ICAE Working Paper Series - No. 90 - February 2019

Unrealistic models and how to identify them: on accounts of model realisticness

Claudius Gräbner

First revision (July 2019)

Institute for Comprehensive Analysis of the Economy
Johannes Kepler University Linz
Altenbergerstraße 69, 4040 Linz
icae@jku.at
www.jku.at/icae

(UN)REALISTIC MODELS AND HOW TO IDENTIFY THEM: ON ACCOUNTS OF MODEL REALISTICNESS*

WORKING PAPER

Claudius Gräbner
Institute for Socio-Economics
University of Duisburg-Essen
&
Institute for the Comprehensive Analysis of the Economy
Johannes Kepler University Linz
&
ZOE. Institute for Future-Fit Economies
claudius@claudius-graebner.com

Revision 1 (July 2019)

ABSTRACT

What determines the realisticness of a model? It is argued that to come up with an account of model realisticness that can answer this question, one has to make strong philosophical commitments to an account of representation, an account of model-world comparisons as well as the ontology of models and their targets. Without such commitments it is impossible to determine the realisticness of a model. Since all these areas are subject to ongoing philosophical debate, it is not feasible to come up with a unique and all-encompassing account.

Based on this observation, one account of model realisticness, which is based on an anti-realist fictional view of models, a commitment to realism about mathematical objects, and the DEKI account of representation, is introduced and discussed. The account aligns well with the practice of applied scientists, who regularly apply validation techniques to assess the realisticness of models. This practice can be accommodated nicely in the proposed account, which suggests a number of promising avenues for further philosophical inquiry.

***K*eywords** Models · epistemology · fictionalism · representation · model-world comparisons

1 Introduction

The present contribution is concerned with the quest of developing an account of model realisticness that helps to understand when we can speak of 'unrealistic' or 'realistic models'. Discussions about the realisticness of models abound in both applied sciences and philosophy. While in philosophy the question of how such models become epistemically meaningful takes centre stage (see e.g. Ylikoski and Aydinonat, 2014), in the applied sciences, the 'realism' of a model is regularly used as a sign for quality in the scientific, but also in the public debate about models: while policy makers regularly demand 'realistic' models to base their judgements upon, within the scientific discourse,

*The author gratefully acknowledges funding by the Austrian Science Fund (FWF) under grant number ZK 60-G27.

new models are often motivated by the fact that they are 'less unrealistic' than previous models, and that problematic assumptions have been replaced with more realistic ones, making the model as a whole 'more realistic'.

Notwithstanding this prominence in both the academic and the wider public discourse, the concrete criteria that determine the degree of realisticness of a model remain ambiguous and contested. This paper tries to address this ambiguity in two steps: first, it is argued that this ambiguity is due to the fact that several ways to determine the realisticness of model exist, and that all of them require us to make certain controversial philosophical assumptions, *inter alia* about the underlying account of representation and an account of model-world comparisons. In a second, more positive, step a contingent but consistent account of model realisticness is proposed. This account aligns well with the actual practice of applied scientists to assess model realisticness. Such direct link of a philosophical account of realisticness with the actual practice of model validation undertaken by applied scientists is a novel contribution that hopefully conduces to a close and fruitful interaction between applied modellers and philosophers specialised in the analysis of models.

Throughout the paper, the following formulation will be used both to illustrate the need for making certain meta-theoretical assumptions, and as a vantage point for the development of the new account of model realisticness below:[2]

Realisticness 1

> "The realisticness of a model M is determined by the quality of the relevant aspects of the representation relationship between M and its target in the real world T."

As such, this formulation leaves open a number of intricate philosophical issues, yet if one is willing to make a number of philosophical assumptions it can be turned into contingent, but consistent and useful account of model realisticness. As will be argued in more detail in the next section, this applies to all possible accounts of model realisticness: any such account must at least refer (1) to an account of representation, (2) to an account of model-world comparisons and, thereby, at least to some extent, (3) to an account of the ontology of models. All of these areas remain contested, and it is not my aim to resolve these controversies. Rather, I argue that any attempt to determine the realisticness of a model must take a stand on these issues. I exemplify this by building upon an antirealist fictional view of models, realism about mathematical objects, and the DEKI account of representation. The reward for these assumptions is a consistent account of model realisticness that aligns well with the actual practice of applied scientists, who regularly use particular *validation techniques* to determine the realisticness of a model.

Thus, this paper contributes to the discussion about model realisticness in two respects: the first contribution is to show the contingency of any account for model realisticness upon the three main assumptions mentioned above. This contingency is nothing that makes model realisticness special as a philosophical concept. Yet, by making the contingencies explicit, the paper may facilitate the future development of more sophisticated accounts, and contribute to a more transparent and effective scientific discourse about model realisticness. The second contribution is to come up with a consistent account of model realisticness that aligns well with the practice of applied scientists, and which hopefully helps to strengthen the dialogue between philosophy and applied sciences in a subject area where such dialogue could be most fruitful.

To reach these goals we proceed as follows: The next section explains why a general account of model realisticness is not feasible and the commitment to accounts of representation and model-world comparisons (and, thus, to the ontology of models) is inevitable. Based on this conclusion I use the guiding example from above to show how a coherent account of model realisticness can be derived from such commitments: section 3 grounds the account in the DEKI account of representation (Frigg and Nguyen, 2016; Nguyen, 2016) and section 4 explains how the relevant aspects of

[2] For an outline of how one could interpret 'unrealisticness' differently see Mäki (2017). This formulation is inspired by the practice of many applied scientists to judge the realisticness of a model by comparing selected parts of the model with reality. While they would probably not frame their views on realisticness in such magniloquent philosophical language and think of models more in terms of things somehow similar to their targets, I will show below in section 6 how such account of model realisticness aligns well with the practice of applied scientists. Alternative formulations will also be discussed below in section 2.4.

the representation relationship between the model and its target can be identified. Section 5 then discusses an account of model-world comparisons that grounds quality assessments between the model and its target. How such comparisons are made in practice, and how this corroborates the overall account is discussed in section 6. Finally, section 7 discusses the results and highlights important avenues for further research.

2 A general account of unrealisticness is impossible

If we wish to determine the realisticness of a model, we must make some important philosophical assumptions in a number of contested areas. This is an inconvenient argument: It would obviously be preferable to have general criteria to determine the realisticness of a model, independent of whether one is willing to subscribe to, say, realism or anti-realism with regard to models. Unfortunately, the search for such general criteria is a *cul de sac*, as will be argued in the next three subsections. At the end of this section I try – without success, unfortunately – to get rid of some dependencies by exploring alternative formulations of REALISTICNESS 1.

2.1 The need for an account of representation

First, to justify the label realisticness, an explicit reference to 'reality' is inevitable: it makes sense to discuss the realisticness of models only if they somehow relate to reality.[3] The nature of the relationship between a model and its target has frequently been characterized as a relationship of *representation* (e.g. Frigg and Nguyen, 2017). Moreover, it makes sense to understand 'realisticness' not as a binary property, but to allow for comparative claims such as "model A is more realistic than model B". To make such claims, and to explain when a model represents a target in reality at all, one has to come up with an account of representation. A number of such accounts have been proposed in the literature; only those with direct bearings for upcoming discussions will be mentioned here (for a more complete review see, e.g., ibid.).

Of continuous popularity are accounts based on the *similarity* between a model and its target (e.g., Giere, 2010; Mäki, 2009; Weisberg, 2012). According to these, models represent their targets in virtue of being similar to them. There are a number of variants of this approach, depending on how 'similarity' is operationalized, and which qualifications to the relationship are imposed: classical approaches are 'naturalistic' in the sense that they assume that similarity can be determined objectively. Against this, more recent conceptions, such as the one by Giere (2010), highlight subjectivity and grant the model user an important role: "Agents (1) intend; (2) to use model, M; (3) to represent a part of the world, W; (4) for some purpose, P.", or in other words "A model M represents target T iff used as a representation of T by an agent A for the purpose P" (but see already Apostel, 1960, for a much earlier account along these lines). Granting subjectivity such a strong place in a conception of similarity seems unattractive, yet it plays a considerable role in most recent accounts of representation based on similarity (such as, e.g., Weisberg, 2012, 2013).[4] Moreover, Weisberg (2012) introduced the idea of a *similarity metric*, which compares the features of the target adequately represented in the model to those that are not. This suggests a clear link between the idea of representation as similarity and the assessment of models realisticness: Weisberg's idea of a similarity metric requires only slight modifications to serve as a more general "realisticness metric", which is not even dependent on the subscription to a similarity-based account of representation. This link will be discussed more detailed below in sections 3 and 5. A general challenge for approaches based on the idea of similarity, is that most models do not literally instantiate the properties of the systems they are meant to represent. It is, thus, not *a priori* clear how they can be similar to them (Salis, 2016). There is, however, a solution – which comes with a number of assumptions – to be discussed in the context of model-world comparisons below in section 5.

[3]This is why the qualifier 'its target in the real world' has been added to the guiding example above. There are many models that represent non-actual systems, such as three-sex populations or buildings to be built in the future. Yet, it is not clear whether "realisticness" is the right property to be discussed in this context. Other labels, such as "plausibility" might be preferable. This question will be taken up in section 7, but will also be an interesting topic for future work.

[4]For critical remarks on this role granted to subjectivity see Rusanen and Lappi (2012).

Another influential approach of representation emphasizes the *inferential role* of models (Frigg and Nguyen, 2017). Such approaches focus on how models enable users to make inferences about their target, thereby linking the representative capacities of a model directly to its ability to meet the surrogative reasoning condition.[5] Suárez (2004), for example, formulates necessary conditions for a model M to represent its target T by requiring M to allow for "competent and informed agents to draw specific inferences regarding [T]" (p. 773). Thus, he simply takes the surrogative reasoning condition and turns it into a necessary condition for a model to represent its target (Nguyen, 2016, p. 123). Another inferential view put forth by Hughes (1997) formulates three conditions: it must denote the target ('D'), it must have an internal dynamic that can be examined by the model user, e.g. by using mathematical rules of inference or by simulating the model in a computer ('D'), and, finally, allows for an interpretation in the light of the target system ('I'). But the conception as such does not offer (and does not seem to be intended to offer) any deeper insights into how the representation relationship gets established. An extended account, based on the notions of denotation, exemplification, keying-up and imputation, has been proposed under the label of 'DEKI' by Frigg and Nguyen (2016), and it is able to solve many challenges an account of representation necessarily faces. It will, therefore, be used section 3 to ground the representation relationship in the guiding example introduced above.

There are many other accounts of representation (for a critical review of the various accounts see, e.g., Frigg and Nguyen, 2017), and no consensus has emerged so far. But it is clear that representation as such is important whenever the realisticness of a model is assessed: *without specifying one's account of representation, it is not feasible to determine the degree of realisticness of a model*. A related question is whether different accounts of representation yield different conclusions with regard to the realisticness of a model. While the answer to this question does not affect the fundamental need for a conception of representation, it does have implications for its relevancy: if our assessment of the realisticness of M is the same whichever conception of representation we use to ground our assessment on, the prospect to find a general conception of realisticness would be improved considerably. Yet, at least in practice, different conceptions of representation come with different implications at least for how the adequacy of a representation (and, thus, its realisticness) is assessed: for example, structural approaches to representation posit the existence of a structure in the target systems, and the adequacy of a representation should be assessed on the level of its overall structure. Conceptions based on similarity between models and targets usually take a more modular approach, where different instances of the model are compared to instances of the target and then, as in the case of Weisberg (2012), aggregated into a similarity score. Finally, inferential conceptions would focus on the claims a model allows its users to make about the target. In practice, this calls for different validation techniques, so the adoption of different accounts of representation are likely to yield different assessments of the realisticness of a model.

2.2 The need for an account of model-world comparisons

Second, apart from a notion of representation, the realisticness of a model necessarily rests upon a comparison between the model and its target in reality.[6] Without any reference to 'reality', how could one make sense of the term 'realisticness'? But comparisons between models and objects in reality are not a trivial matter: models are mostly non-actual, i.e. they do not exist as concrete physical objects, while their targets in reality are mostly actual entities. Thus, models often do not literally instantiate the properties of the systems they represent (see e.g. Godfrey-Smith,

[5] Since there is a wide consensus that model thinking should enable users to make informed claims about the targets of these models, Frigg and Nguyen (2017) introduce the term *surrogate reasoning condition* to stress that every account of model representation should explain how these claims about targets can be derived from model thinking.

[6] Some models represent target systems that are not (yet) real: Knuuttila and Koskinen (2017), for example, discuss examples of models in synthetic biology that represent organisms that not yet exist, but which could be designed by using a model. The same holds for models in engineering or architecture, which often represent buildings that have not yet been built. But these cases are not of foremost importance for an account of model realisticness, and will be discussed not before section 7.

2009). Any solution to this problem requires an explicit account of model-world comparisons, an area of ongoing philosophical debate (see, e.g., Salis, 2016; Weisberg, 2013).[7]

Proposals in the literature range from conceptualizing model-world comparisons as hypotheses about abstract and actual entities (e.g. Giere, 1990), over claiming that we can compare imagined model systems with actual targets in a fictionalist context (e.g. Frigg, 2010) to the denial that such comparisons are made at all (e.g. Toon, 2012). A thorough discussion of the pitfalls of these proposals, as well as a constructive proposal that is taken up in section 5, can be found in Salis (2016). In the end, the core issue is that model systems do not literally instantiate properties of their targets and comparisons are thus necessarily false (ibid.). But as explained in section 5, this problem can be circumvented, for example, by committing oneself to the realism of mathematical objects. In any case, if one wishes to come up with a complete account of model realisticness, one has to build upon a convincing account of model-world comparisons.

2.3 Necessary commitments on ontology

Third, as argued in Salis (ibid.), any successful accounts of model-world comparisons themselves require explicit assumptions about realism of models, mathematical objects or other objects in the target systems. Thus, whenever we are concerned with the 'realisticness' of models, we must be willing to make certain ontological assumptions with regard to the nature of models and their targets. The ontology of models is heavily debated in philosophy: prominent accounts include the view that models are akin to the work of *fiction* (see e.g. Frigg and Nguyen, 2016; Salis, 2018), others consider them to be *abstract structures* (e.g. Weisberg, 2013) or *epistemic artifacts* (e.g. Knuuttila, 2011).[8] Adherents to the idea of models as abstract entities consider them to be "socially constructed entities", which have "no reality beyond that given to them by the community of scientists" (Giere, 1990, p. 78). More recently, the move towards more subjective accounts of representation has affected accounts for model-world relationships as well: for Weisberg (2013), models consist of a structural part, but as also a construal, which summarizes the intentions of the model user in the form of the intended target and scope as well as the (subjective) fidelity criteria. Yet such accounts still face the problem that abstract entities do not instantiate the properties of their targets, which renders them incompatible with any realist conception of models (see also Salis, 2016). The same challenge is faced by (indirect) fictionalist accounts, such as the one by Frigg (2010).

The latter example shows that even within a certain view of models, such as fictionalism, various interpretations exist: aside from subscribing to direct or indirect variants of fictionalism, one might also be a realist with regard to fictions, or an antirealist. The same is true with regard to the entities in the target system, and mathematical objects, which play an important role in many models both in the natural and social sciences (see e.g. Nguyen and Frigg, 2017). Thus, even if one subscribes to the idea that models are akin to the work of fiction and endorses a particular ontology of fictions, such as the one of Walton (1990), one still has several interpretations at one's disposal. For the present purpose, it is important to remain as general with regard to one's ontology of models as possible, but to be specific enough to make sense of the model-world comparisons inherent to any account of model realisticness such as in *Realisticness 1* above. To this end, however, an explicit stand on the ontology of models is necessary – at least to the extent one's account of model-world comparisons requires this. In this sense, the necessary reliance of an account of model realisticness on the commitments on the ontology of models is less direct than its reliance on an account of representation and model-world-comparisons.

2.4 An alternative account of realisticness?

In all, the elaborations above make clear why developing a general account of the unrealisticness of models is difficult and maybe even impossible: any account of model realisticness depends on assumptions in currently contested areas

[7] One might debate whether, for example, the GDP of an economy is actual. If not, one would then compare non-actual models with non-actual entities. For these kinds of comparisons basically the same problems must be solved and we do not need to distinguish between these cases here (for details see Salis, 2016).

[8] For a recent survey see, e.g., Gelfert (2017).

of philosophy, and as long as philosophers have not settled these issues definitely, a plurality of different, contingent and plausible accounts of the realisticness of models is unavoidable (see already Mäki, 1998). But this must not be an excuse for unwarranted specificity: one should try to develop accounts of a model's realisticness as general *as possible*. Therefore, we shall meticulously reconsider the possibilities to eliminate unnecessary dependencies in REALISTICNESS 1.

We might do so be reformulating REALISTICNESS 1 in the following way:

Realisticness 2

> "The realisticness of a model M is determined by the accuracy of relevant claims (or imputations) it makes about a target T."

REALISTICNESS 2 would drastically reduce the dependencies on contested claims by shifting the the burden of explanation to the claims about the target made based upon the model. This is potentially attractive but there are considerable costs associated with the alternative formulation:

First, models are often judged as unrealistic because of their *assumptions* and not because of their resulting imputations. There is an asymmetry between assumptions and imputations since models are used to produce imputations from their assumptions, and it makes little sense to treat the simplifying and idealising assumptions for, say, the decision making of economic agents in a model as imputations regarding actual economic actors. Nevertheless, models can be and are criticized (or at least labeled) as unrealistic because of their assumptions. This is hard to reconcile with an account of model realisticness that determines realisticness exclusively in terms of claims or imputations based on the model.

Second, REALISTICNESS 2 is less general than REALISTICNESS 1: claims made based on a model are only one of several ways to judge the realisticness of a model. I discuss in section 6 the various types of model validation which are common in applied science. All of them align well with REALISTICNESS 1, but how could we make sense of, for example, input validation with REALISTICNESS 2? One would need to contend that the assumptions a model makes on the input side can be translated into claims about the target. But, as argued above, this is unconvincing: assumptions are not imputations and models are about deriving imputations from assumptions; collapsing the distinction between them would do no good.[9]

Finally, establishing the link between the model and the claims based on the model is not as trivial as it first may seem. There must be a clear relationship between the model and the claims. It would make little sense to determine a property a model, realisticness, by reference to claims that are not directly related to the model. The most immediate link would be that of justification: we make a claim a about target T because of our analysis of model M, which represents T. Thus, the claim a must be dependent upon our analysis of model M. In this case, one first has to justify the adequacy of M for the situation at hand. Only then it makes sense to rationalize claims about T by reference to M. As also argued by Salis (2016), a necessary prerequisite for this conclusion is a comparison of some sort between M and T, i.e. a model-world comparison. Therefore, some account of model-world comparisons is necessary. As argued by Salis (ibid.) and in section 5, these conceptions themselves hinge upon a certain ontological conception of models.

Thus, it seems that getting rid of any of the three dependencies is not straightforward and REALISTICNESS 1 is of acceptable generality. In the remainder of this paper, I wish to demonstrate how a coherent account of model realisticness can be based on REALISTICNESS 1 if one takes a definitive stand on all the areas discussed above.

3 Establishing the representation relation

> "The realisticness of a model M is determined by the quality of the relevant aspects of the **representation relationship** between M and its target in the real world T."

[9] At the same time, *Realisticness 2* can be thought of as a special case of *Realisticness 1* when one is mainly concerned with output and, to some extent, process validation.

The first part of our account of model realisticness refers to the "relevant aspects of the representation relationship between the model and its target". This begs two questions: (1) what is meant by a 'representation relationship' and (2) what are the relevant aspects of this relationship? The first question will be discussed in this, the second in the next section. To answer both of them, an account of 'representation' is required.

As discussed above, several answers to the question of when a model M represents a target T have been proposed. Frigg and Nguyen (2017) not only provide an overview about various accounts of representation, but also formulate a number of demands that every account of representation should meet. In the present case I will mainly rely on the DEKI account of representation, as developed by Nguyen (2016). Not only does DEKI meet the most of the requirements as discussed in Frigg and Nguyen (2017), it is also straightforward to align it with the actual practice of applied modelers to relate their models to reality (Gräbner, 2018).

The DEKI account formulates conditions that a model M must satisfy in order to count as a representation of a target T:

1. The model M must **denote** its target T.
2. The model M must be a Z-representation **exemplifying** properties $P_1, ..., P_n$.
3. The model M must come with a **key** K that indicates how the properties $P_1, ..., P_n$ can be translated into a set of features $Q_1, ..., Q_m$.
4. The mode M must **impute** some features on its target T.

The first condition requires the model user to make clear *that* she uses the model as a means for representation, and she must also specify *what* the target system of the model is. The second condition requires the model to be a Z-representation. A Z-representation is a one-place predicate and Z is a place-holder for the kind of the target of the model. For instance, a picture p of a woman is a *woman*-representation: $WOMREP(p)$. This does not necessarily mean that it represents one particular woman – the women could also be fictional (such as Wonder Woman). Referring to Z-representations in this context is important for any general account of representation because one would like it to account for models that represent general (such as cities in general instead of a particular city) or hypothetical (such as a three-sex population) targets as well. An arbitrary object becomes a Z-representation when it is interpreted as such. A Z-representation exemplifies certain properties $P_1, ..., P_n$, if it instantiates them by directly referring to them. This represents an important challenge for many accounts: since most models are not physical systems they cannot really instantiate the properties of their actual targets. As will become clear in the next section, here lies a crucial challenge: since most models are not physical, they cannot instantiate many properties of their targets. A mathematical model of the economy does not instantiate any stock of capital, yet the true economy clearly does. Solutions to this challenge will be discussed in the next section when scrutinising model-world comparisons, but for now we follow Frigg and Nguyen (2017) and assume that instantiation must not necessarily be interpreted literally, but that an explicit interpretation of the model structure instantiating certain properties is sufficient.

The third condition asks for a key that links the just discussed properties of the model, P, to the relevant properties of the target system. How this is to be done must be specified by a *key* (or a 'dictionary'), the simplest example for a key being the legend of a map. This key turns the properties of the model P into features, Q, which are interpretations of P in the light of the target system. The final condition requires the model to impute at least one feature Q on the target system. This step is crucial when we assess the 'realisticness' of a model, because one particularly relevant way in which a model can be unrealistic is by making false imputations: a model can make a prediction about how the target behaves, but it behaves differently. But a model that makes false predictions about its target does not stop representing this target - it just *mis*represents the target because it imputes the properties wrongly on the target.

The DEKI account is summarized in figure 1. As we can see, there are a number of links between the model and its target in the shaded area, and the assessment of any of these links can make us claim the model to be 'unrealistic'.

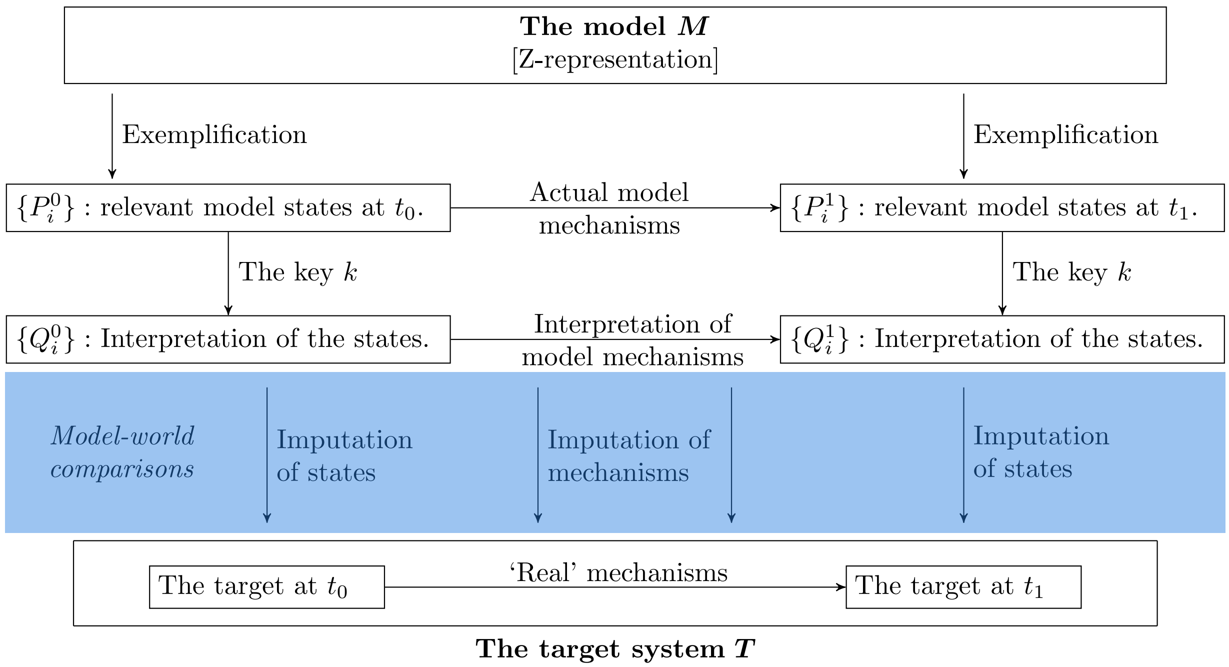

Figure 1: A visualization of the DEKI framework. Note that the denotation of from the model to the target is not visualized for reasons of readability. Moreover, the account is also applicable to models that are not meachanistic, for which the remarks on interpretations of model mechanisms can be omitted.

Borrowing the notation from Weisberg (2012)'s weighted feature matching we can say that the overall realisticness of the model M with target T, $R(M,T)$, depends on the weighted assessments of these links:[10]

$$R(M,T) = \alpha\left[\phi\left(f(Q^0T^0)\right)\right] + \beta\left[\phi\left(f(Q^mT^m)\right)\right] + \gamma\left[\phi\left(f(Q^{t>0}T^{t>0})\right)\right], \tag{1}$$

where α, β and γ are the weights placed on the assessment of the various links between the model and reality, $f(\cdot)$ are the weighting functions that weight the features of the model and the target within this link, Q^0 is the feature set of the model in $t = 0$, i.e. the initial setup of the mode, $Q^{t>0}$ is the state of the model in some timestep $t > 0$, Q^m are the interpretations of the mechanisms built into the model, and ϕ is the assessment function that compares the model and its target. If the comparison takes the form of Weisberg's feature matching, then ϕ is the ratio of shared features. For the first part of function (1):

$$\alpha\left[\phi\left(f(Q^0T^0)\right)\right] = \alpha_1 f\left(Q^0 \cap T^0\right) + \alpha_2 f\left(Q^0 - T^0\right) + \alpha_3 f\left(T^0 - Q^0\right) \tag{2}$$

We will expand this formulation to capture the various ways in which the links can be accessed in practice in section 6. For now, the more important question is: which link is the decisive one when it comes to the realisticness of the model? In other words, where do the weights α, β and γ in the formula come from? This directly relates to the second question posed at the beginning of this section: what are the "relevant aspects of the representation relationship" of which the quality needs to be assessed?

[10]Without loss of generality we take $\alpha = \{\alpha_1, \alpha_2, \alpha_3\}$ to allow for potential differences in the weights for shared and unshared features.

4 Selecting the relevant aspects of the representation relation

> "The realisticness of a model M is determined by the quality of **the relevant aspects of the representation relationship** between M and its target in the real world T."

In the previous section we have clarified how the representation relationship in our conception of model realisticness can be grounded on the DEKI account of representation. Now we can build upon these insights and concern ourselves with how the "relevant aspects" of this relationship can be determined. As can be seen in the shaded area in figure 1, there are a number of links between the model and its target. Thus, the question becomes: which of the links is decisive for our assessment of the realisticness of the model, and which aspects of the chosen link should be of central importance? Or, in terms of equation (1), how should we choose the weights α, β and γ, and the specific form of the weighting functions $f(\cdot)$?

It does not make sense to assess the quality of all of the links at once, i.e. to use $\alpha = \beta = \gamma > 0$ and $f(\cdot)$ assigning equal weights to all aspects of these links. The reason is that modelers usually do not mean to design a model such that it matches all aspects of their targets. Rather, they focus on particular aspects of the latter (Mäki, 2009). For Mäki, these aspects are specified in a model comment, for Godfrey-Smith (2006) and Weisberg (2007, 2013) in the 'construal' of a model: Weisberg argues that any model must come with such a 'construal', which is determined by the model user's intentions and is composed of four parts: an assignment that provides "the specification of the phenomenon in the world to be studied" (Weisberg, 2007, p. 219), and essentially plays the role of denotation in the DEKI account. The intended scope clarifies those parts of the target that should be represented by the model, while the fidelity criteria specify the "standards theorists use to evaluate a model's ability to represent real phenomena" (ibid., p. 219).

Against this backdrop it seems inadequate to specify *a priori* which links between a model and its target in the shaded area of figure 1 are relevant to assess its realisticness: if one would consider all aspects on equal footing, almost all models would qualify as 'unrealistic'. If not, one would need to come up with an all-encompassing criterion. Such claims a la Friedman (1966)'s "only the predictions of a model matter" have been proven to be of little value. Thus, it makes sense to leave the decision about which links between a model and its target should be placed under scrutiny to the model user, but to require the latter to be explicit about the *kind of realisticness* she is assessing: if one scrutinizes the mechanisms that are designed in the model, and compares them with the mechanisms in the target, one is assessing the realisticness of a model with respect to its mechanisms (corresponding to a relatively large value for β in equation (1)). If one compares the output of the model for some time in the future with those values observed in the target, one is assessing the realisticness of a model with respect to its predictions, using a comparatively large value for γ in equation (1). It makes little sense to speak of an overall realisticness of a model, because it is by no means clear how the scores for the various aspects of the relationship between the model and its target should be aggregated. Surely, the purpose for which the model is to be used provides for some guidance, since it is less important to assess the mechanisms in a model that is used for purely predictive tasks, but even here it is not clear whether one should rely *only* on the output of the model, and, if not, how to weight the various assessments of the model realisticness.

The only solution to this problem I can envisage is to abolish the idea of a 'general realisticness', but to qualify the kind of realisticness one is talking about, and to choose the weights in equation (1) on a case-by-case basis by refering to the construal of the model. Such a solution is internally coherent and helps to avoid communication problems such as researcher A furiously criticising the unrealisticness of researcher B's model by pointing to the unrealistic assumptions, but researcher B actually meant to develop the model for entirely different circumstances and only for the purpose of prediction.

5 The model-world relationship: theoretical issues

> "The realisticness of a model M is determined by the quality of the relevant aspects of the representation relationship **between M and its target in the real world T**."

In the previous section we have clarified what is meant by the 'the relevant aspects of the representation relationship'. We will now be concerned with the question of how to assess their quality, or, in terms of equation (1), how to determine the realisticness values for the various dimensions (i.e. the function $\phi(\cdot)$). Again, two steps are necessary: the first one is theoretical and is concerned with the assumptions we have to make in order to reasonably make comparisons between a model and its target at all. Such comparisons are necessary to 'assess the quality' of the relationship. The second step is more practical and asks how such comparisons can be and are made in practice. This helps aligning the general account with the actual practice of scientists and, luckily, to corroborate it further. The first step will be taken in this, the second step in the next section.

As argued above, the realisticness of a model refers to its alignment with its target, and the DEKI account of representation as introduced above as well as our interpretation of Weisberg (2012)'s weighted feature matching in equation (1) already hint towards several links between a model and its target. But assessing these links is not straightforward: aside from a few exceptions, models do not literally instantiate the properties that are compared to the properties as instantiated by the target (e.g. Godfrey-Smith, 2009; Salis, 2016).[11] Thus, it is *a priori* unclear how statements about a model translate into (potentially) true statements about the real target and, thus, how the quality of the model-world relationship can be assessed: it is not clear how a statement based on any of the links in the shaded area in figure 1 can be be true (or 'good' or 'realistic').

A potential solution has been proposed by Salis (2016). Her account of model-world relationships not only produces the intuitive truth and quality conditions for model-world relationships but also helps to circumvent the problem of uninstantiated properties in model systems.[12] At the same time the account comes with two assumptions to be discussed in more depth at the end of the section: one has to subscribe to realism with regard to mathematical objects and anti-realism with regard to models.

Salis' account, which will be followed in the rest of this paper, works as follows: First, one develops and analyzes a model within a purely imaginary context. Similar to Frigg (2010), she builds upon Walton (1990) and considers models to be props in a game of make-belief. So far, we are operating exclusively in the upper part of figure 1, i.e. the model.

Second, one refers to what she calles an *extended imagination*, where one extends the game (in Walton's sense) involving the original model by an extended fiction inspired by the real target system. Here, one compares - within one's imagination - properties of the (fictional) model with the (fictional) target. Now we are operating within the full picture of figure 1, but consider the lower part (the target) only in our imagination (which is, of course, still inspired by our observations of the real target). In particular, the model-world comparisons are conducted for two imagined systems.

Third, to leave the imaginative sphere and to consider the lower part of figure 1 not only in our imagination, but as the real system it is,[13] we must 'build a bridge' between the fictional world of the model and the real world of the target system. If we assume that mathematical objects truly exist, they can provide for this bridge: we can say there exist mathematical values (say, x_1 and y_1) describing the state of the target, and mathematical values (say, x_2 and y_2) describing the states of the model we interpret to represent the states of the target. As shown in Salis (2016, p. 256), a comparison of these values, i.e. comparing x_1 with x_2 and y_1 with y_2 produces the intuitively correct results. We may illustrate this procedure via the example visualized in figure 2.

[11] An important thing to remember is that in figures 1 and 3 P refers to the states in the model and Q to the interpretations of these states with respect to the target. Q, as the interpretation of a model state, is ontologically different to the target. The challenge is to nevertheless compare them.

[12] Her account also provides an answer to the critique of fictionalism by Odenbaugh (2015) who fears that "if modeling is a form of make-believe, then this scientific success is make-believe as well" (p. 285). Thus, it also allows us to remain more agnostic with regard to our ontological stand on models, which is attractive for any account of model realisticness.

[13] When we compare models with reality in practice we often do not access reality directly, but use data which itself has been created by models: data on the unemployment rate or the GDP of an economy, which is often used to assess the realisticness of a model, is often not a result of a direct observation, but itself the outcome of a model. Thus, this final step is not trivial. This point will be taken up again in section 6.

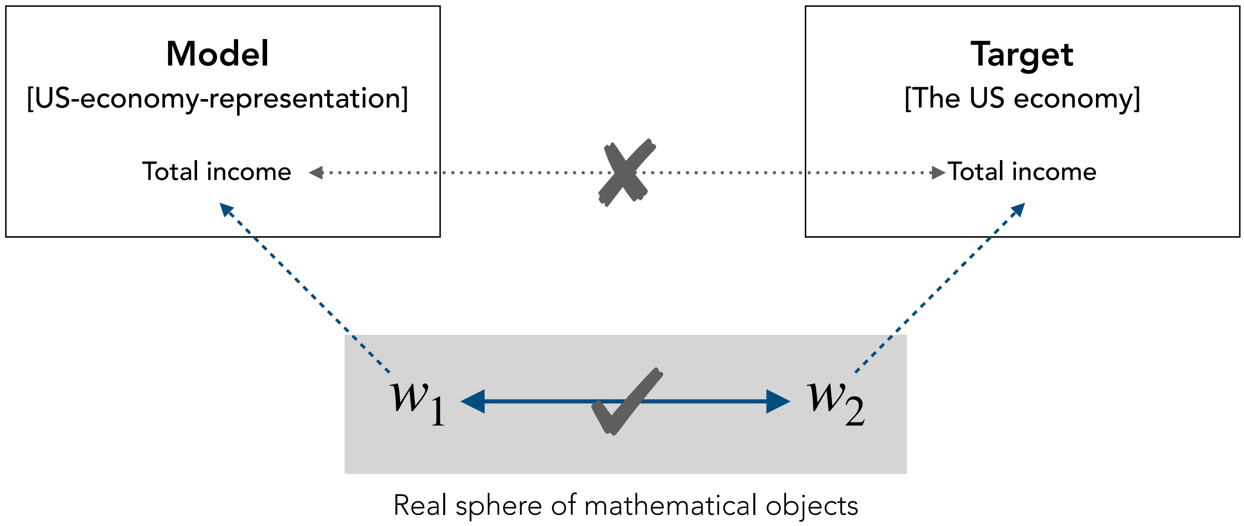

Figure 2: An illustration of how model-world comparisons work once realism for mathematical objects is assumed. The upper comparison between the model and the target cannot work since the model cannot instantiate the property 'total wealth' as the actual US economy. The comparison of the two degrees of total wealth within the sphere of mathematical objects is feasible, however.

Suppose we have a model that we use to explain the aggregated income in the economy of the USA. Assuming that the real economy in the USA instantiates a property 'total income', we cannot compare the state variable of the model (P_Y) that we interpret to represent aggregate income in the model economy (Q_Y) with the true values, since the model does not really instantiate aggregate income. We can conduct this comparison only within an extended game, in which we imagine that our model economy instantiates 'total income' in the same way as the real target economy does. But we can assume the existence of some degrees for aggregated income w_1 and w_2 (which we consider to be mathematical objects), such that $w_1 \approx w_2$. By subscribing to realism with respect to mathematical objects, these values truly exist and can serve as a bridge between the model and the real world.[14] We can then claim that the aggregated income in the USA is w_1, and, according to our model M, the US model system has aggregate income w_2. This latter statement obviously has truth conditions, and it can be evaluated against the states of the real world.

This way, we can directly compare the values and evaluate them against the fidelity criteria as attached to our model M and fill in the respective scores in formula (1). Depending on the outcome of this assessment, we can make statements about the realisticness of the model if the model shares many of those instantiated properties with its target (similar to Weisberg, 2013). By this, we have resolved the meta-theoretical issues in assessing the relations in the shaded area of figure 1. But before we turn to the practical difficulties inherently attached to these comparisons in the next section we must discuss the ontological commitments that we have to subscribe to for our theoretical solution to be viable:

First, one has to take an antirealist stand on models as fictions. Salis (2016) formulates her account within a fictional view of models, according to which models are props in a game of make-belief a la Walton (1990), just as literal fiction (see Frigg, 2010). If one took a realist view on models one would assume that the fictional entities in the models exist. But if they do, it is not clear how they could instantiate the properties their actual counterparts have. One might, of course, *imagine* them to instantiate these properties, but in this case an antirealist account of models, according to which we imagine the fictional entities in the models right from the start seems much more plausible (Salis, 2016).

[14]We must subscribe to realism here only because these values could otherwise not 'bridge' the fictional world of the model with the real world of the US economy. This does not mean that one can only compare two numbers if one is a realist with regard to mathematical objects.

Second, one has to assume realism with regard to mathematical objects. The reason is that mathematical objects represent the 'bridge' between the model and the target: by assuming that there exist mathematical values that represent the degrees of properties to be compared between the model and the target, we ensure that we compare two instantiated properties. In other words: if mathematical objects should serve as the 'bridge' between model and reality, we must assume this bridge to exist - otherwise one could not walk upon it.

6 The model-world relationship: practical issues

> "The realisticness of a model M is determined by **the quality of the relevant aspects** of the representation relationship between M and its target in the real world T."

Now that the (meta-)theoretical issues about the relationship between models and their targets have been discussed we can turn to the practical issues involved. This contributes to our philosophical undertaking in three ways: first, we see whether our account is consistent with scientific practice; second, we might even be able to distill further philosophical lessons from this practice; third, we might find inspiration to formulate new challenges for further philosophical scrutiny.

When applied sciences relate their models to reality they regularly employ *validation techniques*. In contrast to *verification* techniques, which are used to assess the internal functioning and the coherence of models, *validation* techniques are used to assess the link between the model and its target (Gräbner, 2018). Scientific practice here lends support to the claim of section 4 according to which there is not a single all-encompassing way to determine the realisticness of models, but various dimensions of realisticness must be distinguished. While there are numerous concrete techniques for model validation – which necessarily differ depending on the modeling framework at hand – four main types of validation can be delineated (Gräbner, 2018; Tesfatsion, 2017): input validation, process validation, descriptive output validation, and predictive output validation.

These different forms – of which the boundaries can be fluid in practice – are compared visually in figure 3, and, after a slight reformulation, it is straightforward to align them with the structure of equation (1): *Input* validation assesses how well the initial model specifications fit the target system and it takes place in the red-shaded area of figure 3. This is the first term in equation (1). For example, in a model of a financial market, one might compare the initial number of traders in the model, and in the target markets. Such assessment of the initial model conditions does not necessarily rely on a one-to-one comparison, but usually includes an 'interpretation': a single agent in the model could be thought of as representing 100 traders in the target market. Nevertheless, the model must provide some 'key' (Frigg and Nguyen, 2016) that facilitates such interpretations and prevents them from becoming arbitrary.

Process validation studies how well the mechanisms built into the model mimic the mechanisms operating in the target system. It takes place in the green-shaded area of figure 3, and refers to the second term in equation (1). While this is important whenever one is interested in the *structural* validity of a model (see e.g. Grüne-Yanoff and Weirich, 2010), it is also notoriously difficult since mechanisms in reality are usually not directly observable.[15]

Descriptive output validation scrutinizes the ability of the model to replicate real data. For example, one might ask how well a macroeconomic model is able to re-create a past time series for GDP growth. This practice has to be distinguished from *predictive output* validation. Although equivalent in classical accounts of scientific explanation (e.g. in the covering law model of Hempel and Oppenheim, 1948), and located in the same blue shaded area in figure 3, the two are exercised very differently: while descriptive output validation simply means to calibrate the free parameters of a model to create an output that matches the time series of interest, predictive output validation means to separate the available data into a training and test set, to use the former to calibrate the models, and then to test its performance out

[15] There are many reasonable ways to assess the question of whether the implemented mechanism A is more or less likely to operate in the real world than mechanism B. These ways include expert and stakeholder validation (or 'participatory validation', Smajgl and Bohensky, 2013), process tracing (Steel, 2008, ch. 9), face validation (Klügl, 2008) and a clever use of experiments (e.g. Bravo et al., 2015).

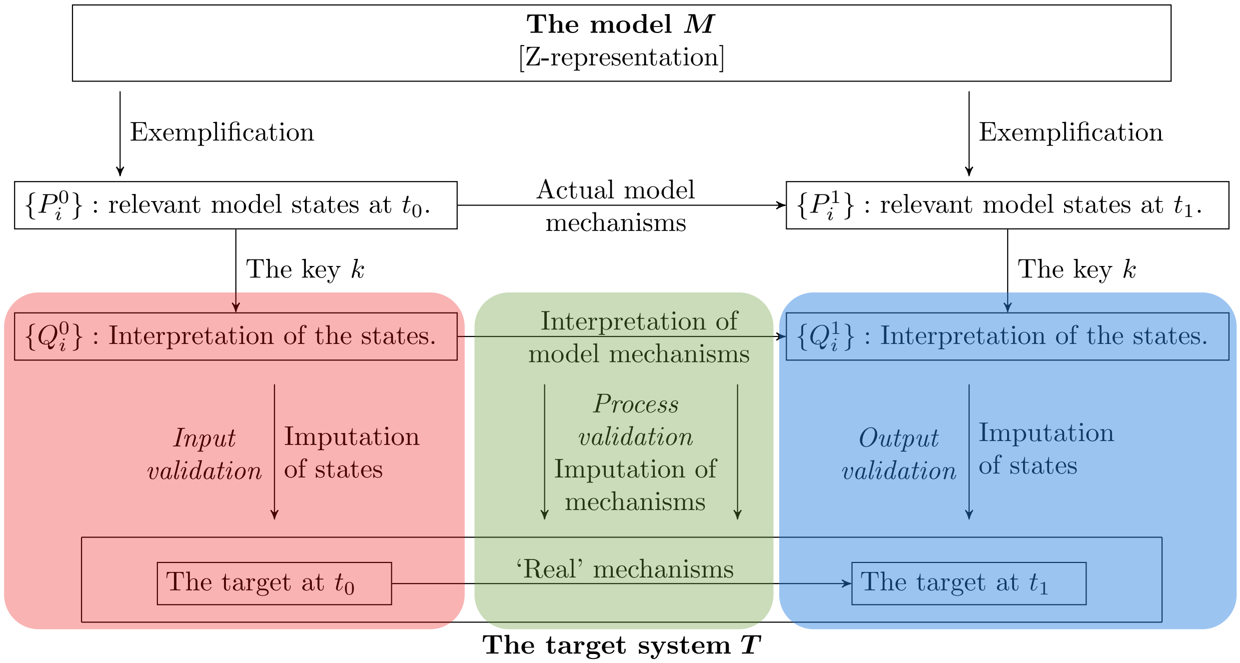

Figure 3: An illustration of where the various kinds of validation techniques take place. Note that there are two different kinds of output validation, *descriptive* and *predictive* output validation. They cannot be distinguished in the figure, but only by the validation practice: for predictive output validation one compares the values the models predicts after being trained on a limited training data set not encompassing the whole time period with those of the target in the whole time period, while *descriptive* output validation only calibrates the model to match the states of the target. For such a calibration, all data is used, and no distinction between test and training data is made.

of the original sample on the test set. It is well-known from statistics that a descriptive output validation supplies best results for models with many free parameters while such models – due to the issue of overfitting – perform poorly when predictive output validity is assessed. Thus, while it cannot be distinguished in all theoretical accounts, the difference between descriptive and predictive output validation in practice is huge and significant. Therefore, equation (1) must be modified to allow for this distinction:

$$R(M,T) \;=\; \alpha\left[\phi\left(f(Q^0T^0)\right)\right] + \beta\left[\phi\left(f(Q^mT^m)\right)\right] + \gamma\left[\phi\left(f(Q^{t>0}T^{t>0}_{\text{full data}})\right)\right] + \delta\left[\phi\left(f(Q^{t>0}T^{t>0}_{\text{test data}})\right)\right] \tag{3}$$

.

In this final formulation, the last term in equation (1) has split, such that for every model we can distinguish between its realisticness in terms of *descriptive* and *predictive* output realisticness.

These four approaches to model validation echo our claim from the previous section according to which there is not one all-encompassing way to assess the realisticness of a model, but rather that a model can be realistic in various dimensions. In practice this means that we may consider a model to be realistic in the sense of its descriptive output capacity whenever it can be calibrated well to existing observations, and we may call it 'realistic' in the sense of its mechanistic adequacy it performs well in process validation. Of course, the same model may score very differently in the various forms, or it may not even be accessible to all of them. Moreover, there are many cases in which we evaluate the realisticness of a model with more than one non-zero weight in equation (3). Some of the resulting combinations are very common, such as the distinction between the realisticness in terms of predictions (with relatively large δ) and the one in terms of mechanisms (with relatively large β), others less so.

In other cases one might also question whether the term 'realisticness' is the right label. For example, some would consider a model that is realistic only in terms of its ability to predict well to be the result of 'mere curve-fitting' – although few would treat a model that is realistic in terms of its mechanisms, but not in terms of prediction, equally snidely. But, I would contend, this reflects informal preferences of research community, which probably differ across disciplines and areas of application, and does not point to a fundamentally different account of realisticness - only to different prioritisation within the account just introduced.

The conclusion that there is not a single supreme form of model realisticness, but that it makes sense to distinguish different dimensions of it enjoys further corroboration by the claim of some that there are some serious practical – maybe even fundamental – trade-offs with regard to the various ways models are related to reality: the idea of trade-offs in modeling design goes back to Levins (1966). More recently, Matthewson and Weisberg (2008), have proven a fundamental trade-off between model precision and generality. With regard to the realisticness of a model, Gräbner (2018) claims that there are at least practical trade-offs faced by applied scientists that prevent a model to perform well in all four kinds of validation simulataneously. And at least for predictive and descriptive output validation, this trade-off is formally well established in statistics (see e.g. chapter 8 in Stachurski, 2016). The existence of such trade-offs makes it even clearer that it is less useful to call a model 'realistic' or 'unrealistic' *per se*. Rather, one must be specific *in what sense* the unrealisticness of a model is assessed.[16]

7 Summary and discussion

This contribution was concerned with accounts of model realisticness. While discussions about the realisticness of models abound in both philosophy and applied scientific practice, coherent accounts that provide for procedures to classify models as 'realistic' or 'unrealistic' have been sparse. The present contribution tried to address this gap in the literature by making two related contributions:

First, it has been argued that any account of model realisticness must be based both on an account of *representation* and an account of *model-world comparisons*. Particularly the latter also requires one to make certain ontological commitments about the realism of mathematical objects and model systems. All these areas are philosophically highly contested, which is why it is not surprising that no single account of model realisticness has emerged, and unless philosophers find a consensus on the subject areas just mentioned – which is unlikely to happen – a general and all-encompassing account of model realisticness remains a chimera.

Second, it has been illustrated how one can reach a coherent account of model realisticness once one is willing to make philosophical commitments in these areas. This has been illustrated by an account that relies on the DEKI account of representation and is built on the assumptions that (1) mathematical object do exist independent of the mind (realism with regard to mathematical objects) and (2) models do not exist independently of the mind (anti-realism with regard to models). The account can be expressed via the simple formulation "The realisticness of a model M is determined by the quality of the relevant aspects of the representation relationship between M and its target in the real world T". It is possible to express this account formally in a way that is very similar to Weisberg's similarity account of representation

[16] While the practice of using various validation techniques aligns well with the account of model realisticness outlined so far, there remain two practical challenges, which are both of considerable philosophical interest: First, a lot of data that is used for model validation is not a collection of pre-existing facts but strongly theory-dependent. This begs the (admittedly, age-old) question of how and whether a model can be related directly to its target in reality at all. Second, models not only represent their targets, they also impact upon them, or are even used to transform the world along their own lines (Boldyrev and Ushakov, 2016). This becomes particularly relevant if a model impacts causally upon its target in such a way that its own predictions are validated. Despite these difficulties, the ways of applied scientists to compare models to their targets corroborates (or at least is consistent) with the central ingredients of our account of the realisticness of models.

(see equation (3)) and the account comes with a number of advantages and interesting philosophical and practical implications, which deserve further attention:

One advantage refers to the flexibility of the account: by relying on the DEKI account of representation, it also does not have problems with models of non-actual systems, such as models of building to be built in the future. While one might question whether 'realisticness' is the best way to assess such models. 'Plausibility' seems to be a more intuitive (and broader) alternative. But the account presented here – as other obvious choices such as Sugden (2000)'s 'credibility' or the MISS of Mäki (2009) – could well serve as a vantage point for an account of model plausibility.

Another advantage of the present account is that it aligns well with the actual practice of many applied scientists. The way realisticness is assessed theoretically can be directly operationalized via the use of established validation techniques, which are commonly used in applied science to determine the 'realisticness' of a model. Such alignment with scientific practice also immediately suggests some further avenues for philosophical scrutiny: first, many validation techniques directly refer to data that have themself been processed or even produced by other models. Data on aggregate production or unemployment in economics are typical examples. Thus, the validation of one model often relies on another model producing adequate data, which presupposes validation of this latter model. The theoretical and practical implications of this 'nested validation' for accounts of model realisticness, but also model-world comparisons in general, are certainly worth exploring in future work.

Another interesting avenue is to relate the present idea of 'realisticness' to other alternatives in the literature (for a small survey see Mäki, 2017). Of direct relevance for the proposal at hand are 'substitute models' as defined by Mäki (ibid.): models that denote targets in reality, but are not explored with respect to their relation to these targets (i.e. they are not validated), but only with respect to their internal dynamics. Such models are widely used in scientific practice, in particular in economics, where they provide 'causal mechanism schemes' a la Ylikoski and Aydinonat (2014) to be used as components for more complex models. As such, they are clearly unrealistic since once they were compared to their targets, the fit would – in accordance with their construal – be poor. The way the resulting 'causal mechanism schemes' are aggregated to more complicated models that are then more 'realistic' in the sense advocated in the present paper would then be an interesting avenue for future research, as would be a comparison with 'minimal models' that lack any target right from the start (Grüne-Yanoff, 2009). Also, the approach taken to realisticness here is a modular one: one assesses the realisticness of various parts of the model and aggregates them to an overall 'realisticness score' (see equation (3)). Alternatively, one could start from a more structural view of models, which would suggest to compare a model to its target as a whole. This would certainly lead to a different account of realisticness, an avenue that seems to be worth exploring. Finally, many scientists do not relate a single model to reality, but use multiple models at once (Aydinonat, 2018). It seem worthwhile to explore how the joint application of models affects the validation of the individual models, or whether such sets of models can have something like a 'joint realisticness'.

In all, the present contribution has introduced one coherent account of model realisticness, which is built upon an anti-realist fictional view of models, realism about mathematical objects, and the DEKI account of representation. It is certainly desirable to develop alternative accounts, not resting on these particular assumptions, but it was also shown that in any alternative account of model realisticness would necessarily rest upon equally strong commitments in the subject areas discussed above.

References

Apostel, Leo (1960). “Towards the formal study of models in the non-formal sciences”. In: *Synthese* 12.2-3, pp. 125–161.

Aydinonat, N Emrah (2018). “The diversity of models as a means to better explanations in economics”. In: *Journal of Economic Methodology* 82.3, pp. 1–15.

Boldyrev, Ivan and Alexey Ushakov (2016). “Adjusting the model to adjust the world: constructive mechanisms in postwar general equilibrium theory”. In: *Journal of Economic Methodology* 23.1, pp. 38–56.

Bravo, Giangiacomo et al. (2015). “Intermediaries in Trust: Indirect Reciprocity, Incentives, and Norms”. In: *Journal of Applied Mathematics* 2015.3, pp. 1–12.

Friedman, Milton (1966). “The Methodology of Positive Economics”. In: *Essays in Positive Economics*. Chicago, IL: University of Chicago Press, pp. 3–43.

Frigg, Roman (2010). “Models and fiction”. In: *Synthese* 172.2, pp. 251–268.

Frigg, Roman and James Nguyen (2016). “The Fiction View of Models Reloaded”. In: *The Monist* 99.3, pp. 225–242.

– (2017). “Models and Representation”. In: *Springer Handbook of Model-Based Science*. Ed. by Lorenzo Magnani and Tommaso Bertolotti. Dordrecht, Heidelberg, London and New York, pp. 49–102.

Gelfert, Axel (2017). “The Ontology of Models”. In: *Springer Handbook of Model-Based Science*. Ed. by Lorenzo Magnani and Tommaso Bertolotti. Dordrecht et al.: Springer, pp. 5–24.

Giere, Ronald N (1990). *Explaining Science: A Cognitive Approach*. Chicago and London: University of Chicago Press.

– (2010). “An agent-based conception of models and scientific representation”. In: *Synthese* 172.2, pp. 269–281.

Godfrey-Smith, Peter (2006). “The strategy of model-based science”. In: *Biology & Philosophy* 21.5, pp. 725–740.

– (2009). “Models and fictions in science”. In: *Philosophical Studies* 143.1, pp. 101–116.

Gräbner, Claudius (2018). “How to Relate Models to Reality? An Epistemological Framework for the Validation and Verification of Computational Models”. In: *Journal of Artificial Societies and Simulation* 21.3.

Grüne-Yanoff, Till (2009). “Learning from Minimal Economic Models”. In: *Erkenntnis* 70.1, pp. 81–99.

Grüne-Yanoff, Till and Paul Weirich (2010). “The Philosophy and Epistemology of Simulation: A Review”. In: *Simulation & Gaming* 41.1, pp. 20–50.

Hempel, Carl G and Paul Oppenheim (1948). “Studies in the Logic of Explanation”. In: *Philosophy of Science* 15.2, pp. 135–175.

Hughes, Robin (1997). “Models and representation”. In: *Philosophy of Science* 64.4, pp. 325–336.

Klügl, Franziska (2008). “A validation methodology for agent-based simulations”. In: *the 2008 ACM symposium*. New York, New York, USA: ACM Press, p. 39.

Knuuttila, Tarja (2011). “Modelling and representing: An artefactual approach to model-based representation”. In: *Studies in History and Philosophy of Science Part A* 42.2, pp. 262–271.

Knuuttila, Tarja and Rami Koskinen (2017). “Synthetic Biology’s Alternative Realities – Turning Fictional Systems into Concrete Ones”. In: *Mimeo*.

Levins, Richard (1966). “The Strategy of Model Building in Population Biology”. In: *American Scientist* 54.4, pp. 421–431.

Mäki, Uskali (1998). “Realisticness”. In: *The Handbook of Economic Methodology*. Ed. by John B Davis et al. Cheltenham, UK, pp. 409–413.

– (2009). “MISSing the World. Models as Isolations and Credible Surrogate Systems”. In: *Erkenntnis* 70.1, pp. 29–43.

– (2017). “Highly unrealistic models: their unrealisticness and their functions”. In: *Mimeo*.

Matthewson, John and Michael Weisberg (2008). “The structure of tradeoffs in model building”. In: *Synthese* 170.1, pp. 169–190.

Nguyen, James (2016). “How Models Represent”. PhD thesis. London: London School of Economics and Political Science.

Nguyen, James and Roman Frigg (2017). “Mathematics is not the only language in the book of nature”. In: *Synthese* 60.1, pp. 135–22.

Odenbaugh, Jay (2015). “Semblance or similarity? Reflections on Simulation and Similarity”. In: *Biology & Philosophy* 30.2, pp. 277–291.

Rusanen, Anna-Mari and Otto Lappi (2012). “An Information Semantic Account of Scientific Models”. In: *Epsa Philosophy of Science Amsterdam*. Ed. by Henk W de Regt, pp. 315–327.

Salis, Fiora (2016). “The Nature of Model-World Comparisons”. In: *The Monist* 99.3, pp. 243–259.

– (2018). “Scientific Discovery Through Fictionally Modelling Reality”. In: *Topoi* 67.3, pp. 380–11.

Smajgl, A and E Bohensky (2013). “Behaviour and space in agent-based modelling: Poverty patterns in East Kalimantan, Indonesia”. In: *Environmental Modelling & Software* 45, pp. 8–14.

Stachurski, John (2016). *A Primer in Econometric Theory*. Cambridge, MA: The MIT Press.

Steel, Daniel P (2008). *Across the Boundaries. Extrapolation in Biology and Social Science*. Oxford, UK, New York, NY: Oxford University Press.

Suárez, Mauricio (2004). “An Inferential Conception of Scientific Representation”. In: *Philosophy of Science* 71.5, pp. 767–779.

Sugden, Robert (2000). “Credible worlds: the status of theoretical models in economics”. In: *Journal of Economic Methodology* 7.1, pp. 1–31.

Tesfatsion, Leigh (2017). “Modeling economic systems as locally-constructive sequential games”. In: *Journal of Economic Methodology* 24.4, pp. 384–409.

Toon, Adam (2012). *Models As Make-Believe: Imagination, Fiction, and Scientific Representation*. Hampshire, UK: Palgrave Macmillan.

Walton, Kendall (1990). Mimesis as make-believe: On the foundations of the representational arts. Cambridge, MA: Harvard University Press.

Weisberg, Michael (2007). “Who is a Modeler?” In: *The British Journal for the Philosophy of Science* 58.2, pp. 207–233.

– (2012). “Getting Serious about Similarity”. In: *Philosophy of Science* 79.5, pp. 785–794.

– (2013). *Simulation and Similarity*. New York, NY: Oxford University Press.

Ylikoski, Petri and N Emrah Aydinonat (2014). “Understanding with theoretical models”. In: *Journal of Economic Methodology* 21.1, pp. 19–36.