Biol Philos (2012) 27:891–902
DOI 10.1007/s10539-012-9329-z

Robustness analysis disclaimer: please read the manual before use!

Jaakko Kuorikoski · Aki Lehtinen · Caterina Marchionni

Received: 15 December 2011 / Accepted: 13 June 2012 / Published online: 21 June 2012

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 et al. (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 Levins (1966), Wimsatt (1981, 2007) and Weisberg (2006), we have argued that in economics this practice has an epistemic rationale in terms of *derivational robustness analysis* (Kuorikoski et al. 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

J. Kuorikoski · A. Lehtinen (✉) · C. Marchionni
Social and Moral Philosophy, University of Helsinki, P.O. Box 24, Unioninkatu 40a, 00014 University of Helsinki, Finland
e-mail: aki.lehtinen@helsinki.fi

formalisations *substantial assumptions*. However, the tractability of models also requires further assumptions (i.e. *tractability assumptions*) that are known to be false or altogether lack a meaningful empirical interpretation. 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 substantial 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 substantial assumptions to a result.

In a recent contribution to this journal, Odenbaugh and Alexandrova (2011) argue against the epistemic virtues of robustness analysis and specifically target our account. They claim that robustness has no epistemic import and '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 robust theorems remains questionable and robust theorems cannot specify causal mechanisms.
(2) *Models are rarely independent*. The epistemic import of derivational robustness hinges 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 as a useful heuristic for formulating templates, that is, 'open formulas' that can be filled in with appropriate causal content in order to provide actual and testable causal explanations. However, their reading of our proposal for the epistemic import of robustness analysis rests on some misunderstandings. 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 the third section we address Odenbaugh and Alexandrova's two main criticisms. The fourth section discusses the 'open formula' account.

Derivational robustness analysis: a restatement of general properties

Robustness analysis and confirmation

Let us start with some 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 concur that the only causal knowledge derivable from formal modelling is implicit in the substantial assumptions of the models, and derivational robustness does not take us beyond our causal background knowledge (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 strictly taken to mean testing a hypothesis by empirical data, 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 robustness a role in the epistemic evaluation of models disagree. Odenbaugh and Alexandrova seem to think that the only epistemically relevant inference is that in which the consequences of the 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 more sense to devote time and resources to empirically test a model-result derived from reasonably credible assumptions than one derived from, say, the patently absurd 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.

We also assent to Odenbaugh and Alexandrova’s claim 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 substantial assumptions are at least reasonably realistic. It is only against the background of some realistic core assumptions that derivational robustness analysis can increase confidence in the robust theorem.

Robust theorems

Now that we have got the terminological red herring of confirmation out of the way, we can concentrate on Odenbaugh and Alexandrova’s examples. Odenbaugh and Alexandrova offer two counterexamples to show that robustness does not guarantee

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

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 given certain conditions, different auction mechanisms yield the same expected revenue. Odenbaugh and Alexandrova rightly point out that this theorem lacks robustness with respect to the assumption that players' information sets, and thus their valuations, are independent. Without this assumption, revenue equivalence does not hold. Odenbaugh and Alexandrova tout this as a challenge to our account. However, within our framework, the robustness of a result is shorthand for 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.

To further illustrate our stance, it is worth mentioning that there are many 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 given various assumptions about the structure of markets, market equilibrium is efficient in that it is not possible to increase someone's utility without decreasing someone else's utility.[2] Many of the assumptions of this theorem are obviously unrealistic: there are no public goods (e.g., lighthouses, bridges, safety in the streets) and no externalities (e.g., pollution) and individuals have complete information. Even scholars (notably, Hahn 1970) who developed the mathematical framework from which the theorem can be derived interpret it as implying that real markets are *not* efficient. The theorem is then used to explain what features of actual markets account 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

[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 (Wimsatt 2007, p. 94–132; see also Hindriks 2008).

strictly speaking false for 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 empirically-supported false theories proves that empirical testing is useless. What it does prove is (1) that robustness alone does not guarantee the truth of a robust theorem[4] and (2) that not all theorems are robust theorems á la Levins.

In their other example, on Horn's Markov chain model of forest succession, Odenbaugh and Alexandrova show that Horn tested 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 explain 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 credible. In our account, derivational robustness analysis is the repeated modification of a model by a community of modellers. It is a process that takes place over time and involves several modelling exercises (i.e. a family of models) but that never reaches the ideal state of completion, the discharging of all "false" assumptions, or what Odenbaugh and Alexandrova term 'absolute robustness analysis'.

This communal aspect is not tangential to the epistemic value of robustness and can be used to illuminate and counter another criticism by Odenbaugh and Alexandrova. They 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 cannot. 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. 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 to try 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.

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

Things to check before applying

Odenbaugh and Alexandrova have two more substantial worries about derivational robustness analysis. They argue that robust theorems cannot specify causal mechanisms, and that models are rarely independent in the way required by robustness analysis. We will now discuss these arguments.

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 untenable presuppositions. Odenbaugh and Alexandrova argue that (1) 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, (2) given that it is rare that all assumptions are de-idealised or discharged, it follows that (3) models are incapable of depicting causal mechanisms. In other words, according to Odenbaugh and Alexandrova, it would be possible to interpret a model as depicting a causal mechanism only if we could either discharge or de-idealise each questionable auxiliary assumption. 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 only rare that every assumption of a model is discharged or de-idealised with 'some true assumption' (p. 764), it is *never* the case. The only model capable of such a feat would be reality itself.[5] Consequently, Odenbaugh and Alexandrova must think that models never depict causal mechanisms—which

[5] Odenbaugh and Alexandrova's notion of absolute robustness is somewhat ambiguous. The distinction between relative and absolute robustness is introduced by appealing to the difference between many versus all assumptions, but in introducing the term 'absolute' they write: 'The only way to remove this worry is to show that there is some true assumption when conjoined with the substantial core that implies the prediction. Call this the ''absolute'' robustness analysis.' (p. 764) If one could attain absolute robustness by replacing each questionable assumption with others that are not necessarily known to be true, then our claim that only reality itself could attain absolute robustness might be false. However, given that they do indeed require that the true assumption is among the tried ones (see also below), we stick to our guns.

indeed seems to be the point of the open formula account of models they endorse (see also Alexandrova 2008, 2009) and that we discuss in the next section.

What would be required for a model to depict a causal mechanism? Given their claim that the epistemic benefits of derivational robustness analysis hinge on finding a true assumption to replace each and every questionable assumption (absolute robustness analysis), Odenbaugh and Alexandrova 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 dubious. 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 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 explicate 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 argue that with a suitable interpretation of the independence of assumptions, derivational robustness analysis can be seen as a limiting case of general robustness analysis à la Wimsatt, 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 account.

To illustrate our notion of the independence of assumptions, let us briefly consider two examples, one from economics and one from biology. Consider first the example of transport costs that we present in KLM. A well-known result in geographical economics is that spatial agglomeration occurs when economies of scale are high, market power is strong, and transportation costs are low. The result holds whether transport costs are assumed to be linear (in distance) or of the 'iceberg' form (where a fraction of the good 'melts' during transit). These two assumptions are not independent with respect to all features: for example, both model transport costs as a positive function of distance—which is a realistic assumption. What should be independent are the lies that the alternative mathematical implementations of this assumption, that is, the 'icebergness' and 'linearity', smuggle into the model. 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 spatial 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.[6] 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, from biology, consider the much-discussed case of the Volterra principle, which states that 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, as well as 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 invariably 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 for any natural 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

[6] Clearly we remain agnostic regarding the validity of the geographical economics' result.

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 probability of "truth" of the conjunction of their assumptions. We do not think it does. If the known falsity of any assumption, no matter how picayune and 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 consisting of only 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, and the sole task is to find the true assumptions. But 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 why independence often fails in practice. First, they worry that false assumptions are rendered dependent by the very fact that they are systematically selected to accommodate the robust theorem of interest. 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. But, if anything, we would wager the opposite. 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 countervails 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).[7] Although 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 never-questioned core assumptions, regardless of their empirical credentials, 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.

[7] 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).

Consumer report: derivational robustness analysis versus 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 is as follows: 'In a situation *x* with some characteristics that may or may not include $\{C_1,\ldots,C_n\}$, a certain feature F causes a certain behaviour B' (Odenbaugh and Alexandrova 2011, p. 769). What Odenbaugh and Alexandrova take to crucially distinguish a template from a causal explanation is that the former 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 what 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, 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. However, 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 are agnostic to 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 neither can nor should replace empirical tests. Neither is it a universal remedy to problems in the modelling practices of economists and biologists. Despite these qualifications, it has an important epistemic rationale.

Acknowledgments This research was financially supported by the Academy of Finland. 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 discussion. We also thank Jani Raerinne and Bradley Turner for their comments on earlier versions of this paper. The usual disclaimers apply.

References

Alexandrova A (2008) Making models count. Philos Sci 73:383–404

Alexandrova A (2009) When analytic narratives explain. J Philos Hist 3:1–24

Calcott B (2011) Wimsatt and the robustness family: review of Wimsatt’s re-engineering philosophy for limited beings. Biol Philos 26:281–293

Cartwright N (1991) Replicability, reproducibility, and robustness: comments on Harry Collins. Hist Polit Econ 23:143–155

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

Forber P (2010) Confirmation and explaining how possible. Stud Hist Philos Sci 41:32–40

Hahn FH (1970) Some adjustment problems. Econometrica 38:1–17

Hindriks FA (2008) False models as explanatory engines. Philos Soc Sci 38:334–360

Kuorikoski J, Lehtinen A (2009) Incredible worlds, credible results. Erkenntnis 70:119–131

Kuorikoski J, Lehtinen A, Marchionni C (2010) Economic modelling as robustness analysis. Brit J Philos Sci 61:541–567

Levins R (1966) The strategy of model building in population biology. Am Sci

Odenbaugh J, Alexandrova A (2011) Buyer beware: robustness analyses in economics and biology. Biol Philos 26:757–771

Teller P (2001) Twilight of the perfect model model. Erkenntnis 55:393–415

Weisberg M (2006) Robustness analysis. Philos Sci 73:730–742

Weisberg M, Reisman K (2008) The robust Volterra principle. Philos Sci 75:106–131
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
Wimsatt WC (2007) Re-engineering philosophy for limited beings. Harvard University Press, Cambridge, Mass, London