Philosophical Problems in Science (2025), **78**, 27–40
(Zagadnienia Filozoficzne w Nauce)
doi:10.59203/zfn.78.709

Philosophical Problems in Science

ARTICLE

Tractability assumptions and derivational robustness, a match made in heaven?

Gareth Fuller

University of Maine at Presque Isle
Email: gareth.fuller@maine.edu

Abstract

The epistemology of modeling in science is a topic that has received no shortage of attention from both philosophers and scientists. Models, with their ubiquitous unrealistic and false assumptions, raise a myriad of interesting epistemic questions. Many have picked up a concern expressed by biologist Richard Levins, which is that, for many models, it is unclear if any particular result is a product of the realistic assumptions, or if the result, is in some critical way, dependent on an unrealistic assumption. His suggestion of a solution to this problem, the use of robustness analysis, has similarly been picked up by many as a plausible solution to this concern.

Since Levins' time, taxonomies of modeling assumptions and types of robustness analyses have been developed. One common connection made is that of tractability assumption, or assumptions introduced for the purpose of mathematical tractability, and derivational robustness analysis, or a robustness analysis that tests the influence of assumptions on the derivation of some result. Derivational robustness analysis is often singled out as being particularly well-suited for resolving the epistemic concerns introduced by tractability assumptions.

In this paper, I argue that tractability assumptions do not present a consistent set of epistemic concerns. Given this, derivational robustness analysis cannot be used to resolve the epistemic concerns raised by all tractability assumptions. I use this to motivate some concerns about how well-suited the taxonomy that includes tractability assumptions is to discussions about the epistemic concerns of unrealistic modeling assumptions.

Keywords: scientific modeling, robustness analysis, tractability assumptions

1. Introduction

The epistemology of modeling in science is a topic that has received no shortage of attention from both philosophers and scientists. Models, with their ubiquitous unrealistic and false assumptions, raise a myriad of interesting epistemic questions. Many have picked up a concern expressed by biologist Richard Levins (1966), which is that, for many models, it is unclear if any particular result is a product of the realistic assumptions, or if the result is in some critical way dependent on an unrealistic assumption. His suggestion of a solution to this problem, the use of robustness analysis, has similarly been picked up by many as a plausible solution to this concern (Weisberg 2006; Weisberg and Reisman 2008; Kuorikoski, Lehtinen, and Marchionni 2010; Raerinne 2013).[1]

Since Levins' paper, taxonomies of types of modeling assumptions have been developed, often distinguishing types of unrealistic assumptions (Musgrave 1981; Mäki 2000; Hindriks 2006; Cartwright 2007; Kuorikoski, Lehtinen, and Marchionni 2010). The taxonomy of modeling assumptions I will work with includes the categories of substantial assumptions, which are realistic assumptions, and then galilean and tractability assumptions, which are unrealistic. Just like modeling assumptions, taxonomies of robustness analyses have also been developed since Levins' paper (Woodward 2006; Weisberg and Reisman 2008; Raerinne 2013). One particular type of robustness analysis, Derivational Robustness Analysis (DR), has been singled out as being particularly useful for resolving concerns with tractability assumptions (Kuorikoski, Lehtinen, and Marchionni 2010; Raerinne 2013). Given this connection between DR and tractability assumptions, I focus on these two specifically.

The main focus of this paper is to argue that understanding when to apply DR should be separated from the taxonomy of modeling assumptions that includes substantial, galilean, and tractability assumptions.[2] Rather, I argue that understanding the kinds of assumptions that DR can play an important epistemic role, particularly that of resolving concerns about unrealistic assumptions, should focus on particular epistemic features of those assumptions. These epistemic features do not align well with the taxonomy of modeling assumptions commonly used and so this taxonomy is not appropriate for understanding the role of DR. I attempt to do this by showing that tractability assumptions do not consistently capture these epistemic features.

My paper will proceed as follows. In section 2 I introduce DR, focusing on the conditions where it can resolve epistemic concerns of particular modeling assumptions. In section 3 I discuss how the literature singles out tractability assumptions as a unique target for DR. In section 4 I look to several examples of tractability assumptions, arguing that they do not fit well with the required epistemic conditions for the use of DR. I also present an example where DR might be applied to other types of assumptions. In section 5 I draw two general conclusions for the ongoing debates about robustness analysis and the epistemology of models. In section 6 I provide some concluding remarks.

2. Strategies for Resolving Some Epistemic Concerns for Models

In his paper *The strategy of model building in population biology*, biologist Richard Levins outlined the ubiquity, and even necessity, of unrealistic "simplifying" assumptions in models (Levins 1966). The inclusion of unrealistic assumptions, even if necessary, raises some epistemic concerns about these models. Levins raised one particular concern, which is that, for many models, "There is always room for doubt as to whether a result depends on the essentials of a model or on the details of the simplifying assumptions" (Levins 1966, p.423). A result that depends on particular details of simplifying assumptions, rather than the realistic ones, gives us reason to question the value of that result or the model. For instance, if we have a model of a ball rolling down an inclined plane, and this model tells us the ball will have certain velocity at the bottom of the ramp, but the model does not include friction, we might question whether or not this result can be applied to real inclined planes. The result may depend on the lack of friction, which would mean the result does not apply to planes with friction. The general concern is that, If some result is dependent on the unrealistic simplifying assumptions, this can give us reason to discount the result or the model itself.

This concern has been raised by many, and one straightforward solution to this problem is to simply replace the unrealistic assumption with a more realistic assumption. In the case of the frictionless plane, it would include adding friction into the model to make the model more realistic. If the result of the model changes, then we know that the result was dependent, in some way, on the lack of friction. This process of making a model more realistic is the process of de-idealization or concretization (Peruzzi and Cevolani 2022).

While de-idealization is straightforward, it has several drawbacks. First, many models include a great many simplifying assumptions, and it could very well be the case that there are simply too many simplifying assumptions to apply this to. Levins nots that, at least with population biology models, to carry out the process of idealization such that the model lacks any simplifying assumptions, would require a model "using perhaps 100 simultaneous partial differential equations" and this would lead to models that would be too insoluble and the results of which would have no real meaning to us (Levins 1966). Further, it is not the case that all assumptions in a model can be de-idealized, since the true vale of some parameters are not known or are vaguely defined (Levins 1966).

Given these concerns about de-idealization, Levins' solution to this problem is a method known as robustness analysis. In robustness analysis,

> We attempt to treat the same problem with several alternative models each with different simplifications but with a common biological assumption. Then, if these models, despite their different assumptions, lead to similar results we have what we can call a robust theorem which is relatively free of the details of the model (Levins 1966, p.423).

Essentially, the strategy of robustness analysis is to develop several models with a shared "core" of realistic assumptions about the causal mechanism of interest, but different unrealistic or simplifying assumptions. If all of these models produce the same result, or a "robust theorem", then this is meant to show that the robust theorem is derived from the shared core assumptions rather than any unrealistic assumption. This avoids the concerns with de-idealization since none of the models used in robustness analysis are completely de-idealized but rather incorporate *different* simplifying assumptions.

De-idealization and robustness analysis differ in several ways. The most obvious one is that de-idealization focuses on creating more and more realistic models, while robustness analysis does not focus on producing more realistic models, but just different, false models. Epistemically, de-idealization works to boost credence in the idealized model because the less realistic model agrees with a more realistic one. Ultimately, however, once the more realistic model has been developed, the less realistic one is no longer needed. When it comes to robustness analysis, the boost in credence stems from the fact that the result of a model does not depend on any particular set of simplifying assumptions, but, rather, seems to be dependent on the shared, realistic, set of assumptions. Each model in the set for robustness analysis is needed.

Since Levins' paper, a variety of types of robustness analysis have been developed (Woodward 2006; Weisberg and Reisman 2008; Raerinne 2013).Woodward has developed an influential taxonomy that includes Inferential Robustness (IR) and Derivational Robustness

(DR) (Woodward 2006). IR is a method applied when there is a data set and its relation to a particular hypothesis. It is often the case that data alone is not enough, but some further assumptions are needed to derive any conclusions. In some cases, it could be that there are several competing sets of further assumptions. IR would be the process of using the data set with each of these different sets of competing assumptions to see if the results stay the same or change. If the result is robust, or the same, across the different competing sets of extra assumptions, then this can give us reason to accept the result. If the result changes across these competing sets, then we have reason to withhold judgment on the particular hypothesis.

Woodward discusses IR in relation to several different disciplines, but considers how it might be applied in climate modeling (Woodward 2006). Someone who is trying to estimate the annual increase in average ocean temperatures due to greenhouse gasses, will require data about ocean temperature increase, but will also need a climate model to use with this data. If there are several competing climate models, IR would be the process of using the data with each if the climate models and comparing results. If the result stays the same (or within some range of similarity) then this should give us reason to believe the result. However, if the data produce different results with the different climate models, *and* there is no reason to accept one climate model as better than the others, then we cannot draw conclusions about the impact of greenhouse gasses on ocean temperatures from the data.

DR, on the other hand, is a method of determining the influence of particular assumptions in a model by altering the assumption to see how this influences the derivation of a result. As Woodward puts it,

> Suppose that we have a model or theory that allows for the derivation of observed facts P. Suppose the model contains some assumption A, which might concern, e.g., the value taken by some parameter x, or the relationship between two parameters x and y, where A figures in the derivation of P. It is natural to ask how sensitive this derivation is to the assumption of A (Woodward 2006).

The strategy of DR, then, is to alter assumption A and to see how this impacts the derivation of P. If altering A still produces P (or a similar enough result), then it can be said that P is a "derivationally robust" result of the model, or at least derivationally robust in relation to A and its alternatives. Woodward's DR varies slightly from Levins' robustness analysis in that it is not only applied to false or simplifying assumptions. It can be applied to any assumption of the model to check its relative influence on some result.

Now, DR is closely related to both de-idealization and IR in a variety of ways. IR and DR can be distinguished in that IR focuses on testing the "inductive warrant" a data set has provides for some conclusion (Woodward 2006). DR, on the other hand, is not about determining inductive warrant of a data set, but rather about checking the relative influence of a certain assumption on the results of a model. While IR focuses on using the same data with different theoretical models, DR focuses on using the same theoretical model but different values of a parameter to check its relative influence.

DR and de-idealization differ in several ways as well. As mentioned above when comparing de-idealization and robustness analysis as Levins presented it, de-idealization requires that some assumption in a model can be made more realistic while robustness analysis does not. This holds for DR, since the focus is simply on altering some parameter, not necessarily making it more realistic. Further, they function somewhat differently epistemically and find different applications. In de-idealization, it is the more realistic model that does the epistemic heavy lifting. However, when it comes to DR, there is no single value of the altered assumption that does the heavy lifting. It is the set that provides an epistemic boost.

This difference can be made clearer by considering the conditions where a result being robust or not provides an epistemic boost. In particular, Woodward held that failure of a result to be derivationally robust only sometimes raises concerns about a model, depending on what else is known about the phenomenon being modeled, while a result failing to carry over through the process of de-idealization does raise concerns. For instance, if it is known that there is only a single realistic value for some parameter in a model, and this value is stable, *and* this value is known, then showing some result is derivationally robust in relation to multiple values of this parameter would not provide much support for the model (Woodward 2006). What we would be most interested in is what the model does with the known, stable value. This sort of situation, where the value of the parameter is stable and known would be more suitable for de-idealization (assuming the parameter of interest was unrealistic to begin with). For DR, it is when there are multiple equally realistic values of some parameter that some results being derivationally robust provides support for a model. Similarly, it is when there are multiple realistic values for some parameter that failure to be derivationally robust in relation to this parameter raises some concerns about the model. When there are multiple, equally realistic values of some parameter, moving from a less realistic to more realistic assumption cannot be done, and so de-idealization cannot be applied.

Woodward provides as an example of when failure of derivational robustness raises some red flags is with social preference models in economics. The success of these models is highly sensitive to values chosen for a few crucial variables (Woodward 2006). However, there are two key aspects about these kinds of models that make failure of robustness concerning. First,

> (a) there is no real evidence that supports these particular choices of parameter values rather than other choices. Arguably, no choices for these values are consistent with all the data and insofar as the particular values chosen are consistent with some range of experimental results, this is equally true of other possible choices (Woodward 2006).

There are multiple equally acceptable values of the parameters which cannot all be captured at once. If the result of a model is tied to a proper subset of these, then this can indicate a concern for the model. Second,

> (b) we have independent reason to think that the parameter values in question (or their distribution) are unlikely to be stable across different people, across contexts, or perhaps even across time for different people and we have no theory which allows us to predict or understand such variation (Woodward 2006).

Not only are there multiple equally acceptable values of the parameters, but it is also believed that, for any individual, the "real-world" value of the parameter varies across time and context. Once again, if the model is incapable of capturing this variation, this can raise some concerns.

As an example of when a failure of robustness raises *no* concern about a model is that of the inverse square law. The correct predictions about the behavior of the planets produced by the inverse square law fails to be robust across a range of values for the gravitational constant and the exponent of the distance term in the law (Woodward 2006). This, however, is of little concern for the model. For instance, the gravitational constant only takes a single value and this value does not vary across individual planets or a range of contexts. Ultimately, there is only a single realistic or appropriate value for the gravitational constant to take, and so the fact that different values of the gravitational constant fail to accurately predict the behavior of planets is of little concern.[3] What we care about is that the correct results are produced with the correct value.

Very generally, the conditions under which failure of derivational robustness raises no concerns are when "we have a great deal of independent evidence both that the parameter is highly stable across a range of different circumstances, that it does indeed take the claimed value and that we can use it, in conjunction with existing theory, to derive a wide range of observed phenomena" (Woodward 2006, p.232). In this case, if the known, stable parameter is incorrect in the model, we would apply de-idealization. Under conditions where there are multiple equally acceptable values of some parameter, and this value does not vary or fluctuate across contexts and individuals, then failure of robustness *can* indicate a concern.

3. Derivational Robustness and Tractability Assumptions

Woodward's account of DR has been influential and has found development and application to models working with certain taxonomies of modeling assumptions (Kuorikoski, Lehtinen, and Marchionni 2010; Raerinne 2013).[4] In this section, I examine two accounts of DR as applied to a taxonomy of modeling that includes three types of modeling assumptions, substantial, galilean, and tractability assumptions. I outline these two accounts to show that they see the same epistemic value

of DR as Woodward and arguments for why DR is applied to just tractability assumptions. In later sections, I will argue that this pairing between DR and tractability assumptions needs some reconsideration.

Kuorikoski et al. (2010) look to extend the discussion of the role of DR, highlighting two different, albeit related, roles. First, they hold that DR can be used to bolster a belief in a model because "it guards against error by showing that the conclusions do not depend on particular falsehoods" (Kuorikoski, Lehtinen, and Marchionni 2010, p.534). This is a concern in line with concerns raised by Levins and captured in Woodward's discussion. However, they also hold that "it confirms claims about the relative importance of various components of the model by identifying which ones are really crucial to the conclusion" (Kuorikoski, Lehtinen, and Marchionni 2010, p.543). The first of these is in-line with the role that Woodward outlines for DR, where it can highlight potential concerns raised by particular assumptions. The second is a more about exploring the causal mechanism of interest.[5]

The first role also closely follows Levins' concern about unrealistic assumptions. It is recognized that models incorporate assumptions that are known to be unrealistic or false. For at least some of these assumptions, there is a concern that falsehoods may play a significant role in the derivation of some result. By altering these false assumptions, and showing that some result is derivationally robust across these alterations, this can assuage some fears about the influence of any *particular* false assumption. The second role is important since it allows modelers to learn more about the causal relations of the real-world target by examining how they interact in the model.

These roles for DR, however, are embedded in a framework that distinguishes different types of assumptions, a taxonomy that is not part of Woodward's focus. In particular, Kuorikoski et al. and their interlocutors share a taxonomy of modeling assumptions that includes three general types; substantial, galilean, and tractability assumptions (Kuorikoski, Lehtinen, and Marchionni 2010; Odenbaugh and Alexandrova 2011; Raerinne 2013; Lisciandra 2017; Harris 2021). In brief, substantial assumptions are the assumptions of a model that are intended to represent the causal features of interest, picking out the features of the real-world target that the modelers are interested in studying. Galilean assumptions are assumptions made for a model that omit certain causal features of the real-world target from the model. These omitted causal features are those that are expected to be "causal confounders" for the substantial assumptions. Galilean assumptions, then, are assumptions about omission of some real causal features so that the substantial assumptions can be better investigated. Finally, tractability assumptions are assumptions made for a model that are simply done for mathematical tractability.

As an example, we can look at a Lotka–Volterra model.

$$\frac{dV}{dt} = rV - (aV)P$$

$$\frac{dP}{dt} = b(aV)P - mP$$

It is a pair of differential equations used to represent the relationship between a predator and prey population. In this model, $V(t)$ stands for the size of the prey population at time t, $P(t)$ the size of the predator population, the constant r stands for the growth rate of the prey population, the constant m stands for the death rate of the predator population, the constant a stands for the predator attack rate, and the constant b stands for the predator conversion efficiency.

An example of a substantial assumption in this model is that the predator and prey populations are negatively coupled. When the prey population size increases, the predator population size increases while when the predator population increases, the prey population size decreases. An example of a galilean assumption is that there are no influences on the size of the prey population, outside of the predator population. This means that factors like environmental carrying capacity are not found in the model, despite there being carrying capacities in the real-world. These sorts of influences are omitted for the purpose of better understanding how predator and prey populations influence each other. Finally, an example of a tractability assumption would be the *specific* functional form chosen to represent rate of prey capture per predator relative to the size of the prey population. For this model, it is assumed that there is a linear increase in rate of prey capture as the prey population increases (Harris 2021). This is a tractability assumption since any chosen functional form is technically false of any real-world target, but necessary to use the differential equations that simplify modeling.

Within this taxonomy, the galilean and tractability assumptions are unrealistic or false assumptions while the substantial assumptions are (hopefully) realistic. Given that substantial assumptions are intended to capture the causal relations that the modelers want to study, it is hoped that these assumptions accurately capture features of the real-world target. Galilean assumptions, on the other hand, are assumptions that omit some real-world features that might get in the way of studying the substantial assumptions and tractability assumptions can misrepresent the target in a variety of ways.[6]

An important aspect of this taxonomy is that it distinguishes different types of unrealistic assumptions, and provides an opportunity to evaluate how these differences might influence the interpretation of models.[7] Galilean assumptions, being made to omit some causal features, are intended to have a clear causal interpretation

(Kuorikoski, Lehtinen, and Marchionni 2010). This clear causal interpretation means that it is clear how to replace galilean assumptions with a more realistic assumption (i.e. de-idealizing the omitted causal feature). However, tractability assumptions are introduced for reasons of mathematical tractability and “in many cases have no empirical merit on their own” (Kuorikoski, Lehtinen, and Marchionni 2010, p.548). Given this, it is not always clear what it would mean to replace a tractability assumption with a more realistic assumption. Further, the falsehoods introduced by galilean assumptions are often suspected of making some difference to results of the model, given that they omit a causal confounder. However, for tractability assumptions, it is hoped that the falsehoods they introduced have no impact on the results of the model (Kuorikoski, Lehtinen, and Marchionni 2010).

It is the fact that models incorporate unrealistic assumptions like galilean and tractability assumptions that leads Kuorikoski et al. to point out that one of the roles of DR is to guard against error by showing that some result does not depend on a falsehood. However, given their differences, this epistemic role is limited to tractability assumptions. As Kuorikoski et al. point out that “it is only the failure of robustness with respect to tractability assumptions that is epistemically problematic” (Kuorikoski, Lehtinen, and Marchionni 2010, p.548). While they hold that DR can be applied to all types of assumptions, a result being derivationally robust, or failing to be derivationally robust, in regards to a substantial or galilean assumption often “suggests a new empirical hypothesis about a causally relevant feature in the modelled system”, which does not, on its own, raise any concerns about the model (Kuorikoski, Lehtinen, and Marchionni 2010).

However, failure of robustness in regards to some tractability assumptions is problematic “because it suggests that the result is an artefact of the specific set of tractability assumptions” (Kuorikoski, Lehtinen, and Marchionni 2010, p.548). The falsehood introduced by tractability assumptions is considered and is treated differently from galilean assumptions because tractability assumptions are supposed to lack a clear empirical interpretation. A failure of derivational robustness cannot be used to generate a new empirical hypothesis about the real-world target because of a failure of robustness with tractability assumptions, unlike galilean assumptions. Further, tractability assumptions are hoped to be negligible to the results of the model, while galilean assumptions are suspected of being causally significant and non-negligible. This is why *failure* of robustness is treated differently between tractability and galilean assumptions, despite both being false.

Similarly, that, “for many tractability assumptions it is often unclear what it would mean to replace them with more realistic ones [...] This is why tractability assumptions are often replaced with assumptions that are also unrealistic, but in a different way” (Kuorikoski, Lehtinen, and Marchionni 2010, p.548), means that they fit into the conditions outlined by Woodward. Since it is unclear what it would mean to replace one tractability assumption with a more realistic one, all potential options are on par, and other strategies like de-idealization are not possible. There are multiple equally acceptable, albeit all false, options. Kuorikoski et al. lay this out, almost exactly as Woodward does, when they state that "For most economic phenomena of interest there might not be a single true functional form, fixed over time, against which the exact form of the assumptions could be compared" (Kuorikoski, Lehtinen, and Marchionni 2010, p.546). It is here that we see agreement between Woodward and Kuorikoski et al on applying DR and its value being in cases where there is no single, best option for the value of a parameter. This, on Kuorikoski et al.'s account captures conditions that are unique to tractability assumptions.

That tractability assumptions are the appropriate target of DR is not only held by Kuorikoski et al. Raerinne, for instance, argues that DR should be considered a method that *solely* focuses on the influence of tractability assumptions, while robustness applied to substantial (and galilean) assumptions is known as "Sufficient Parameter Robustness" (SPR) (Raerinne 2013). This distinction, once again, comes down to epistemic differences. Tractability assumptions, being false, can introduce error into a model. However, in some disciplines, such as biology which is Raerinne's target, there is often redundancy of causes or causal mechanisms, where multiple mechanisms might produce the same effect. The point of SPR is to find a sufficiently abstracted set of parameters that are able to capture the shared effect across these redundant causal mechanisms (Raerinne 2013, p.299). Once again, tractability assumptions are well suited for DR because of the unique epistemic threat they are supposed to pose.

In the literature on DR, tractability assumptions are often singled out as unique. While there are several differences between tractability and galilean assumptions, the significant difference is that tractability assumptions present the conditions that Woodward outlines as those where DR can be used to identify epistemic concerns. For tractability assumptions, it is often understood that there are multiple, equally acceptable, options and no clearly most realistic opportunity. In the next section, I challenge the extent to which this is true.

4. Applying Derivational Robustness to Tractability Assumptions

As discussed in section 3, DR can play two related roles, one to guard against error and the other to determine which assumptions are crucial to the derivation of some result. While these are related, the concern about "guarding against error" is focused specifically on tractability assumptions. However, in order for DR to play this role, the assumption being tested must meet certain conditions. In particular, there must be no, single, correct variation of the assumption.[8] If there is a single correct variation of the assumption, then the model can be de-idealized to ensure that the false assumption is not causing problems. In this section, I intend to show that tractability assumptions, as a category, do not neatly fit into this condition. I will do so by examining examples of tractability assumptions from the literature, showing that several examples of tractability assumptions do not fit this condition. The upshot, in section 5, will be that we should separate discussions of the epistemic import of DR from the common taxonomy of substantial, galilean, and tractability assumptions.

I start with Kuorikoski et al.'s own example of a tractability assumption, that of *specific* functional forms. I will then examine another common example, that of representing populations continuously rather than discreetly (Colyvan and Ginzburg 2003; Colyvan 2013;

Lisciandra 2017). Finally, I consider Hindrik's example of a tractability assumption, the "single planet assumption" in Newton's models of the solar system.

We can start with Kuorikoski et al.'s own example of a tractability assumption, that of a *specific* functional form chosen for a model. However, Kuorikoski et al. present this under conditions that very closely match those required by DR. As quoted in section 3, they draw on examples from economics, and state that, "For most economic phenomena of interest there might not be a single true functional form, fixed over time, against which the exact form of the assumptions could be compared" (Kuorikoski, Lehtinen, and Marchionni 2010, p.546). This lines up well with what Woodward has to say about social preference models.[9] In these conditions, the exchanging or altering of a functional form does not make a model more or less realistic. Rather, each functional form is limited in how much of the real-world target it captures. Given this, applying DR to the functional form of a model seems to be a prime example of such models.

Yet it is not clear that this carries over to all such examples of functional form as a tractability assumption. Harris, discussing the Lotka-Volterra model while raising a different concern about robustness, claims that:

> Although it may be true that any assumed functional form will strictly speaking be false for any real-world predator-prey system, there is certainly a sense in which one particular functional form might be more adequate to describe the rate of prey capture than another, despite both of them being strictly false (Harris 2021, p.14592).

That all tractability assumptions might, technically, be false of a real-world target does not mean that they are all "equally unrealistic". There could, very well, be a functional form that *best* captures the relationship being modeled among the variety of options for functional forms available. For instance, in the example of the Lotka-Volterra model from section 2, it was pointed out that the functional form for rate of prey capture per predator was linear. For any real-world population, it is unlikely that the rate of prey capture will be exactly linear. However, a linear rate of prey capture might best capture the trend or most closely approximate real-world populations.

Now, if this is the case, that there may be some instances where there is a best or most adequate functional form, this does not immediately undercut the value of DR in this case. However, if it is also *known* what the most adequate functional form is, even if it is false, then we would not need or want to apply DR. Rather, modelers would want to apply de-idealization, if the most adequate functional form was not used originally. In Woodward's example of the inverse square law, the fact that some result is lost when the value of the gravitational constant is changed raises no concerns because the model has been made less accurate (Woodward 2006, p.232). Rather than being able to chalk up the loss of the result to its dependence on a particular, false, assumption, we can chalk up the loss of the result to the fact that the new value deviates from the correct value. If there is a most adequate functional form, then the fact that some correct result is lost does not necessarily indicate a fault in the model. Rather, one could chalk up the loss of the correct result to the fact that an inferior functional form was chosen, just like altering the value of the gravitational constant.

This concern about some tractability assumptions does not end with a subset of functional forms. There are a variety of other types of tractability assumptions which are not amenable to the application of DR. Another common example is that of representing populations continuously rather than discreetly (Colyvan and Ginzburg 2003; Colyvan 2013; Lisciandra 2017).[10] Colyvan and Ginzburg discuss the Lotka-Volterra models, and, as they say in footnote 9:

> The use of differential equations in population dynamics is clearly an idealisation; populations are discrete and so the appropriate mathematical machinery is really *difference* equations. Differential equations, however, continue to be used, mainly for their convenience (Colyvan and Ginzburg 2003).[11]

The use of differential equations introduces a tractability assumption. It is an unrealistic assumption in the model, that populations are continuous, that is made for the purposes of mathematical tractability or simplicity.

However, notice that the conditions surrounding this assumption do not meet the conditions outlined for DR to play the guarding role it is intended to. When it comes to this sort of misrepresentation, there is very clearly a more appropriate or correct assumption to make, namely one that allows populations to be modeled discreetly. The conditions where there are multiple, equally unrealistic, options and no single best option is not met by this situation. This kind of tractability assumption simply does not fit into the required conditions.

Another example of a tractability assumption comes from Hindriks (2006). Hindriks' example is that of Newton's model of the solar system. In particular, "Newton assumed that there were no interplanetary gravitational forces. Given his theory, a planet moves around the sun in an ellipse under the assumption that it is the only planet orbiting the sun" (Hindriks 2006). Now, at first this may appear to be a galilean assumption because it is omitting some causal feature. However, the purpose for this omission is not to remove a causal confounder. Rather, the reasons for this assumptions was that:

> The mathematics needed to take account of the gravitational effects of other planets was only developed some time after Newton's death. So Newton was not even in a position to do without the single-planet assumption. This means, I suggest, that the problem of determining the movements of a planet in relation to the sun was intractable for Newton without the single-planet assumption (Hindriks 2006).

The single-planet assumption was one made out of mathematical necessity, since the model could not be constructed otherwise.

Now, just like the example of continuous populations, there is a clear way to correct this assumption. There is a most appropriate assumption to go in place of the single-planet assumption, namely multiple planets and their respective gravitational forces. As far as meeting the conditions laid out for when DR can play a guarding role, this sort of tractability assumption fails. Further, just like the example of functional form where there is a most adequate, a variation of results between the continuous population and a discrete population, or the single planet assumption and one that captures inter-planetary gravitational forces, can be chalked up to the fact that one assumption is more accurate than the other, rather than it simply being an "artifact" of a particular tractability assumption. Essentially, the conditions that Woodward outlines are necessary for DR to play the role of guarding against error, and these conditions are not captured by all tractability assumptions.

5. Why Does This Matter?

Given what I have stated above, there is a very fair question to ask at this point. Without even debating over my proposed examples of tractability assumptions that do not fit the mold, it is fair to ask what the point of this is. I have not established that DR *cannot* play the role of guarding against error and artifacts, because I have accepted Kuorikoski et al.'s point that there are cases where different functional forms are all equally adequate. I do not even want to limit this possibility to just functional forms. Those who argue that DR can be used to assuage fears about tractability assumptions can, with little concern, accept that not all tractability assumptions can be subjected to DR, but DR can still play the role they have laid out for many tractability assumptions.

This pushback is fair, but I think there are two important points to highlight from my discussion. First, and what I see as a less significant point, is that it is always helpful to get a bit more clarity on the methods being used. Knowing that DR cannot be applied to *all* tractability assumptions is useful to know, so that correct applications of DR is carried out. I want to be clear, my intention has *not* been to argue that DR is incapable of playing the epistemic role outlined by Woodward, Kuorikoski et al., or Raerinne. I fully accept that it can play this sort of role. This, however, leads to my more significant point about the value of the taxonomy of modeling assumptions commonly used, and the close connection between DR and tractability assumptions.

This more significant point regarding the use of the class of tractability assumptions. The development of tractability assumptions came out of an exchange across several decades and authors, where the focus was on categorizing the different types of unrealistic assumptions that might show up in a model (Musgrave 1981; Mäki 2000; Hindriks 2006). However, tractability assumptions, as they are understood in Kuorikoski et al. (Kuorikoski, Lehtinen, and Marchionni 2010) and their interlocutors are primarily categorized by the reason that they were introduced into a model, namely for reasons of mathematical tractability.[12] In this way, it makes perfect sense to group the various assumptions together, since this can provide insight into why certain unrealistic assumptions are included into a model and even why models often include so many unrealistic assumptions.

However, categorizing assumptions by the reason they were introduced into a model does not necessarily line up with the epistemic concerns raised by those assumptions. An important difference in the epistemic concerns raised by assumptions seems to be tied to whether or not they can be corrected or de-idealized, if the correct value is known, or if there is a most adequate or correct value for the assumption. Tractability assumptions, as I have tried to argue, are split between some that have a most adequate or correct value and some that do not. In this way, they do not raise consistent epistemic concerns. Given that tractability assumptions do not form a consistent epistemic set in this regard, they may be ill-suited as a category for debates about the epistemic practices of modeling.

This concern about using tractability assumptions can extend to the overall taxonomy that includes substantial and galilean assumptions. Once again, I do not want to deny that this taxonomy has any use. Classifying assumptions by the role they play in a model, such as representing core causal features or omitting causal confounders, can be helpful for better understanding the method of constructing models, particular methods used within specific sciences, as well as why models often end up with certain kinds of misrepresentations. However, once again, this does not mean that the categorizations line up with the epistemic concerns raised by the type of assumption. In particular, the epistemic conditions under which DR can play an important epistemic role are not ones that are clearly a part of any single type of modeling assumption. Rather, talk of the value of DR and when it should be applied should be divorced from the taxonomy of modeling assumptions that is commonly used.

When discussing the epistemic concerns raised by unrealistic assumptions, a taxonomy that better aligns with the features that seem to raise distinct epistemic concerns is superior. In particular, a taxonomy that better captures the difference between the assumptions where there are multiple, equally acceptable options and those where there is a single option.

6. Conclusion

In this paper I have argued that the close connection between DR and tractability assumptions is somewhat tenuous. I have not denied that DR can be used to assuage fears about tractability assumptions. Rather, I have generally looked to establish where DR can be usefully applied to tractability assumptions, as well as evaluated the common taxonomy of assumption types that includes tractability assumptions. I have argued that tractability assumptions do not make an epistemically consistent set, and that the taxonomy that includes tractability assumptions may be of limited use in assessing the epistemic concerns raised by models.

Notes

1 It has also faced its fair share of criticism. See: (Orzack and Sober 1993; Odenbaugh and Alexandrova 2011, 2011; Lisciandra 2017; Harris 2021).
2 I am not arguing that this taxonomy has no value, but when it comes to the epistemology of modeling, I believe it fails to capture what is important.
3 Woodward goes on to point out that this is, in fact, an advantage of the model.
4 Kuorikoski et al. (2010) highlights the role of DR in economic modeling while Raerinne (2013) focuses on biological modeling.
5 Raerinne, despite providing a role for DR, wants to distinguish a different type of robustness analysis for the second role that is called "Sufficient Parameter Robustness". I will discuss this a bit later in this section.
6 I will discuss a few examples in section 4.
7 Musgrave (1981), while not using this taxonomy, provided distinctions of different types of unrealistic assumptions to pursue some debates about economic models, most prominently Milton Friedman's defense of the use of unrealistic assumptions in economic models (Friedman 1966).
8 A further condition to *apply* DR is that we must be able to actually alter the assumption in question. If we cannot alter the assumption, false or not, then we cannot apply DR.
9 An example from economics, to be clear.
10 Raerinne discusses the opposite where "In order to make a model more manageable, we may choose to model a process that appears to change continuously over time, such as population growth, by means of a discrete equation instead of a differential equation" (Raerinne 2013). This example from Raerinne will still run into the same issues that I discuss below.
11 Italics in the original.
12 Lisciandra (2017), in footnote 8 on p. 86, notes Hindriks has a view of tractability assumptions that requires that they remove aspects of the real-world target that are suspected of being non-negligible.

References

Cartwright, Nancy. 2007. *Hunting Causes and Using Them: Approaches in Philosophy and Economics.* Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511618758.

Colyvan, Mark. 2013. Idealisations in normative models. *Synthese* 190 (8): 1337–1350. https://doi.org/10.1007/s11229-012-0166-z.

Colyvan, Mark, and Lev Ginzburg. 2003. The Galilean turn in population ecology. *Biology and Philosophy* 18 (3): 401–414. https://doi.org/10.1023/A:1024121002194.

Friedman, Milton. 1966. The Methodology of Positive Economics. In *Essays In Positive Economics,* 3–43. Chicago: University of Chicago Press.

Harris, Margherita. 2021. The epistemic value of independent lies: False analogies and equivocations. *Synthese* 199 (5): 14577–14597. https://doi.org/10.1007/s11229-021-03434-8.

Hindriks, Frank A. 2006. Tractability assumptions and the Musgrave–Mäki typology. *Journal of Economic Methodology* 13 (4): 401–423. https://doi.org/10.1080/13501780601048733.

Kuorikoski, Jaakko, Aki Lehtinen, and Caterina Marchionni. 2010. Economic Modelling as Robustness Analysis. *The British Journal for the Philosophy of Science* 61 (3): 541–567. https://doi.org/10.1093/bjps/axp049.

Levins, Richard. 1966. The strategy of model building in population biology. *American Scientist* 54 (4): 421–431.

Lisciandra. 2017. Robustness analysis and tractability in modeling. *European Journal for Philosophy of Science* 7 (1): 79–95. https://doi.org/10.1007/s13194-016-0146-0.

Mäki, Uskali. 2000. Kinds of Assumptions and Their Truth: Shaking an Untwisted F-Twist. *Kyklos* 53 (3): 317–335. https://doi.org/10.1111/1467-6435.00123.

Musgrave, Alan. 1981. Unreal Assumptions in Economic Theory: The F-twist untwisted. *Kyklos* 34 (3): 377–387.

Odenbaugh, Jay, and Anna Alexandrova. 2011. Buyer beware: Robustness analyses in economics and biology. *Biology & Philosophy* 26 (5): 757–771.

Orzack, Steven Hecht, and Elliott Sober. 1993. A critical assessment of Levins's The strategy of model building in population biology (1966). *The Quarterly Review of Biology* 68 (4): 533–546.

Peruzzi, Edoardo, and Gustavo Cevolani. 2022. Defending De-idealization in Economic Modeling: A Case Study. *Philosophy of the Social Sciences* 52 (1-2): 25–52. https://doi.org/10.1177/00483931211049759.

Raerinne, Jani. 2013. Robustness and sensitivity of biological models. *Philosophical Studies* 166 (2): 285–303. https://doi.org/10.1007/s11098-012-0040-3.

Weisberg, Michael. 2006. Robustness Analysis. *Philosophy of Science* 73 (5): 730–743. https://doi.org/10.1086/518628.

Weisberg, Michael, and Kenneth Reisman. 2008. The Robust Volterra Principle. *Philosophy of Science* 75 (1): 106–131. https://doi.org/10.1086/588395.

Woodward, Jim. 2006. Some varieties of robustness. *Journal of Economic Methodology* 13 (2): 219–240. https://doi.org/10.1080/13501780600733376.