

ROBUSTNESS ANALYSIS DISCLAIMER:
PLEASE READ THE MANUAL BEFORE USE!

Jaakko Kuorikoski*, Aki Lehtinen* and Caterina Marchionni*

*Social and Moral Philosophy, University of Helsinki

Corresponding author:

Aki Lehtinen, Social and Moral Philosophy, University of Helsinki

P.O. Box 24 (Unioninkatu 40a), FI-00014 Helsinki, Finland

Email: aki.lehtinen@helsinki.fi

ROBUSTNESS ANALYSIS DISCLAIMER:
PLEASE READ THE MANUAL BEFORE USE!

Abstract. Odenbaugh and Alexandrova (2011) provide a challenging critique of the epistemic benefits of robustness analysis, singling out for particular criticism the account we articulated in Kuorikoski, Lehtinen and Marchionni (2010). Odenbaugh and Alexandrova offer two arguments against the confirmatory value of robustness analysis: robust theorems cannot specify causal mechanisms and models are rarely independent in the way required by robustness analysis. We address Odenbaugh and Alexandrova's criticisms in order to clarify some of our original arguments and to shed further light on the properties of robustness analysis and its epistemic rationale.

Keywords Modelling, Robustness, Idealisation, Causal mechanisms

Introduction

Instead of trying to include ever more realistic information about the system under study, theoretical modellers in biology and economics often simply replace one set of unrealistic assumptions with another in order to derive the same or similar results. Can this practice be justified from an epistemic point of view? Along with Richard Levins (1966), William Wimsatt (1981, 2007) and Michael Weisberg (2006), we have argued that in economics this practice has an epistemic rationale in terms of

derivational robustness analysis (Kuorikoski, Lehtinen and Marchionni 2010, henceforth KLM).

Our argument for the epistemic virtues of derivational robustness can be briefly summarised as follows. Given that theoretical modellers typically aim to capture the functioning of some causal mechanism of interest, they try to represent it by mathematically formalising its key elements. We call these representational formalisations *substantive assumptions*. However, the tractability of models also requires introducing further assumptions (a.k.a. *tractability assumptions*) that are known to be false or lack a meaningful empirical interpretation altogether. Derivational robustness analysis, the comparison of models with slightly different assumptions, provides information on which assumptions are really necessary for a given result: if a result can be derived with different tractability assumptions while retaining the substantive assumptions, this suggests that the result does not depend on the problematic tractability assumptions. In this way, robustness analysis justifiably increases our degree of confidence in the ‘robust theorem’ connecting the substantive assumptions to a result.

In a recent contribution to this journal, Jay Odenbaugh and Anna Alexandrova (2011) argue against the epistemic virtues of robustness analysis and specifically target our account. They claim that robustness has no epistemic import and ‘it is best regarded as a method of discovery rather than confirmation’ (p. 757). In particular, they offer two arguments against the confirmatory status of derivational robustness analysis.

(1) *Robust theorems cannot specify causal mechanisms.* Given that some idealisations always remain undischarged by robustness analysis, the empirical status of ro-

bust theorems remains questionable and robust theorems cannot specify causal mechanisms.

(2) *Models are rarely independent.* The epistemic import of derivational robustness hinges crucially on the independence of models, but reports of independence have been exaggerated.

According to Odenbaugh and Alexandrova, the only rationale for derivational robustness analysis is that it may serve a useful heuristic purpose in formulating schemas, ‘open formulas’ that can then be filled in with appropriate causal content in order to provide actual and testable causal explanations.

Their reading of our proposal for the epistemic import of robustness analysis rests on some misunderstandings, however. This reply discusses both their criticism of our account and their own positive proposal. The aim is to shed further light on the properties and functioning of derivational robustness analysis and its epistemic rationale.

We proceed as follows. In the next section we restate the general features of derivational robustness analysis. In particular, we clarify the relationship between robustness analysis and confirmation and the notion of a ‘robust theorem’. In Section 3 we address Odenbaugh and Alexandrova’s two main criticisms. Section 4 discusses the ‘open formula’ account. Section 5 provides the conclusions.

Derivational robustness analysis: a restatement of general properties

Robustness analysis and confirmation

Let us start with some of the claims made by Odenbaugh and Alexandrova with which we agree. First, any respectable empiricist would want to steer clear of the

unqualified claim that a modelling exercise without any new empirical data could amount to a ‘method of empirical confirmation’. Otherwise modelling would be akin to a form of mystical divination. We do not hold such a view, and nowhere in KLM is it stated that derivational robustness analysis confers empirical confirmation – quite the opposite (see KLM, p. 543). We also agree that the only causal knowledge derivable from formal modelling is that already implicit in the substantial assumptions of the models (cf. Odenbaugh and Alexandrova, p. 762). One can of course restrict the use of the term ‘confirmation’ to refer only to the comparison of a model with data. If confirmation is taken to mean the severe testing of a hypothesis by an empirical datum, then derivational robustness analysis is obviously not confirmatory. Our point is that such testing is not the only epistemically important inferential step in model-based reasoning. Modelling is a form of fallible inference from prior causal assumptions to new causal conclusions facilitated by external inferential aids (i.e., models)¹. The epistemic import of derivational robustness lies in making these inferences more reliable: if we discover that our conclusions depend less on assumptions already known to be problematic, we become justifiably more confident about the conclusions. Whether or not one calls this rational increase in our degree of belief ‘confirmation’ is largely a terminological matter.

What is not a terminological matter is that such inferences are epistemically important. Those who deny that robustness has anything to do with the epistemic evaluation of models disagree about this. Odenbaugh and Alexandrova seem to think that the only epistemically relevant inference is the one in which the consequences of the

¹ This view is further elaborated in Kuorikoski and Lehtinen (2009).

model are tested against data. If such a viewpoint is taken seriously, then prior assessment of how realistic the assumptions are does not matter. But this is not sensible. It surely makes much more sense to test a model-result derived from reasonably credible assumptions than one derived from, say, the assumption that cats are remotely controlled Martian robots. It is not that Odenbaugh and Alexandrova explicitly advocate recklessly neglecting unrealistic assumptions, quite the contrary, but if one takes empirical testing to be the only epistemically relevant modelling activity, such neglect necessarily follows.

Odenbaugh and Alexandrova posit that ‘there is no reason to believe a proposition just because a set of false beliefs imply some proposition’ (p. 761). This is certainly true in the sense that *mere* robustness should not make us believe a theorem in the absence of empirical support for any of the model’s assumptions. Indeed, robustness analysis is sensible only when the substantive assumptions are at least reasonably realistic. It is only against the background of some realistic core assumptions that confidence in the robust theorem increases through derivational robustness analysis.

Robust theorems

Odenbaugh and Alexandrova offer two counterexamples to show that robustness does not guarantee truth: the revenue equivalence theorem(s) in auction theory and Horn’s Markov chain model of forest succession (for details, see Odenbaugh and Alexandrova 2011, pp. 766-768). The revenue equivalence theorem proves, roughly, that different auction mechanisms having certain properties yield the same expected revenue. The problem Odenbaugh and Alexandrova identify with this theorem is that it is known to lack robustness with respect to the assumption that players’ infor-

mation sets, and thus valuations, are independent. Without this assumption, revenue equivalence does not hold. Within our framework, talking of the robustness of a result is shorthand for talking about the robustness of the connection between the substantial assumptions and the outcome jointly predicted by the alternative models. If the substantial assumptions on which the outcome depends are false, what we learn from robustness analysis is precisely that the predicted outcome depends on assumptions known to be false. Therefore, the revenue equivalence theorem does not constitute a counterexample to our account of robustness because the theorem is *known* to depend crucially on an assumption that is known to be false. In such a case, derivational robustness does not license confidence in a theorem despite its robustness with respect to a whole host of other assumptions.

In addition to the revenue equivalence theorem, there are several theorems in economics and biology that are robust with respect to a variety of tractability assumptions and yet false in the sense that the outcome never occurs in reality. However, in many such cases what is actually robust is that the consequents of the theorems (the predicted outcome) can no longer be derived if key assumptions known to be false are replaced by more realistic ones. There is thus a robust result, viz. that the theorem no longer follows if some substantial assumptions are changed. Rather than intended to be descriptions of actually occurring phenomena, such theorems are often treated as useful theoretical baselines that provide a point of comparison with the reality.

For example, the first theorem of welfare economics roughly states that if various assumptions about the structure of markets hold, market equilibrium is efficient (i.e. it is not possible to increase anybody's utility without decreasing someone else's

utility).² Many of these assumptions are obviously unrealistic: it is assumed that there are no public goods (e.g., lighthouses, bridges, safety in the streets) and no externalities (e.g., pollution) and that individuals have complete information. Even the scholars who developed the theory (notably, Frank Hahn 1970) interpret the theorem as implying that real markets are *not* efficient. The theorem is then used to explain what features of actual markets are responsible for their inefficiency.³ Similarly, in biology the Hardy-Weinberg equilibrium is not robust with respect to the inclusion of factors like selection, drift and meiotic drive and is therefore strictly speaking false of all natural populations. Nevertheless, the Hardy-Weinberg theorem is used as a theoretical baseline in model-based investigations of the effects of those factors omitted from the theorem.

The existence of such apparently ‘false’ theorems no more proves that derivational robustness analysis is useless or non-confirmatory than the existence of false theories that have received empirical support proves that empirical testing is useless. What it does prove is (i) that robustness alone does not guarantee the truth of a robust theorem⁴ and (ii) that not all theorems are robust theorems à la Levins.

² Other prominent examples are the Modigliani-Miller theorem in the theory of finance and the Heckscher-Ohlin theorem in the theory of international trade.

³ Wimsatt’s discussion of neutral models captures these features of such ‘false theorems’. Neutral models are neutral in the sense that they provide a ‘null’ baseline situation which is rarely or never actualised; however, they are useful in providing contrasts for models and explanations (see also Hindriks 2008, Wimsatt 2007, p. 94-132).

⁴ Indeed Wimsatt (see also Calcott 2011, 1981, 2007), who also has argued for the epistemic benefits of robustness, has emphasised this.

In Odenbaugh and Alexandrova's other example, Horn's Markov chain model of forest succession was tested for robustness with respect to one false assumption, viz. that dead trees are immediately replaced by new ones. However, Horn either neglected or was unable to perform the same trick with other false assumptions. As we explained in KLM, confidence in a result comes in degrees and depends on the extent to which it has actually been subjected to robustness analysis. Checking the robustness of the result with respect to a single assumption in the presence of other assumptions that are known to be false is certainly not enough to render a theorem justifiably believable. In our account, derivational robustness analysis is conceived as the repeated modification of a model by a community of modellers - a process that takes place over time and involves several modelling exercises (i.e. a family of models) but which never reaches the ideal state of completion, or what Odenbaugh and Alexandrova nicely term 'absolute robustness analysis'.

Odenbaugh and Alexandrova also pose the following related question: 'So how can the fact that your standard toolbox yielded a couple of different tractability assumptions that do imply the result provide genuinely independent evidence?' (p. 763) Of course it does not. As noted above, confidence in a theorem often increases when several modelling exercises are performed by scientists other than those who originally proposed the model. This communal aspect is not tangential to the epistemic value of robustness. First, variation in background knowledge and modelling skills makes it more likely that diverse kinds of auxiliary assumptions are tried out, thereby alleviating worries of lack of independence. Second, as we will see below, the communal aspect of robustness analysis plays a role in ensuring that attempts are also made to prove the *lack* of robustness of a theorem.

The causal claims represented in robust theorems should be independently empirically tested whenever possible (e.g., KLM, p. 549). However, this cannot always be done. In economics and biology (and especially in ecology), much of the data are purely observational, rather than experimental, and modelling is the only way of trying to make causal sense of it. Derivational robustness analysis can be useful precisely in those cases where direct, model-independent empirical testing of model-based causal claims is not feasible.

Things to check before applying

Now that we have got the terminological red herring of confirmation out of the way, we can concentrate on Odenbaugh and Alexandrova's more substantial worries about derivational robustness analysis: that robust theorems cannot specify causal mechanisms, and that models are rarely independent in the way required by robustness analysis.

Robust theorems cannot specify causal mechanisms

Odenbaugh and Alexandrova claim that since derivational robustness analysis can never defuse all unrealistic assumptions, robust theorems 'cannot specify a causal mechanism'. In their own words:

Kuorikoski et al. ... argue the common structure of a family of models is a representation of a causal mechanism that explains the relevant phenomenon (2010, 14). But in order to defend this interpretation of robust theorems, the

advocate needs to claim that all the false assumptions of the family of models in question have been discharged ... Our worry reduces to this: if we are worried about idealizations simpliciter we need absolute robustness analyses and when we are worried only about certain idealizations relative robustness analyses are sufficient. However, robustness analyses only have confirmatory value when we can perform absolute robustness analyses. (ibid., p. 764)

The best way of refuting their argument is to break it down and reveal what we regard as questionable presuppositions. Odenbaugh and Alexandrova argue that (i) unless all of a model's assumptions are de-idealised or discharged by absolute robustness analysis, derivational robustness analysis cannot capture a causal mechanism. However, (ii) given that it is rare that all assumptions are de-idealised or discharged in such a way, it follows that (iii) models are incapable of depicting causal mechanisms. In other words, according to Odenbaugh and Alexandrova, it would be possible to interpret the model as depicting a causal mechanism if we could ensure that one of the various auxiliary assumptions we try were the true one and if we could do this for all questionable auxiliary assumptions. However, since we can *rarely* do this, derivational robustness has little or no epistemic relevance.

We highlight the word 'rarely' because this is what we fundamentally disagree with: it is not that it is rare that every assumption of a model is discharged or de-idealised, rather it is *never* the case that all assumptions are either de-idealised or discharged. The only model that could do so would be reality itself. Consequently, Odenbaugh and Alexandrova must think that models never depict causal mechanisms

– which indeed seems to be the point of the open formula account of models they endorse (see also Alexandrova 2008, 2009) and that we will discuss in Section 3.

What would be required for a model to depict a causal mechanism? Given that Odenbaugh and Alexandrova claim that the epistemic benefits of derivational robustness analysis hinge on the ability to test the robustness of each and every assumption (absolute robustness analysis), they apparently endorse something like the ‘perfect model model’ (Teller 2001), according to which, every part of a model has to be true for it to be an adequate representation of anything.

The perfect model model is highly questionable. If even a single idealisation is sufficient to render a model representationally inadequate, then no model can be used to represent causal mechanisms. However, this applies not only to models but to representation more generally. If the slightest distortion of the thing represented means the failure of that representation, then it would be simply impossible to represent causal mechanisms at all. Such a view immediately leads to absurdities. It implies, for example, that if a model uses the gravitational constant $G=6.653645 \times 10^{-11}$ rather than the correct 6.67384 in representing gravitation, it does not represent a causal force at all. Note also that, using the logic of the argument requiring absolute robustness analysis, it would be equally plausible to say that de-idealisation does not get us any closer to representing a mechanism either, unless one is able to remove *all* the idealisations included in the model.

Models are rarely independent

Odenbaugh and Alexandrova’s second substantial worry is that false modelling assumptions are in fact never independent in the way required by our account. As a

reminder, in KLM we explicated independence in Levins' famous slogan 'our truths are the intersection of independent lies' as requiring that, from the point of view of the modellers, the alternative tractability assumptions induce different kinds of falsities in models. We argued that with a suitable interpretation of independence of assumptions, derivational robustness analysis can be seen as a limiting case of general robustness analysis in Wimsatt's sense, that is, as a form of triangulation between independent means of determination. Our independence argument however is in no *other* way directly dependent on a flimsy analogy between derivational robustness and multimodal triangulation. In fact, we fully agree with Odenbaugh and Alexandrova that the case of multiple independent sources of empirical support is a much stronger form of, well, empirical support (p. 762), but this is not inconsistent with our argument.

To explain our notion of independence of assumptions, let us briefly consider two illustrations, one from economics and one from biology. Consider first the example of transport costs that we used in KLM. A given result in geographical economics (that spatial agglomeration occurs when economies of scale are high, market power is strong, and transportation costs are low) was shown to be robust with respect to whether transport costs were assumed to be linear (with respect to distance) or of the 'iceberg' form (where a fraction of the good itself is assumed to 'melt' during transit). These two assumptions are definitely not independent with respect to all their features: for example, both model transport costs by assuming that costs depend positively on distance - which is a realistic assumption. What should be independent are the lies that the alternative specific mathematical implementations of this positive dependence between cost and distance bring into the model, i.e., their 'icebergness'

and linearity. Icebergness and linearity are independent if there are no (prior) reasons to think they have a similar mathematical and empirically interpretable impact on the modelling result over and above that of the positive dependence between cost and distance. Whether or not agglomeration occurs depends not on the specific details of transport costs (i.e., whether they are linear or of the iceberg-form) but on the size of those costs.⁵ Note that being able to formulate such judgments does not require knowing in advance how ‘far’ the two assumptions are from the ‘true’ assumption. In fact, in many cases there may be no such thing as the one true assumption (KLM, p. 543).

As another example consider the much-discussed case of the Volterra rule or principle, according to which a general biocide introduced to a predator-prey system will increase the relative abundance of prey. According to Weisberg and Reisman (2008), this general rule can be derived from a class of Lotka-Volterra models with different functional forms for the rate of prey capture per predator and for the relation between predator births and the number of prey captured respectively, with or without population stochasticity, and with or without the possibility of prey or predator satiation. Whereas many of these modifications are naturally seen as possible substantial assumptions (e.g., whether or not there is prey satiation is surely a plausible candidate for an empirically relevant determinant of overall population dynamics), the way in which they are implemented always introduces tractability considerations. As in the case of transportation costs, any functional form for, say, the rate of prey capture per predator (i.e., for the functional response) is strictly speaking false of any natural

⁵ Clearly we remain agnostic regarding the validity of the geographical economics’ result.

population. So, according to our account two alternative functional forms for the functional response are independent if there are no prior reasons to believe that the exact forms of these functions mathematically affect the model result in a similar way.

Weisberg and Reisman (2008) also discuss a way in which practically all the tractability assumptions can be expected to be independent: the derivation of the robust theorem in a completely different modelling framework. Whereas the class of Lotka-Volterra models described above are sets of differential equations relating population aggregates, the Volterra principle can also be demonstrated using agent-based computational models. Such models represent the *same* core causal mechanisms, albeit describing them at an individual level. However, the radical difference in the modelling framework means that the tractability assumptions, although still unavoidable, are of an altogether different kind: they relate to the behavioural rules of individuals and the spatial representation of their environment, rather than to population-level generalisations as in the original Lotka-Volterra models.

Odenbaugh and Alexandrova's general counterargument against the relevance of independence is that the probability of the conjunction of a set of fully independent claims in which even one is false remains zero. Their analysis is similar to Cartwright's (1991) early view in the sense that it rests on the idea that there must be a correct true assumption for everything that we want to represent in a model. This argument presupposes that it makes sense to evaluate the truth-status of models by considering the probabilities of "truth" of models considered as conjunctions of assumptions. We don't think it does. If a known falsity in any assumption, no matter how slight and how irrelevant for the result, renders the probability of the whole

model zero, then all models have a probability of zero – which is an unacceptable conclusion. Their argument thus only makes sense if model development is considered a process that approaches a model that only contains true assumptions. If one presumes that each unrealistic assumption can and should be replaced with a single unique true assumption, there is no point in conducting robustness analysis in the first place. The task is rather that of finding true assumptions – and this is clearly not feasible in the majority of cases. Robustness analysis is about coping with unavoidable falsity rather than finding *the* truth.

Odenbaugh and Alexandrova also offer two reasons for thinking that independence always fails in practice. First, they worry that the very fact that false assumptions are systematically selected so as to accommodate the robust theorem renders them dependent. We do not know whether there is a systematic modelling or publication bias towards similar models that establish rather than discredit the robustness of important theorems. If anything, we would wager the opposite to be true. Although individual modellers might have only a limited set of assumptions in mind and a strong bias for proving robustness, the incentive to demonstrate the fragility of the results of rival models acts as a counterweight to this tendency. This worry is therefore another reason why we stress the communal nature of derivational robustness analysis.

Second, Odenbaugh and Alexandrova fear that many theorems are robust precisely because of the presence of a common modelling framework (p. 763).⁶ Although

⁶ A similar worry was voiced by Nancy Cartwright regarding economics. She argued that the common economic modelling framework “overconstrains” the modelling results in a detrimental way (Cartwright 2009).

this empirical worry does not amount to a counterargument to our general analysis of the epistemic import of derivational robustness, we do share this concern - at least to some extent. If modelling within a field systematically (and dogmatically) rests on a set of assumptions that is never questioned, regardless of their empirical performance, no amount of robustness analysis can get us closer to confirming or even discovering causal mechanisms. Nevertheless, we do not think that modelling in economics or biology represents such a degenerate case. The wider use of different kinds of modelling frameworks (e.g., agent-based modelling, as in the case of the Volterra principle) might very well be advisable in ecology and economics, but different modelling frameworks just mean different, but not necessarily fewer, idealisations, and hence more work for derivational robustness analysis.

Consumer report: Derivational robustness analysis vs. the perfect model model and open formulae

If the presence of any false idealisation indeed implies that a model is an inadequate representation of a causal mechanism, what is the cognitive value of highly idealised theoretical models in biology and economics? According to Odenbaugh and Alexandrova, even though models do not represent causal mechanisms, they serve as templates, which can in turn be filled in with context-dependent causal details, thereby producing causal explanations. A template for a causal claim would be as follows: ‘In a situation x with some characteristics that may or may not include $\{C_1, \dots, C_n\}$, a certain feature F causes a certain behaviour B .’ (Odenbaugh and Alexandrova 2011, p. 769) What Odenbaugh and Alexandrova take to be the crucial difference with a

robust theorem is that the template neither quantifies over the situation nor specifies the conditions of application.

In their view, the heuristic import of robustness analysis manifests itself in the process of constructing such templates. Performing derivational robustness analysis is a means to identify the template itself, namely what kind of F (possibly) causes what kind of B, and in which possible situations. This view is similar to Patrick Forber's take (2010) on the role of robustness analysis in biology, which, according to him, serves only to delineate the very general limits of what is biologically possible, whereas confirmation happens only when specific alternative models are contrastively compared against data. According to Odenbaugh and Alexandrova, however, a template is turned into a causal explanation only when the situation and conditions are specified, and the model itself does not provide any information about this. If the template does not have any causal interpretation on its own, what principles guide this specification? Fit with the data is certainly one such criterion, but it cannot be the only one, since theoretical modelling and statistical curve fitting are quite different enterprises: the point of theoretical models is not just to capture observed relations between variables, but also to show why they are so related. If the template itself does not have any causal content, it remains a mystery how by specifying the template *new* causal knowledge is created over and above that which is already assumed in the specification. If instead the specification of the model does not produce any new causal conclusions, what is the point of using the model in the first place?

We are not proposing that a model's derivational robustness determines whether or not it is causal. Our suggestion is that theoretical model templates provide causal explanantia if their structure roughly represents the causal mechanism expected to be

in operation in some circumstances. Economists use models based on utility maximisation and equilibrium because they believe that the phenomena they are modelling are, by and large, realised by actors pursuing their ends in a strategic environment (either intentionally, due to selection, or as if institutionally programmed). Similarly, ecologists use templates based on different populations and environmental circumstances because they believe that their complex interactions are responsible for the phenomena under study and that some model templates capture these interactions better than others.

Whether these rough representations amount to a specification of a mechanism depends on what one means by ‘specification’. Even Odenbaugh and Alexandrova admit that templates can rarely, if ever, be specified totally, i.e., the scope of x and all the conditions can never be exhaustively delineated. This means that there always remains an inductive gap in using even filled-in templates to generate causal explanations of particular phenomena. There is no reason, however, why such partial specifications cannot be regarded as perfectly adequate causal explanations. Until Odenbaugh and Alexandrova clarify exactly what is meant by ‘specification’ and why modelling cannot achieve it, we do not really see a substantive difference between our positions.

Conclusions

Our account of derivational robustness analysis makes rational sense of the practice of building and comparing models that differ only with respect to a few, often equally unrealistic, assumptions. This should not be taken as a blanket defence of such

modelling practice, however. Derivational robustness analysis cannot and should not replace empirical tests. Neither is it a universal remedy to the problems in the modelling practices of economists and biologists.

Acknowledgments. Parts of this reply have been presented at the workshop “Economic models: their nature and significance”, University of East Anglia (UK); at a seminar jointly organized by the Department of Economics and the Department of Mathematics, University of Florence, and at the Philosophy of Science seminar, University of Helsinki. We thank the participants in these events for their insightful comments and discussions. The usual disclaimers apply.

References

Alexandrova A (2009) When analytic narratives explain. *Journal of the Philosophy of History* 3:1-24

Alexandrova A, Northcott R (2009) Progress in economics: Lessons from the Spectrum Auctions. In: Kincaid H, Ross D (eds) *Oxford Handbook of Philosophy of Economics*. Oxford University Press, Oxford

Alexandrova A (2008) Making models count. *Philosophy of Science* 73:383-404

Calcott B (2011) Wimsatt and the robustness family: Review of Wimsatt’s Re-engineering Philosophy for Limited Beings. *Biology and Philosophy* 26:281-293

Cartwright N (2009) If no capacities then no credible worlds. But can models reveal capacities?. *Erkenntnis* 70:45-58

Cartwright N (1991) Replicability, reproducibility, and robustness: Comments on Harry Collins. *History of Political Economy* 23:143-155

Forber P (2010) Confirmation and explaining how possible. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 41:32-40

- Hahn FH (1970) Some adjustment problems. *Econometrica* 38:1-17
- Hindriks FA (2008) False models as explanatory engines. *Philosophy of the Social Sciences* 38:334-360
- Kuorikoski J, Lehtinen A (2009) Incredible worlds, credible results. *Erkenntnis* 70:119-131
- Levins R (1966) The strategy of models building in population biology. *American Scientist* 54:421-431
- Odenbaugh J, Alexandrova A (2011) Buyer beware: robustness analyses in economics and biology. *Biology & Philosophy* 26:757-771
- Teller P (2001) Twilight Of The Perfect Model Model. *Erkenntnis* 55:393-415
- Weisberg M (2006) Robustness analysis. *Philosophy of Science* 73:730-742
- Wimsatt WC (2007) Re-engineering philosophy for limited beings. Harvard University Press, Cambridge, Mass. ; London
- Wimsatt WC (1981) Robustness, reliability and overdetermination. In: Brewer MB, Collins BE (eds) *Scientific inquiry and the social sciences*. Jossey-Bass, San Francisco, pp 124-163