as we think mechanisms should be understood. Conceptions of mechanism vary, of course (see Glennan 1996, 2002, Machamer, Darden and Craver 2000, Machamer 2004, Bechtel and Abrahamsen 2005, Bechtel 2011, Psillos 2011, Illari and Russo 2014: ch. 12). Some understand mechanisms on a reductive basis; that is, they seek explanations of wholes in terms of the activity of their parts. There is also a necessitarian strand to the understanding of mechanisms: that, if functioning correctly and under some specific conditions, they bring about their effects as a matter of necessity. We reject a number of these connotations, as should become clearer in what follows. If we simply say that the mechanism is the process leading from A to B, then we leave open the possibility that explanations can come from a higher level than that which they explain and not just a lower level. This will be important in social explanation, for instance, but it also allows for emergent, top-down causation generally. An obvious candidate for top-down causation in medicine is the placebo effect.

An exclusively reductive explanation could bring a loss of information, since it would fail to account for top-down causation, but also for the interaction of the parts. For instance, two drugs taken together could interact in a synergistic or an antagonistic way, where their effect in combination is not simply the sum of their individual effects, had they been taken separately. In this sense, the whole is not simply the sum of the parts.

3. The nature of causation

We have set out some of our basic concepts that will be useful in what follows. We explained what we meant by causal science and by deep causal knowledge, which amounts to knowledge of mechanisms. We will be arguing for what we call causal failure as a significant opportunity for the acquisition of deep causal knowledge: knowledge showing us the *how*, and not just the knowledge *that*, of causation.

We have further groundwork to lay out, however, for our point is to be made within a particular understanding of causation. What we say is built upon the main claims of causal dispositionalism and its application to scientific method (Mumford and Anjum 2011). Such a view accepts an ontology of dispositions, or real causal powers (developed in Harré and Madden 1975, Mumford 1998, Molnar 2003, Cartwright 1989, Martin 2008). It is not necessary to embrace every detail of this metaphysics in order to see the appeal of the features of causation to which we attend. In any case, those who support the ontology often disagree over the details. However, that causation works in a tendential and context-sensitive way, as we argue here, can be motivated by reasons that are independent of the powers ontology, even though we think such an ontology is the best way to understand how all these features are united in causation. We will explain them briefly and then, in later sections, say how they relate to discovery of deep causal knowledge through causal failure in post-marketing pharmacological settings.

a. *Intrinsicality*

We understand causes as powerful properties of things, which we sometimes call dispositions, and these manifest themselves in their effects. Crucially, this is a non-relational view of causation that stands in opposition to accounts that make it extrinsic. In which accounts would it be extrinsic? A regularity theory of causation (Hume 1739, Psillos 2002), for instance, claims that particular event *a* causes particular event *b* only if there is a general pattern in which every instance like *a* is followed by an instance like *b*. In other words, there could not be a unique case of causation. The defender of intrinsicality rejects this as counterintuitive, and instead supports a form of causal singularism. For a singularist, the fact of whether *a* causes *b* is not ontologically dependent on what happens at other times and places. There could, then, be a case of causation that only happened once in the history of the universe. If a person takes a pill and then experiences an allergic episode, there is no ontological requirement, therefore, in order to say that the pill caused the episode, that other people took the same kind of pill and had a similar allergic reaction. What happens to other people, at other times and places, is instead a matter of the epistemology of causation.

Taking causation as intrinsic and singular is important in the medical context. David Hume, in his extrinsic account of causation, was committed to the principle that 'like objects, plac'd in like circumstances, will always produce like effects' (Hume 1739: I, iii, 8, p. 105). The singularist is not logically committed to Hume's principle. Under singularism, it does not matter whether there are any such 'like circumstances' in order to claim causation. Arguably, there are no 'like circumstances' in biology, since no two individuals are identical. When it comes to pharmacological cases, there seems to be an acknowledgment that this is prima facie the case; namely that there can be individual variation and even medical uniqueness (Smith et al. 2012). Indeed, that the same pharmacological intervention can have different effects in different people is acknowledged within

the structure of trial design. It is the reason why RCTs are conducted on at least two randomised large groups. There is an acceptance that not everyone in each group will have identical responses but, instead, with large-enough groups we can gather trend data on incidences of outcome.

b. *Tendencies*

A singularist still has to account for why there is a degree of order and stability in the world. Unaccompanied, singularism is consistent with a cause having any kind of effect at all: that is, anything could follow (be caused by) anything. This does not seem to describe our world, which is regular enough and predictable enough for science to be a worthwhile enterprise. But in rejecting Hume's principle, above, that like objects will always produce like effects, we should not find ourselves defending the equally absurd principle that like objects could produce *any* effect.

Causal dispositionalism tells us that a cause tends towards a particular type of effect, where *tendency* is to be considered an irreducible modality that is less than necessity but more than pure contingency (Mumford and Anjum 2011: ch. 8). A cause need not result in the same effect always, therefore, but will tend to produce some types of effect more than others. We can see that pharmacological interventions prima facie work this way. For example, SSRIs tend to improve depression but they don't in every case. The taking of an SSRI bears a more than contingent connection with a lessening of depression: it will tend to produce that outcome. Such a tendency can be adequate grounds for prescribing such an intervention even if there is no guarantee of the desired outcome.

Tendencies come in different degrees of strength. For example, oral contraception is causally linked both to prevention of ovulation (strongly) and producing venous thromboembolism (weakly). The former happens in >99% of cases, while the latter in around 15 in 100,000 cases. A weak tendency can still be a very important matter if it is towards a drastic outcome; for example, we should aim to thwart even a weak tendency towards death.

There are reasons why causes do not necessitate their effects. The simplest is that causes can be counteracted and prevented from manifesting their effects by other causal factors. For instance, oral fexofenadine tends to an antihistaminic effect, but this tendency is counteracted by grapefruit juice, which reduces its bioavailability (Bailey 2010). A deeper reason is that causes simply work in this dispositional kind of way. A radioactive radium atom has a tendency to decay by a certain time, specified by its half-life. But it only has a disposition to do so. There is no necessity that it does. Indeed, the specification is in the form of a half-life because that is the point at which there is a 50/50 chance of decay having occurred. If it fails to decay, it is not because anything stopped it, under at least one plausible interpretation. It is simply that its tendency did not manifest. A similar example from biology could be the disposition of ion channels to be either open or closed. Whether an individual channel is open or not is dependent on the presence or absence of a drug but only in a

way that can be described in a dose/response curve for opening/closing. Whether or not it does is not, for all we know, a matter of necessity.[1]

The ontological commitment to tendencies has some empirical basis. It is the way we experience the world as working. Indeed, in our own actions, we know that efforts to achieve a certain outcome are never guaranteed to succeed. But it is still worth the while because our actions will at least tend towards those outcomes. The commitment also fills any gap left when we reject a regularity theory of causation in favour of singularism. Singularism does not have to mean that anything goes.

c. *Complexity*

Dispositionalism accepts complexity as essential to the proper understanding of causation. A complexity view is also a rejection of monocausality: the idea that an effect has just one cause. Instead, we argue that an effect typically occurs when a number of different causes form a mutual manifestation partnership (see Martin 2008), which produces something together than none of the component causes could have produced alone. Ampicillin tends to cure staphylococcal infections, for instance, but only in interaction with a non-resistant bacterium and proper conditions for being absorbed by the patient's body and being distributed to the site of the infection. The drug in isolation can do no causal work. Another example is how some hallucinogenic drugs (e.g. N,N-dimethyltryptamine) are inactive when ingested, but once combined with a trigger (e.g. beta-carbolines) they become bioavailable. On a dispositionalist account, both these substances should count as causes of the hallucinogenic effect.

Despite the implausibility of monocausality, philosophical theories often begin with an attempt at understanding what it is to be a cause (singular), abstracted from the world's complexity. Such an abstraction is relatively harmless if complex causes are nothing but the additions of single causes. However, this is often not the case; for instance, in cases of synergistic and antagonistic interactions. In contrast, philosophical analysis frequently works by abstracting a cause away from the messy reality in which it is located, to see how a 'pure' or 'perfect' instance works, and using that insight to understand the regular, impure cases.

This is a dubious approach in the case of causation, however. For one thing, we have already explained how causes can prevent each other from manifesting. Unless we see that a number of different factors are usually in play, we may fail to see that causes only tend towards their effects rather than guarantee them. But there is a further and even more important reason why we have to acknowledge complexity, which, again, we have encountered. Effects do not simply add together. They often combine in nonlinear and sometimes emergent ways. Emergence occurs where causes that are brought together interact, change each other, and in so doing form a differently empowered whole (for the details on our account of emergence, see Anjum and Mumford 2017). Chemical bonding typically works in nonlinear and emergent ways with respect to causal powers, where the molecule can have very different properties and powers from each element taken separately. For instance, the power of water to put out fires is an emergent power, which is formed

[1] Thanks to Jeffrey Aronson for this example.

when its parts undergo a chemical bonding. None of the parts – hydrogen and oxygen atoms – individually has a power to put out fires. Indeed, they would do the opposite. A pharmacological counterpart to this example is benorylate, which is a chemical compound of aspirin and paracetamol. But unlike aspirin, this compound does not cause gastric erosion, because the chemical bond remains stable while in the gut (Wright 1976).

Now it might be thought that the idea that powers can change each other when they interact looks incompatible with a commitment to intrinsic tendencies. But these are not inconsistent. There is nothing to say that intrinsic features cannot change through their causal interactions. Nor should the facts of emergence lead one to doubt that those intrinsic tendencies are there in the components. They all contribute something of their own towards the interaction, which might not have occurred if any of the contributors were absent.

Causal complexity, in particular nonlinear interactions (such as synergisms and antagonisms), is a crucial feature in pharmacology, and, more generally, in chemistry. In this sense, causal dispositionalism grounds the view that knowledge of a substance's full causal potential can never be exhausted, since it depends on the near-infinite combinations of other entities that it might encounter (Ruthenberg 2016). Accordingly, complexity also represents a well-known challenge concerning risk, in the so-called cocktail effect. The method of RCT tests one intervention at a time and pronounces on its singular effect. However, when one or more drugs work in combination with a particular individual, new causal powers or side effects could emerge. These effects might only come to light in the post-marketing phase, however, once drugs are used in combinations that have never before been tested.

d. *Context-sensitivity*

The final feature of causation to which we draw attention is context-sensitivity, which might be considered the natural consequence of the three previous features, already outlined. The current point tells us that a causal intervention is capable of producing different effects according to the different mutual manifestation partnerships into which it enters. In the case of pharmacology, we can say that different contexts will contribute to bring out different causal powers of a drug.

For example, antiarrhythmic drugs can calm irregular heartbeat (by fast sodium channel inhibition) but in some patients, in a different context, they can worsen it (by potassium channel inhibition) (Smith et al. 2012). Context-sensitivity might be taken as the single notion encapsulating the idea that causes enter into complex partnerships, mutually manifesting their intrinsic tendencies or, in some cases, their joint tendency. The upshot is that we cannot expect the same effect always to come from the same cause. For each person with whom the drug interacts, different manifestation partners are involved, many of which are unique to that person.

Acceptance of the context-sensitivity of causation is a rejection of the epistemic basis of the regularity theory. Hume talks as if we infer causal connections from constant conjunctions of cause and effect. And this assumption is accepted by many of Hume's opponents, who think that in addition to the constant conjunction of cause and effect, causation also involved the necessitation of

the effect by the cause. We can note that this necessitarian claim is built on, and indeed requires, the principle of *same cause, same effect*. But regularity is not our experience, as most empiricists recognise. Instead of moving to a more tendential account that acknowledges causation's context-sensitivity, however, there are usually attempts to bypass it. One strategy is to think of the cause as larger, including also the background conditions under which the effect would be produced. We call this strategy 'causal expansion'. Another strategy is to shield off the cause from possible interferers, mainly by abstraction. We call this strategy 'causal isolation'. Both these are attempts to get a situation where we can assume that the effect is guaranteed to follow the cause under otherwise like circumstances. Indeed, we saw with Hume's statement of the *same cause, same effect* principle, that it was supposed to apply only to causes in the same circumstances. Presumably, then, causes in exactly the same situation would result in the same.

There is a problem with the assumption that causes are conjoined with their effects in this way, however. The only way to demonstrate it empirically would be to create identical circumstances and check whether the same cause would produce the same effect. But this is not practically possible, even in lab settings. Acknowledging that each patient is different, one will never meet the requirements of identical conditions in medicine or pharmacology, since the patient is part of the causal interaction (as a mutual manifestation partner). What is left, then, for the causal necessitarian, is to make a thought experiment or a theoretical model in which one did have identical conditions. For instance, one might infer that if the patient had an identical twin, with the same diet, lifestyle, history, and so on, then if both got the same intervention, the same effect would be guaranteed. But since this is a theoretical presumption, and not something that can be demonstrated empirically, we can only conclude that causal necessitation is an assumption about the nature of causation. Such necessity is thus philosophically motivated and based on an abstraction from our experience, rather than an accurate report of it. Causal necessitarianism – the view that the way in which a cause produces its effect is by necessitating it – has a long tradition in philosophy. It is supported, for instance, by Aristotle (*Metaphysics* Θ 5), Spinoza (1677: I, axiom III), Kant (1781: II, ii, 3, second analogy), Bergson (1889: 199), Ducasse (1924: 55) and many others (for a detailed discussion, see Mumford and Anjum 2011: chapter 3.) It can be found even in Mackie's (1980: 62) INUS-condition account, which might otherwise be thought to allow much of what we want. Even here, however, an effect is produced only when a collection of factors in jointly *sufficient* for the effect, which is to say necessary.

There are thus two consequences from assuming a tendential view of causation, as we promote here. One is that the same cause does not have to produce the same effect under some alleged like circumstances. This is what it means to say that the cause *tends* and no more than tends towards its effect. The second conclusion is that the same effect does not require that the circumstances were alike. We see this in pharmacology, where the same effect can occasionally be produced in very different patients. Accordingly, there is not a one-to-one relationship between cause and effect in dispositionalism, even if the cause includes a large set of background conditions.

4. The dispositional nature of mechanisms

We said that mechanistic understanding is a requirement for what we called deep causal knowledge. But what is a mechanism? One could say that a causal mechanism is what takes the cause to the effect: a process between the intervention and the outcome. In its simplest form, therefore, one could understand a mechanism as something that is added between the intervention and the outcome; such as when the drug D causes its effect E via the mechanism M:

D (→ M) → E

This would be both causal knowledge-*that* and causal knowledge-*how*. For instance, ibuprofen reduces inflammation by inhibiting the COX2 enzyme and synthesis of inflammatory prostaglandins:

Ibuprofen (→ COX2 → inh prostaglandins) → inh Inflammation

However, if we think of mechanisms in the *dispositional* sense – as tendential, complex and context-sensitive – then this cannot be the full causal picture. Understood in this way, the arrows are just tendencies. This means that they can be counteracted or amplified by contextual interferers. In addition to the basic mechanism, we thus need to understand also which factors could interact with the causal process, including how the patient contributes to the outcome. For instance, we would want to know *how* ibuprofen could provoke problems with the cardiovascular system. Moreover, higher-level factors such as stress, anxiety and expectations can affect any pharmacological treatment. Because of the many potential mutual manifestation partnerships, a drug is capable of far more variety of action and different causal interactions than what is described by the basic model of mechanism.

From the dispositionalist perspective, we cannot learn everything about mechanisms simply by many repetitions of the effect. There is a diminishing return on positive evidence: a new corroboration can increase our confidence in the hypothesis further, but beyond a certain point, the extent to which confidence is increased is less than the previous instance. In contrast, cases of causal failure, when something unexpected happens – or there is a discrepancy – could show us where to look to expand our causal understanding. Potentially, causal failure tells us that there is something about a mechanism that we did not understand, that it was capable of another causal pathway, partnered within a certain context, towards a different result than those previously known. We might also uncover more of the causal mechanism of the drug itself: some previously ignored aspect of the molecule, for example, or reaction to the molecule.

Pre-clinical research and clinical experimentation are poorly fit for uncovering causal mechanisms through failure, since they are based on isolation from causal interferers. This leaves the post-marketing drug monitoring as a valuable source of systematically generated mechanistic hypotheses

through studying instances of failure of treatment in individuals. Furthermore, discovery through post-marketing causal failure has empirical considerations in its favour, as we will now go on to show. The approach that we present here, although starting from an ontology of causal dispositionalism, is in line with work by those who promote the epistemological value of single case reports for pharmacovigilance (see for instance Aronson 2003, Aronson 2005, Hauben and Aronson 2006, Karimi et al. 2014, Härmark et al. 2016).

5. The knowledge potential of post-marketing monitoring

Recently, voices have been raised that vindicate the importance of causal irregularities and failures for the process of knowledge-building in the fields of pharmacological research (London and Kimmelman 2015, Ioannidis 2016). These authors encourage the scientific community to abandon the traditional designation of 'negative results' in favour of the more constructive concept of 'unrecognised opportunities' (London and Kimmelman 2015: 28). Such a shift, they argue, will result in the advance of mechanistic understanding, basic research, translation and ultimately healthcare in general. Interestingly, the argument starts from considerations that strike us as close to the dispositional view:

> …drugs alone are not therapeutic agents. Drugs on their own are substances […] that have the *capacity* to interact with particular physiological systems. When they are used in the absence of the knowledge and capacity for applying them properly, they are likely to be toxic. (London and Kimmelman 2015: 29, our emphasis)

Capacities, or dispositions, are intrinsic properties that make something causally powerful. Drugs are effective precisely because they have a causal power toward a curative outcome. However, the therapeutic effect is complex: it emerges from the mutual manifestation of such dispositions within a specific context, namely the patient. A thorough knowledge of the context of a drug's action, therefore, maximises benefits and minimises harms of the treatment. Confirmatory research about the targeted effect of the drug does not alone fulfil this need. An analysis of the patient's context might sometimes start from confirmatory research, for instance by analysis of covariates statistically associated to a certain outcome (see Mannsmann, this volume). However, this approach alone cannot tackle the issue thoroughly, since confirmatory studies work with stringent inclusion criteria and therefore offer information only about some precise patient groups (Rothwell 2005).

After almost 60 years under the shadow of the thalidomide disaster, the link between clinical monitoring of drugs and basic medical understanding is undeniable. Before thalidomide was proved able to interfere with normal foetal development, the mechanisms by which the drug could pass the placental barrier were not accounted for or even considered (McBride 1977). Yet, basic medical knowledge is still largely associated with experimental isolation. When basic researchers investigate real-world complexity instead of lab isolation, however, their efforts do not go unnoticed. An example of this is a study that joined data from the post-marketing monitoring of a large number of

marketed drugs, and used the information to gain understanding of the mechanism by which drugs interact with human molecules (Campillos et al. 2008). In particular, they grouped drugs according to similarity in their spectra of undesired effects, and used this comparison to detect common molecular targets. Through this approach, 20 possible unexpected drug-drug relations were identified, between drugs with dissimilar structures and different therapeutic indications. The authors emphasised that their hypotheses could be formulated only by observing the complexity of interactions between drugs and the human body, that is, outside experimental isolation. Their approach, they suggested, hinted at 'new uses of marketed drugs' (Campillo et al. 2008: abstract).

Once we accept that information from clinical failures is valuable for basic research, more questions follow. Are all types of information equally useful for the purpose of deep causal knowledge? No, we argue. As suggested above, evidence of failure, unlike repeated corroboration of positive effects, can lead to the *how* of causation, as long as it also provides detailed information about the specific causal processes, and of how they can be counteracted.

This view reflects well the reality of post-marketing monitoring of drugs. Large observational studies are useful to give an overview of a drug's impact on the population. Certainly, such surface causal knowledge can potentially improve clinical use of some drugs, for example by awareness of risk factors (or susceptibility factors), and by identifying the reference class to which the patient belongs. Optimising the impact of a treatment for the single patient, however, requires the deepening of mechanistic understanding. Susceptibility to rare adverse reactions, for instance, is hard to predict with population data. This is why genetic tests are being developed in order to make precise individual predictions (see for instance the detection of HLA B*5701 polymorphism as a predictor of allergic reactions to abacavir in the treatment of HIV/AIDS). Clearly, genetic tests require that we understand the molecular underpinning of the adverse reaction. A further reflection is that the mere knowledge of patient's risk factors might inform us about the cases in which a treatment should be avoided, but does not easily serve the purpose of optimising it. Some vaccination programs, for instance, fail to work in obese patients (Wiedermann et al. 2016). However, without the understanding of the molecular processes underlying this failure, it is hard to find a way to overcome the problem.

While evidence of mechanisms can come from a variety of sources, it is easier to formulate informed hypotheses from details of specific clinical cases than from frequentist data alone. We can illustrate this with an example.

Valproic acid was marketed in the late 70s as an anti-epileptic drug. Soon after, its use was broadened to the treatment of migraine, bipolar affective disorder, alcohol withdrawal and more (Sitarz et al. 2014). Although this drug is usually well tolerated, it can cause a range of important adverse reactions, including rare but potentially deadly liver injury. In 1984, a first report on liver damage associated with valproic acid counted 42 cases with fatal outcomes, of which most were paediatric cases and cases of polypharmacy (Powell-Jackson et al. 1984). The same report hinted at hypotheses for the mechanism of harm, together with some evidence in lab animals, but the total evidence of mechanism at the time was inconclusive. Two years later, the incidence of hepatic fatalities had decreased nearly fivefold, despite an overall increase in the use of the drug (Dreifuss 1989). This improvement was probably due to awareness among prescribers, and decreased

prescription in high-risk patients, such as young children, because of the availability of population data. Despite this improvement, fatal liver injury persisted. In 1996, 29 additional cases were reported (Bryant and Dreifuss 1996). Since the causes of such reactions were still mysterious, it was extremely difficult to predict a fatal outcome in an otherwise healthy patient.

A new development in the mechanistic knowledge occurred when detailed evidence about valproic acid's victims began to be gathered. Histological analyses showed some specific modifications, suggesting that the principal cause of harm started from interference with mitochondrial processes (Fromenty and Pessayre 1995). A number of observations confirmed the fact that valproic acid affects mitochondrial functions. For instance, a detailed analysis of urine composition in treated children, compared with untreated children, revealed that the drug influences mitochondrial metabolism (Price et al. 2011). Anecdotal reports pointed out that valproic acid-induced liver failure was fulminant in patients with genetic mitochondrial diseases (Krahenbuhl 2000). Ultimately, detailed molecular characterization of valproic acid's victims revealed that the risk of liver toxicity is determined by point mutations in the gene POLG, which codes for a mitochondrial protein (DNA polymerase γ) (Stewart et al. 2010).

We see, then, that different contexts of interaction brought out new causal powers of the drug. Even if the mutual manifestation partnerings for some of those effects are rare, they are nevertheless valuable for understanding how the drug does its causal works. Such a precise determination of the mechanism of harm is valuable for two reasons. First, the patient's tendency to be harmed by valproic acid might be revealed in advance through genetic screening. Secondly, the discovery triggered a new line of research aimed at a deeper causal understanding: why do patient with POLG mutations have a tendency to develop valproic acid-induced liver injury? This question was addressed, for instance, by culturing cells from patients carrying a mutation of the POLG gene, and observing how these cells react to valproic acid administration (Sitarz et al. 2014).

The example illustrates well the knowledge potential of post-marketing monitoring. In pharmacology, patterns of harm detection are intermingled with patterns of causal discovery. Such mingling is facilitated by evidence of causal failure.

6. How much failure do we need?

Having more data is usually thought to provide better evidence of causation, which is why quantitative research is given epistemological priority over for instance case studies. The question is what we mean by more data. More of what sort of data? And how much is enough?

Kimmelman et al. (2014) argue for the necessity of distinguishing two phases of pre-clinical research: exploratory and confirmatory. The first is focused on generating hypotheses, and the second on testing them through robust evidence. The authors point out that the two stages should use different and complementary types of study, and urge the need of a 'greater volume of confirmatory investigation'.

We maintain that this is a useful distinction, and that the process of corroborating a hypothesis needs a large amount of experimental confirmation. On the other hand, the exploratory phase might

benefit more from failure than from confirmation. Think for instance about the role of serendipitous discoveries in the history of science.[2] The failure of reproducing the expected result can generate new hypotheses, and expand knowledge far beyond what a mere confirmation would have allowed. However, if we need a reasonable number of confirmations in order to establish knowledge, how much failure do we need in order to explore it? For cases of serendipity, the success of a discovery seems to rely greatly on 'catching' the unexpected as soon as it is manifested. The key for discovery is thus the quality, rather than the frequency, of an unexpected observation. This is a general truth for all causal science, and pharmacology is no exception.

Unexpected effects of drugs might be useful for research, even if they do not happen frequently. Since causation is intrinsic and singular, as described above, we cannot deny causation only because it happens too rarely. If the drug has causal powers that almost never manifest themselves, because the mutual manifestation partner is rare, it is nevertheless a power of the drug that ought to be considered a real risk. For example, Zolpidem, a GABAergic hypnotic used to treat insomnia, has provided the medical literature with a number of surprising anecdotes (Rocca 2016). Patients under its effect were reported to sleepwalk, sleep-cook, sleep-eat and sleep-drive. But also, some patients suddenly recovered after stroke, re-gained mobility after brain injury or improved from post-traumatic states shortly after taking the drug (Hoque and Chesson 2009). Although being only circumstantial, and unsuitable for a frequentist analysis, these cases were not ignored by basic research. Quite the contrary. They led to the hypothesis that the brain might have some redundant patterns of self-recovery, which Zolpidem could unmask.

Unexpected effects of drugs can thus function as a springboard for basic research, even when they are very rare (or maybe even because of this). We will now explain why we say that 'more data' ought to mean more data about the *single* cases of causal failure, and then propose a practical way forward for systematically collecting and using such data.

7. Reporting suspected causal failure: a community effort

The problem of improving knowledge translation has gained increasing attention in the medical literature over the last years (Ioannidis 2016, Edwards and Hugman 1997). Following the inefficiency of pre-clinical and clinical research for the ultimate improvement of healthcare, legitimate questions have been asked about mistakes in the information flow from research to clinical practice. This step in the process is often called bench-to-bedside, or knowledge-to-action translation.

However, these concerns cover only half the problem. Indeed, a community that acknowledges the central role of real-world complexity for the advance of medical research must make sure that clinical observations feed research with hypotheses in the most fruitful ways. In other words, not only knowledge-to-action, but also action-to-knowledge translation must be optimised for improvement of the process as a whole. In the specific case of harm detection, how and when

[2] Thanks to Samantha Copeland for bringing this point to our attention.

should clinical observation of off-target effects occur? And how should it be reported and communicated to best fulfil its potential for advancing knowledge?

In the view of causal dispositionalism, deep causal understanding happens by analysing the individual causal process in a detailed way. In the context of harm detection, this view results in a central focus on the clinical observation of side effects. Recognition of the epistemological value of the individual clinical observation has to be seen as an advantage in the context of post-marketing drug monitoring. Indeed, what can be identified as the systematic trigger of pharmacovigilance, if not clinical assessment of suspected side effects? Note that with this we are not denying that large population studies have an important role in monitoring harm. However, we agree with Edwards when he points out that 'any consideration of causality based only on epidemiological approaches is not logically defensible and certainly not socially acceptable' (Edwards 2012: 51). A population study needs sound hypotheses and a good experimental design in order to be fully reliable. Historical examples show that, in the field of pharmacovigilance, population data should not be produced or interpreted in isolation from other types of evidence, including single suspected clinical diagnoses (Edwards 2012: 42 and 51).

The type and amount of information available about a suspected effect determines which hypotheses can be made about its biological underpinning. In other words, mechanistic hypotheses depend on the quality of the side effect reports. In order to gain a deep causal understanding of harm, beyond the assessment that a certain effect occurred, also including the understanding of how it occurred, it is useful to consider the side effect report from a dispositionalist perspective. That is, the main input of the report is a set of conditions (context) that in combination with the drug provoked a certain outcome (Rocca 2016).

So although some effects are rare, the causal processes leading up to them must nevertheless be considered and understood. In this sense, drug agencies would benefit most from receiving as many details as possible, to facilitate the detection of signals of harm, the performance of comparative analyses, and ultimately the formulation of causal hypotheses (see also Aronson 2011, Gagnier et al. 2014).

Consider 'classic' case reports, which marked the beginning of important lines of medical research. For instance, we can recall the first case-series describing patients affected by AIDS (Gmerk 1990). Clearly, the way in which doctors reported their observations influenced the immediate research hypotheses. Clinicians selected some common features of the patients that they suspected to be somehow causally connected with their condition (e.g. the fact that they were all young and homosexual). However, they did not discard other details, although they had no explanatory backing for including them (e.g. cytomegalovirus infection, pneumonia, use of inhalant drugs like 'poppers', acquaintances). With such details at hand, researchers moved immediately in several directions in order to verify causal relations and discard red herrings. This feature, of containing unfiltered details that at the time of the observation are not organised in a causal picture, is an important epistemological value of clinical reports (for a complete analysis of this case, see Ankeny 2010).

Is this expectation feasible in the reality of clinical life? After all, the doctor's task is to provide healthcare, not to write medical literature. Doctors spend an increasing amount of time with administrative duties, and in the last decades studies confirmed that they dedicate less than half

their working day to patients, resulting in a decline in overall productivity (Røhme and Kjekshus 2001, Sinsky et al. 2016). Not all side effects, therefore, can be reported in detail. In fact, under-reporting is a factor that delays pharmacovigilance (Hazel and Shakir 2006). Studies show that health practitioners tend to select the suspected side effects also by novelty, hindering the possibility of getting a realistic overview of the impact of a treatment on the population (Haramburu et al. 1997). A side effect or failure, indeed, might not be novel in itself, but the description of the context of failure – that is, the particular process – might contain information that is important for a deep causal understanding.

How, then, should clinicians prioritise reporting in a way that is feasible with their routine, if not only based on novelty? This problem becomes manageable once we consider that it does not pertain to practitioners in isolation, but rather to a community of practitioners, carers, patients and researchers. From this perspective, science ought to provide clinicians with the criteria they need for prioritising the cases that at a specific time or place are worth more effort in reporting. The priority should not be determined by strict or permanent rules, but should be adaptable over time and in relation to different practices. It should be adjusted based on the current state of knowledge and uncertainties about the specific pharmacological treatments.

We thus suggest an improvement of *action-to-knowledge* translation based on an integrated system of close collaboration between post-marketing monitoring and all phases of research, as described in Rocca (2016). Pharmacological developments should not be seen as straight patterns starting from basic research and ending with the clinical treatment, as depicted in the standard model of clinical translation. Rather, information should flow in a continuous loop among all phases of research and clinical practice. Some authors have recently exhorted this development (Kimmelman 2015). Their motivation, however, remains improvement of *knowledge-to-action* translation. We extend these considerations to the whole process of information flow: from the lab bench to the patient's side, *and back*, in a continuous loop. Reports of suspected side effect should be written with the intent of feeding back to basic research with hypotheses for deep causal knowledge. For this, statistical analyses of correlations is only a part of the process. The reports should also be analysed by detailed comparisons, with the aim of providing basic research with mechanistic hypotheses of harm (for more about this information loop, see Rocca 2016: figure 1).

If the report of suspected harm should not be produced by a clinician in isolation, but by a clinician as part of an integrated system, two practical considerations follow. First, the production of high quality reports must be promoted actively by the system. A strategy for this could start from medical education, with a shift to a type of communication that is focused on patients and their context. Models for medical interview based on bio-psychosocial, patient-centred communication have already been developed and are being taught in some medical schools (Larivaara et al. 2001). Second, the side effect report can be seen as a common effort of one or more practitioners together with the patient. It could be the result of an integrated network of clinical data collected about the patient in several instances and by different practitioners (general practitioner, specialist, physiotherapist, pharmacist, surgeon, et cetera). The utility of data from health registries for statistical population studies is at the moment still under debate (Grimes 2015, Norén et al. 2010). Some see it as a goldmine for statistical analysis, while others are sceptical of using data that were not collected following a specific design and research question. However, a network of registries can

surely give more thorough and detailed information about *single* cases, and be highly valuable for exploratory research. Confirmation belongs to another stage of the loop.

8. Conclusion

While arguing that we need better systems in place to learn from causal failure, we do not, of course, encourage failure itself. In pharmacological cases, such failures can be catastrophic. But when they do occur, we must acknowledge them and adjust our theories accordingly. There seems to be no point in protesting that current approaches to drug safety, based upon the large-scale statistical techniques of evidence based medicine, are adequate as they stand. Adverse drug reactions constitute, on some accounts, the fourth biggest cause of death in the United States (FDA Report 2016). Not all of these deaths are down to basic errors, such as dosage mistakes. Rather, it seems that something fundamental needs improving in how we assess and understand the effects of drugs.

We have provided a philosophical framework that explains some of the difficulties: causes should be understood as complex, intrinsic, tendential and context-sensitive. Applied to drug use, this means that the post-marketing phase represents the reality better than the methods seeking to isolate causes in artificial conditions. Real-life contexts of causal failures – unexpected results – is where new hypotheses can be developed and, with due consideration, deep mechanistic understanding can grow.

Acknowledgements

Thanks to Jeffrey Aronson, Ralph Edwards and Barbara Osimani for useful comments and feedback on the paper. Some of the examples used were suggested to us by Jeffrey Aronson. The research in this paper was conducted with the support of the Norwegian Council for Research (NFR) FRIPRO scheme.

References

Anjum, R. L. and Mumford, S. (2017) ‘Emergence and Demergence’, in M. Paoletti and F. Orilia (eds), *Philosophical and Scientific Perspectives on Downward Causation*, London: Routledge, pp. 92-109.

Ankeny, R. (2011) ‘Using Cases to Establish Novel Diagnoses: Creating Generic Facts by Making Particular Facts Travel Together’, in P. Howlett and M. Morgan (eds), *How Well Do Facts Travel? The Dissemination of Reliable Knowledge*, Cambridge: Cambridge University Press, pp. 252-72.

Aristotle, *Metaphysics*, H. Lawson-Tancred (trans.), London: Penguin, 1998.

Aronson, J. K. (2011) 'Guidelines and a Checklist for Reporting Suspected Adverse Drug Reactions Anecdotally in Journals', in J. Talbot and J. K. Aronson (eds), *Stephens' Detection and Evaluation of Adverse Drug Reactions: Principles and Practice*, 6th edn, Oxford: Wiley-Blackwell: 699-708.

Aronson, J. K. (2003) 'Anecdotes as Evidence', *British Medical Journal*, 326: 1346.

Aronson, J. K. (2005) 'Unity from Diversity: the Evidential Use of Anecdotal Reports of Adverse Drug Reactions and Interactions', *Journal of Evaluation in Clinical Practice*, 11: 195-208.

Bailey, D. G. (2010) 'Fruit Juice Inhibition of Uptake Transport: a New Type of Food-Drug Interaction', *British Journal of Clinical Pharmacology*, 70: 645-55.

Bechtel, W. (2011) 'Mechanism and Biological Explanation', *Philosophy of Science*, 78: 533-57.

Bechtel, W. and Abrahamsen, A. (2005) 'Explanation: a Mechanist Alternative', *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 36: 421-41.

Bergson, H. (1889) *Time and Free Will: An Essay on the Immediate Data of Consciousness*, F.L. Pogson (trans.), New York: Dover Publications Inc., 2001.

Bryant, A. E. and Dreifuss, F. E. (1996) 'Valproic Acid Hepatic Fatalities. III. US Experience Since 1986', *Neurology*, 46: 465-69.

Campillos, M., Kuhn, M., Gavin, A. C., Jensen, L. J. and Bork, P. (2008) 'Drug Target Identification Using Side-effect Similarity', *Science*, 321: 263-66.

Cartwright, N. (1989) *Nature's Capacities and Their Measurement*, Oxford: Oxford University Press.

Chalmers, A. (1999) *What is This Thing Called Science?*, 3rd edn, Maidenhead: Open University Press.

Dreifuss, F. E., Langer, D. H., Moline, K. A. and Maxwell, J. E. (1989) 'Valproic Acid Hepatic Fatalities. II. US Experience Since 1984', *Neurology*, 39: 201-06.

Ducasse, C. J. (1924) *Causation and the Types of Necessity*, New York: Dover, 1969.

Edwards, I. R. (2012) 'Considerations on Causality in Pharmacovigilance', *International Journal of Risk & Safety in Medicine*, 24: 41-54.

Edwards, I. R. and Hugman, B. (1997) 'The Challenge of Effectively Communicating Risk-Benefit Information', *Drug Safety*, 17: 216-27.

FDA Report (2016) 'Preventable Adverse Drug Reactions: A Focus on Drug Interactions', online: http://www.fda.gov/Drugs/DevelopmentApprovalProcess/DevelopmentResources/DrugInteractionsLabeling/ucm110632.htm (accessed 8.2.17)

Fromenty, B. and Pessayre, D. (1995) 'Inhibition of Mitochondrial Beta-oxidation as a Mechanism of Hepatotoxicity', *Pharmacology & Therapeutics*, 67: 101-54.

Gagnier, J. J., Kienle, G., Altman, D.G., Moher, D., Sox, H., Riley, D.; the CARE Group (2014) 'The CARE Guidelines: Consensus-based Clinical Case Report Guideline Development, *Journal of Clinical Epidemiology*, 67: 46-51.

Glennan, S. (1996) 'Mechanisms and the Nature of Causation', *Erkenntnis*, 44: 49-71.

Glennan, S. (2002) 'Rethinking Mechanistic Explanation', *Philosophy of Science*, 69: S342-53.

Gmerk, M. D. (1990) 'History of AIDS: Emergence and Origin of a Modern Pandemic', R. C. Maulitz and J. Duffin (transl.), Princeton: Princeton University Press.

Grimes, D. A. (2015) 'Epidemiologic Research with Administrative Databases: Red Herrings, False Alarms and Pseudo-epidemics', *Human Reproduction*, 30: 1749-52.

Haramburu, F., Bégaud, B., Moride, Y. (1997) 'Temporal Trends in Spontaneous Reporting of Unlabelled Adverse Drug Reactions', *British Journal of Clinical Pharmacology*, 44: 299-301.

Härmark, L., Raine, J., Leufkens, H., Edwards, I. R., Moretti, U., Sarinic, V. M. and Kant, A. (2016) 'Patient-reported Safety Information: A Renaissance of Pharmacovigilance?', *Drug Safety*, 39: 883-90.

Harré, R. and Madden, E. H. (1975) *Causal Powers: A Theory of Natural Necessity*, Oxford: Blackwell.

Hauben, M. and Aronson, J.K. (2006) 'Gold Standards in Pharmacovigilance: the Use of Definitive Anecdotal Adverse Drug Reactions as Pure Gold and High Grade Ore', *Drug Safety*, 30: 545-55.

Hazell, L., Shakir, S. A. (2006) 'Under-reporting of Adverse Drug Reactions: a Systematic Review, *Drug Safety*, 29: 385-96.

Hoque, R. and Chesson, A. L. (2009) 'Zolpidem-induced Sleep Walking, Sleep Related Eating Disorder, and Sleep-driving: Fluorine-18-Fluorodeoxyglucose Positron Emission Tomography Analysis, and a Literature Review of Other Unexpected Clinical Effects of Zolpidem', *Journal of Clinical Sleep Medicine*, 5: 471-6.

Howick, J. (2011) *The Philosophy of Evidence-Based Medicine*, Chichester: Wiley-Blackwell.

Hume, D. (1739) *A Treatise of Human Nature*, L. A. Selby-Bigge (ed.), Oxford: Clarendon Press, 1888.

Illari, P and Russo, F. (2014) *Causality*, Oxford: Oxford University Press.

Ioannidis, J. (2015) 'Translational Research May Be Most Successful When It Fails', *Hastings Center Report*, 45: 39-40.

Ioannidis, J. P. (2016) 'Why Most Clinical Research is Not Useful', *PLoS Medicine*, 13: e1002049.

Kant, I. (1781) *Critique of Pure Reason*, N. Kemp Smith (trans.), London: MacMillan, 1929.

Karimi, G., Star, K., Lindquist, M. and Edwards, I. R. (2014) 'Clinical Stories are Necessary for Drug Safety', *Clinical Medicine*, 14: 326-7.

Kimmelman, J., Mogil, J. S. and Dirnagl, U. (2014) 'Distinguishing Between Exploratory and Confirmatory Preclinical Research will Improve Translation', *PLoS Biol*, 12: e1001863.

Krähenbühl, S., Brandner, S., Kleinle, S., Liechti, S. and Straumann, D. (2000) 'Mitochondrial Diseases Represent a Risk Factor for Valproate-induced Fulminant Liver Failure', *Liver*, 20: 346-48.

Larivaara, P., Kiuttu, J. and Taanila, A. (2001) 'The Patient-Centred Interview: the Key to Biopsychosocial Diagnosis and Treatment', *Scandinavian Journal of Primary Health Care*, 19: 8-13.

London, A. J., Kimmelman, J. (2015) 'Why Clinical Translation Cannot Succeed without Failure', *Elife*, 4: e12844.

Machamer, P. (2004) 'Activities and Causation: The Metaphysics and Epistemology of Mechanisms', *International Studies in the Philosophy of Science*, 18: 27-39.

Machamer, P., Darden, L. and Craver, C. F. (2000) 'Thinking About Mechanisms', *Philosophy of Science*, 67: 1-25.

Mackie, J. L. (1980) *The Cement of the Universe*, Oxford: Oxford University Press.

Martin, C. B. (2008) *The Mind in Nature*, Oxford: Oxford University Press.

McBride, W. G. (1977) 'Thalidomide Embryopathy', *Teratology*, 16: 79-82.

Molnar, G. (2003) *Powers: A Study in Metaphysics*, S. Mumford (ed.), Oxford: Oxford University Press.

Mumford, S. (1998) *Dispositions*, Oxford University Press.

Mumford, S. and Anjum, R.L. (2011) *Getting Causes from Powers*, Oxford: Oxford University Press.

Norén, G. N., Hopstadius, J., Bate, A., Star, K. and Edwards, I. R. (2010) 'Temporal Pattern Discovery in Longitudinal Electronic Patient Records', *Data Mining and Knowledge Discovery*, 20: 361-87.

Powell-Jackson, P. R., Tredger, J. M. and Williams, R. (1984) 'Hepatotoxicity to Sodium Valproate: a Review', *Gut*, 25: 673-81.

Price, K. E., Pearce, R. E., Garg, U. C., Heese, B. A., Smith, L. D., Sullivan, J. E., Kennedy, M. J., Bale Jr, J. F., Ward, R. M., Chang, T. K. H., Abbott, F. S. and Leeder, J. S. (2011) 'Effects of Valproic Acid on Organic Acid Metabolism in Children: a Metabolic Profiling Study', *Clinical Pharmacology & Therapeutics*, 89: 867-74.

Psillos, S. (2002) *Causation and Explanation*, Chesham: Acumen.

Psillos, S. (2011) 'The Idea of Mechanism', in P. M. Illari, F. Russo and J. Williamson (eds), *Causality in the Sciences*, Oxford: Oxford University Press, pp. 771-88.

Rocca, E. (2016) 'Bridging the Boundaries between Scientists and Clinicians: Mechanistic Hypotheses and Patient Stories in Risk Assessment of Drugs', *Journal of Evaluation in Clinical Practice*, DOI 10.1111/jep.12622

Rothwell, P. M. (2005) 'External Validity of Randomised Controlled Trials: "To Whom do the Results of This Trial Apply?"', *The Lancet*, 365 (9453): 82-93.

Ruthenberg, K. (2016) 'About the Futile Dream of an Entirely Riskless and Fully Effective Remedy: Thalidomide.' *HYLE–International Journal for Philosophy of Chemistry*, 22: 55-77.

Røhme, K. and Kjekshus, L. E. (2001) 'When Time Counts: Physicians' Work Assignments and Use of Time', *Tidsskriftet den Norske Lægeforening*, 121: 1458-61.

Sinsky, C., Colligan, L., Li, L., Prgomet, M., Reynolds, S., Goeders, L., Westbrook, J., Tutty, M. and Blike, G. (2016) 'Allocation of Physician Time in Ambulatory Practice: a Time and Motion Study in 4 Specialties', *Annals of Internal Medicine,* 65: 753-60.

Sitarz, K. S., Elliott, H. R., Karaman, B. S., Relton, C., Chinnery, P. F. and Horvath, R. (2014) 'Valproic Acid Triggers Increased Mitochondrial Biogenesis in POLG-deficient Fibroblasts', *Molecular Genetics and Metabolism*, 112: 57-63.

Smith, S. W., Hauben, M. and Aronson, J. K. (2012) 'Paradoxical and Bidirectional Drug Effects', *Drug Safety*, 35: 173-89.

Spinoza, B. (1677) *The Ethics*, R. H. M. Elwes (ed.), New York: Dover, 1955.

Stewart, J. D., Horvath, R., Baruffini, E., Ferrero, I., Bulst, S., Watkins, P. B., Fontana, R. J., Day, C. P. and Chinnery, P. F. (2010) 'Polymerase γ Gene POLG Determines the Risk of Sodium Valproate-induced Liver Toxicity', *Hepatology*, 52: 1791-96.

Wiedermann, U., Garner-Spitzer, E., Wagner, A. (2016). 'Primary Vaccine Failure to Routine Vaccines: Why and What to Do?*, *Human Vaccines & Immunotherapeutics*, 12: 239-243.

Wright, V. (1976) 'A Review of Benorylate. A New Antirheumatic Drug', *Scandinavian Journal of Rheumatology*, 4 (sup13): 5-8.