

Austinian model evaluation

Philippe van Basshuysen

Manuscript accepted for publication in the PSA2022 Proceedings in *Philosophy of Science*

Abstract. Like Austin’s “performatives”, some models are used not merely to represent, but also to *change* their targets in various ways. This paper argues that Austin’s analysis can inform model evaluation: if models are evaluated with respect to whether they are adequate for particular purposes (Parker 2020), and if performativity can in some cases be regarded as a model purpose (a proposition that is defended, using mechanism design as an example), it follows that these models can be evaluated with respect to their “felicity”, i.e. whether their use has achieved this purpose. Finally, I respond to epistemic and ethical concerns that might block this conclusion.

Acknowledgments. For helpful comments and discussions, I thank Markus Ahlers, Elinor Clark, Mathias Frisch, Donal Khosrowi, Lucie White, and Jannik Zeiser.

1. Introduction. “I apologise...”, “I bet...”, and “I promise...” are, following J. L. Austin, “performative utterances”, or “performatives” for short (1962, 6). According to Austin, what distinguishes performatives is that “to *say* something is to *do* something” (1962, 12, emphasis in original), rather than to describe it. Much like Austin’s performatives, scientific models can also sometimes be used to *do* something; not only to represent and describe but to *change* their targets in important ways. Sociologists and philosophers of science have discussed this phenomenon, especially regarding models in economics and finance (Callon 1998; MacKenzie 2006a, b; Guala 2007; Mäki 2013; Boldyrev and Ushakov 2016), and, more recently, regarding epidemiological models (van Basseches et al. 2021; Winsberg and Harvard 2022). Borrowing from Austin, many have called models’ capacity to bring about or change their targets “performativity”.¹

This paper argues that Austin’s analysis of performatives can also inform the evaluation of models. According to Austin, while grammatically resembling constative utterances, performatives cannot be true or false like the latter, but can rather be “happy” or “unhappy”, that is, their user can achieve, or fail to achieve, whatever it is she used them for. We can thus

¹ Not all of the cited authors have. Uskali Mäki criticizes the use of “performativity” in the context of economic models as a misrepresentation of Austin’s notion because, unlike Austin who was interested in phenomena being *constituted* by performatives, MacKenzie et al. typically describe cases where an aspect of economics *causally influences* a target (2013). While I agree with Mäki that there are limits to the analogies between Austin’s performatives and models in science, I wish to defend here a specific analogy, namely that models may be *assessed*, similar to performatives, with respect to their felicity. And while I am agnostic as to whether the use of “performativity” is itself felicitous in this context, my usage targets what is now a considerable tradition that does so.

criticize a performative by asking whether its user has achieved her intended purpose, or has failed to do so due to various kinds of “infelicities” (1962, 14, 25). I wish to defend two claims: first, that performativity can sometimes be regarded as a model purpose and second, that models can be evaluated in such cases with respect to their felicity, that is, whether their use has achieved or failed to achieve this purpose.

My argument draws on the view that models should be assessed with respect to whether they are adequate for particular purposes, which may be epistemic or practical (Parker 2020) (section 2). Can performativity sometimes be regarded as a practical model purpose? By zooming in on a subdiscipline of economics, mechanism design, I shall argue that an important purpose of models can be to change their targets in desirable ways, and accordingly, that performativity can be seen as a benchmark for evaluating the adequacy of the model (i.e. whether it allows for a felicitous use) in these cases (section 3). I then consider two objections to this Austinian kind of model evaluation. The first is that performativity is a spurious model purpose that can be reduced to epistemic purposes. However, it turns out that the appearance of reducibility is deceiving (section 4). The second objection is that performativity, even when treated as a model purpose by scientists, is not a legitimate model purpose due to ethical concerns. However, it will be argued that the concern here is not performativity as a model purpose but rather the means of achieving it, which may (as with other purposes) be legitimate or not (section 5).

2. The Adequacy-for-Purpose View, but What Are Adequate Purposes? According to an “adequacy-for-purpose” view, to which many philosophers and scientists subscribe, “model quality is to be assessed relative to a purpose; model evaluation seeks to learn whether a model is adequate or fit for particular purposes” (Parker 2020, 458). According to Wendy Parker, what it means for a model to be adequate for a given purpose is that its use results in the achievement of the purpose in a particular case, or reliably brings about the purpose

across various cases. She contends that model purposes can be epistemic or practical; but while she provides a catalogue of epistemic purposes, such as explanation, prediction, and teaching, she leaves practical purposes largely unspecified.²

It should be clear that, if performativity can be a model purpose at all, we should find it among the practical purposes. This is because, similar to Austin's notion that we ought to evaluate a performative by asking whether it is felicitous, rather than true, we would evaluate the model's adequacy by asking whether its use has changed its target in the intended way, rather than accurately explained or predicted. This is not to say that models that are performative in some ways cannot also, or even exclusively, have epistemic purposes. For instance, if a model is used to predict a bank failure, which in turn leads to a bank run and thus contributes to the bank actually failing, the model has certainly performed its target, but it has also provided a true prediction. Here, performativity might not be regarded as a purpose of the model at all (prediction might be) but might rather be seen as an adverse side effect, even if one believes that performativity can sometimes be a legitimate model purpose.

So might performativity be considered a practical model purpose, and if so, under what conditions? Even though the adequacy-for-purpose literature hasn't dealt with this question to date, it may provide some general guidance as to how we might determine what should count as a proper model purpose. Importantly, and unlike some previous accounts that exclusively regard models as representational tools, the adequacy-for-purpose view is informed by and aims to account for scientific practice. Thus, some practices that would seem puzzling when

²Parker also notes that, even when stated model purposes are practical, the underlying contribution of the model is often epistemic, such as when the practical goal is flood control, but the epistemic role of the model is to project the severity of flooding for different possible policies (2020, 460). I come back to this point below.

regarding models as aimed at representation only, for instance that modelers sometimes intentionally *misrepresent* a target, can easily be explained from the perspective of adequacy-for-purpose: in such cases, the modelers aim to achieve other epistemic or practical purposes with respect to which the adequacy of the model should then be evaluated (Parker 2020, 459). If the modelers use a model to predict, then prediction should count as one of, or even the main purpose, at hand; but they could sacrifice predictive accuracy if they set themselves different purposes. This suggests that what is to be regarded as a model purpose is typically the decision of the scientists who do the modelling.

That scientists should determine what is to be regarded as model purposes also squares with the fact that they are usually the ones who know the capacities and limitations of their models best, and who can thus best gauge for what purposes their models have a realistic chance of being adequate. While model purposes may be motivated by policy makers' needs, the latter will often lack this knowledge. For instance, Parker (2009) argues that climate models are adequate for coarse-grained predictions, but are not adequate for more demanding predictive aims that many policy makers would hope could be achieved. This does not rule out the policy makers' aims as potential model purposes, as something may well be regarded as a purpose even if it cannot currently be achieved. But it does bolster the view that, to learn about what should count as a model purpose, we should primarily consult scientific practice. This doesn't mean that we should uncritically accept any purported purpose that some scientists have set themselves. But if standard scientific practice is to treat something as a model purpose, this provides a rationale for regarding it as such, unless there are independent reasons to question its legitimacy. So, to make progress with respect to performativity as a potential model purpose, we should consult a concrete scientific practice. We will turn to this now.

3. Mechanism Design, and Performativity as a Model Purpose. While a weather forecast does not change the weather, economic forecasts may influence people's behavior and can thereby change economic outcomes. For Oskar Morgenstern, the famous economist and founder (along with John von Neumann) of game theory, this seemed to pose a serious threat to the possibility of predictions in economics (Leonard 2010, 101). Perhaps ironically, game theory subsequently became instrumental in turning this vice into virtue, by enabling the design of social institutions (cf. Guala 2007).

An important step in this transformation was to model markets and other institutions as “mechanisms”, that is, rules that determine social outcomes as a function of individual actions, where individuals are assumed to be rational maximizers that will anticipate outcomes and choose an optimal action, given their preferences and beliefs (Hurwicz 1973; Wilson 2021). Using game theory, the social outcomes that different mechanisms bring about (assuming the behavioral expectations are correct) could then in principle be calculated as equilibria (“in principle”, because this would require knowledge about private information, as explained below). Because individuals are explicitly modelled as anticipating outcomes and adapting to a given mechanism, the effects of a mechanism on their behavior are “endogenized” in the model, which, if successful, helps overcome Morgenstern's concern with spontaneous behavioral changes. Accurate prediction would, however, require a detailed knowledge, both of the precise institutional rules that are in play and of individuals' beliefs and preferences, which are their private information. If there is uncertainty about the functioning of the institution or about people's private information, the predictive task might fail.

Crucially, in mechanism design, these informational burdens are overcome with a methodological shift: rather than designing a model to explain or predict some target, the goal is to *make the target resemble the model* in relevant aspects (cf. Guala 2005, 164). This kills

two birds – uncertainty concerning institutional rules and concerning private information – with one stone. Uncertainty concerning the institutional rules is diminished by instituting rules in the target that are equivalent to the designed mechanism. Uncertainty concerning private information is diminished by devising “clever” mechanisms (more formally, these are “incentive-compatible” mechanisms), which make it a dominant strategy for agents to reveal their private information and to act according to the plan.

Thus, in mechanism design, the modelling purpose is to create or change institutions in a way that brings about social outcomes that are deemed desirable,³ by enticing information sharing and aligning individual incentives with social goals. Whereas the focus of early mechanism design was predominantly theoretical, roughly since the 1990s, the focus has shifted to applied market design, where economists’ role has been, in Alvin Roth’s words, akin to that of an “engineer” (2002). Economic engineers (including Roth himself) have since been engaged in tasks such as designing spectrum auctions to increase efficiency and raise revenue for the government, assigning seats in schools and universities, making labor markets fairer and preventing them from unravelling, and matching organ donors with recipients to increase transplantation rates and save lives.⁴ Because the purpose of such design processes is to create or change a target in desirable ways, the designed mechanism can be evaluated with respect to its “felicity” – according to whether it is adequate for this purpose. Thus, reflecting on this practice, Roth writes, “the real test of our success will be not merely how well we understand the general principles that govern economic interactions, but how well we can

³ The question of how we might determine what is to be regarded as desirable in a democratic society will be considered below.

⁴ See Kominers et al. (2017) for an overview.

bring this knowledge to bear on practical questions of microeconomic engineering” (1991, 113).⁵

Recall that, according to Parker, what it means for a model to be adequate for a purpose is that its use results in the achievement of the purpose in a particular case (we might call this the “token reading”), or that it brings about the purpose across various cases (“type reading”). Which meaning of adequacy-for-purpose is apt in the context of mechanism and market design? Arguably, the token reading captures model evaluation, especially in applied market design, better. For instance, Ken Binmore and Paul Klemperer claim, concerning auction design, that this “is a matter of ‘horses for courses’, *not* one size fits all; each economic environment requires an auction design that is tailored to its special circumstances” (2002, C94, emphasis in original). Their argument is that some European countries had copied the successful 3G spectrum auctions in the UK without having paid sufficient attention to local differences, which led to failure of these auctions. So, since every design effort is unique and the details of particular institutional arrangements matter, the felicity of a mechanism will mainly be assessed by asking whether its use has achieved the desired result in a particular case at hand.⁶

To sum up the argument so far: mechanism and market design form a specific scientific practice where performativity is explicitly regarded as a model purpose. Thus, if scientific

⁵ While Roth doesn’t use the term “performativity” in this passage, in a recent paper with Michel Callon (one of the leading performativity theorists), there seems to be agreement that the aim of market design is a kind of performativity (Callon and Roth 2021).

⁶ This doesn’t rule out that some general features, such as stability in matching markets, are important for the well-functioning of mechanisms across many cases (e.g. Roth 2002).

practice should guide what is to be regarded as a model purpose, and if we evaluate models with respect to their adequacy-for-purpose, it follows that models can sometimes be evaluated with respect to their “felicity”; to whether their use has achieved or failed to achieve the intended performative effects. Let’s next discuss two kinds of worries with respect to this conclusion.

4. Is Performativity a Spurious Model Purpose? The first worry is that performativity might be dispensable as a model purpose because, even when the ultimate goal is to bring about a target or change it in certain ways, models are used for epistemic purposes within this process. To illustrate, suppose that a model is used to provide a number of conditional projections: if policy A is implemented, outcome x results; if policy B is implemented, outcome y results, and so on, thus enabling the policy maker to pick a policy that will likely result in a desirable outcome. It would seem that, while the model has aided the policy maker in bringing about this outcome, its function here is primarily epistemic, enabling informed judgements about the consequences of different possible policies (see Parker 2020, 461 for a similar case). We might thus evaluate the model with respect to whether it is adequate for this epistemic purpose, praising it if the chosen policy leads to its predicted outcome, and criticizing it if not. Thus, performativity would appear to be merely a spurious, or “reducible” model purpose; its felicity could be assessed by asking whether it is adequate for predictive or other epistemic purposes.

However, while a model may in such a case *also* have epistemic purposes, the appearance of reducibility to epistemic purposes is deceptive. It is useful here to again consult Austin’s analysis. Austin noted that certain relations hold between performatives and constatives; for instance, if the performative “I apologize” is felicitous, the statement that I am apologizing is true (1962, 53). There are similar relations between performative model uses and predictions that are based on those models. Let’s focus, for the moment, on model uses (such as in

mechanism design) that are intended to make a target resemble the model in relevant aspects.⁷ In these cases, certain logical relations will hold; in particular, if the use of a model is successful in performing its target, then a prediction based on that model will come out true. Note that the converse is not the case: a true prediction does not imply a felicitous model use, as the prediction could be true for other reasons. For example, the Black-Scholes-Merton model came to be used by stockbrokers to set option prices and to buy and sell options, which in turn generated price patterns resembling those predicted by the model (see MacKenzie 2006b). It was a serious possibility, however, that the model was simply very good at predicting price patterns in the market (i.e. patterns that would have materialized without the practical use of the model), and it took considerable efforts to work out that this was not so (Rubinstein 1985). So, if a purpose of the model was to align the practice of option pricing, the achievement of this purpose could not have been assessed by asking whether it was adequate for predicting market outcomes.⁸

Similarly for the policy case above: whether the use of the model has causally contributed to a specific outcome cannot be assessed by asking whether the model accurately predicted this,

⁷ In MacKenzie's terminology, this kind of performativity is labelled "Barnesian" (2006a, 17); below, we'll come to different kinds.

⁸ This is not to deny that the accuracy of a model's predictions can sometimes serve as evidence that its performative use was successful. For instance, it could be shown that there were large discrepancies between price patterns before the publication of the Black-Scholes-Merton model and what the model would have predicted then, and these discrepancies decreased rapidly thereafter, thus confirming that the model was successfully used to perform option pricing (see Rubinstein 1985). Of course, this is different to the claim that performative capacities could be reducible to prediction.

as the latter does not imply the truth of the former. For example, suppose the aim is to have a certain percentage of the population vaccinated against a communicable disease, and the modelers, using a rational choice model, advise the policy maker to provide financial incentives to reach this goal. Suppose that the policy maker heeds their advice, but even though the critical vaccine uptake is indeed reached, this wasn't due to the implemented policy; rather, the outcome would have come about anyway, say, because people care for the health of their loved ones. So, unbeknownst to the modelers, even though they provided a true prediction, their use of the model was unhappy – it did not change the observable outcome.

At this point, it might be objected that the modelers' prediction was true for the wrong reasons – that people's motivation to get vaccinated was altruistic rather than selfish (as assumed in the model). Perhaps, then, we should instead look at the explanation that the model provided, which was false. But suppose the explanation were true; for instance, suppose people really did get vaccinated because of the incentives provided, which the model correctly describes. However, the policy led to a crowding out of their altruistic incentives, and had the mechanism not been implemented, the resulting vaccine uptake would still have been identical. What the example shows is that true prediction doesn't imply felicitous model use, and neither does true explanation.⁹ The appearance of reducibility to these epistemic purposes is thus deceiving.

⁹ More generally, the relation between performativity and explanation is more complex than the one between performativity and prediction: happy performance doesn't imply a true explanation as a model could have performative effects without, for instance, identifying relevant causal factors required for successful explanation; and true explanation doesn't

Finally, not all performative model uses are of the kind that we have so far considered. In cases of “counterperformativity”, the use of a model influences an outcome in such a way that it resembles it *less* (Mackenzie 2006a). For instance, some dire, worst-case projections of the course of an epidemic might lead to individuals practicing increased social distancing, which may in turn prevent the worst case from materializing (van Basshuysen et al. 2021). In such cases, it is even less plausible that a model’s felicity could be reduced to epistemic purposes; that a model was designed in order to provide a false prediction, for instance, seems to be an absurd proposition. Counterperformativity, however, gives rise to a concern with respect to the legitimacy of performativity as a model purpose, to which we turn now.

5. Is Performativity Illegitimate? The second concern is that performativity *shouldn’t* be regarded as a model purpose, even when scientists in fact do. Thus, Eric Winsberg and Stephanie Harvard claim that performativity “is never a legitimate purpose for a model” as this “would be a serious threat to democratic decision making” (2022, 4). An important function of models within democratic decision making, according to them, is to allow estimating the costs and benefits associated with possible policies, thus enabling decision makers and the public to determine whether a policy’s benefits are worth its costs. Thus, “[i]t is only by accurately measuring costs and benefits and knowing what sacrifices we wish to make that we can determine whether performativity (be it defined as changes to scientific advising, policy decisions or individual behavior) would be desirable or not” (2022, 4). But if a model is used to influence outcomes, this would seem to make it impossible to conduct an impartial cost benefit analysis. Rather, scientists, by constructing the model in this way, would *presuppose* what the desirable outcome is that the use of the model is supposed to

imply happy performance as we might fail to bring about a desired outcome, even when we understand relevant causal factors at play.

bring about, which would go against the democratic procedure. Furthermore, they argue, since scientists themselves would rely on the biased model results, “not only are they deciding on behalf of everyone else without consulting their values, they are doing so blind” (2022, 4).

As Winsberg and Harvard note, performativity can operate through different “channels”. In its most direct form, a model, by becoming public knowledge, directly influences individual behavior. Examples include the Black-Scholes-Merton model which, as we saw above, came to guide stockbrokers in setting option prices, which in turn generated price patterns resembling those in the model. In other cases, the influence on individual behavior is more indirect, through scientific advising and/or policy making (cf. Guala 2007).¹⁰ Mechanism design is an example where performativity works indirectly, as the designed mechanism is imposed as an institutional rule.

As long as the modelers don’t aim to intentionally deceive policy makers and the public (more on this below), Winsberg and Harvard’s worry that performativity might impair the ability to accurately estimate costs and benefits is unfounded. To see this, let’s consider the indirect channel first. Here, model results are presented to a policy maker,¹¹ who, after evaluating costs and benefits, takes a decision on the policy. Implementing the chosen policy

¹⁰ Guala distinguishes “indirect”, or “spurious performativity”, where models are used for policy making, from “genuine performativity” where economic theory becomes a norm that directly shapes individual decisions, for instance, aligning them with those of rational economic agents.

¹¹ I use “policy maker” as a placeholder for a democratically legitimated authority (such as an elected government, or even a citizenry deciding through a direct democratic procedure).

will then change the outcome, if successful, by bringing about the model projections that are deemed desirable (“Barnesian performativity”) or avoiding projections deemed undesirable (“counterperformativity”). Importantly, however, at the point in time at which the decision is taken, performative effects through this channel do not impair the cost benefit analysis as the channel hasn’t yet been activated. Similarly for scientific advising: since any performative effects haven’t yet kicked in when the modeler designed the model and used it to advise the policy maker, her advice is not blurred by any performative effects.

In the direct channel – where, when a model is made public, people spontaneously adapt their behavior – the sequence of events might appear more complex. This is because individuals may continuously adapt how they conceive of the risks and benefits of different choices. For example, take a model that describes the risk of contracting a disease as a function of the percentage of the population that is vaccinated against that disease. Someone learning about this model may initially estimate the risk of contracting the disease as high enough to get vaccinated, but this may reverse once a critical percentage is reached at which this risk is outweighed by the risks and discomfort of being vaccinated. In this case, when an individual deliberates her options, performative effects may have already influenced outcomes and thus her evaluation of costs and benefits. It doesn’t follow from this, however, that her evaluation is inaccurate, as it may correctly reflect costs and benefits at the point in time at which she is taking the decision. Thus, even if the model was constructed and published with the purpose of reaching a high vaccine uptake, this does not impair people’s ability to accurately estimate costs and benefits – on the contrary, in this case, it will assist them in reaching accurate estimates.

This example can be used to illustrate when performativity is achieved in legitimate ways. For suppose a modeler, in order to reach an even higher vaccine uptake, were to exaggerate the risks of contracting the disease, thus skewing individuals’ estimates in favor of the

vaccination. Even though more social welfare might be gained in this way, the modeler would thereby impose her values on the public without being democratically legitimated, and therefore Winsberg and Harvard would be justified in rejecting this use of the model as illegitimate. It doesn't follow, however, that performativity "is never a legitimate purpose for a model" (2022, 4); rather, their concern should have been attributed to the possibility that a model, in order to achieve this purpose more effectively, is used insincerely or deceptively. But it shouldn't come as a surprise that legitimate purposes might sometimes be pursued in illegitimate ways (e.g. think of someone inventing data points for an extrapolation). Democracy puts limits on the means through which social outcomes may be reached (and may delimit the amount of social welfare that may be achieved), but it doesn't preclude performativity as a model purpose.

6. Conclusion. This paper has defended the view that performativity can in some cases be regarded as a model purpose and that in these cases, models can be evaluated with respect to their "felicity", that is, whether their use has achieved this purpose. I have argued that performativity as a model purpose cannot be reduced to epistemic purposes, and it is not in tension with democratic decision making; in fact, performative models may provide important input into democratic decision making when scientists construct them truthfully. If these arguments are felicitous, models will be widely evaluated in the Austinian way.

References

- Austin, John L. 1962. *How to Do Things with Words*. Oxford: Clarendon Press.
- Binmore, Ken, and Paul Klemperer. 2002. "The Biggest Auction Ever: The Sale of the British 3G Telecom Licences." *The Economic Journal* 112 (478): C74-C96.
- Boldyrev, Ivan, and Alexey Ushakov. 2016. "Adjusting the Model to Adjust the World: Constructive Mechanisms in Postwar General Equilibrium Theory." *Journal of Economic Methodology* 23 (1): 38-56.
- Callon, Michel. 1998. *The Laws of the Markets*. Malden, MA: Blackwell.
- Callon Michel, and Alvin E. Roth. 2021. "The design and performance of markets: a discussion." *AMS Review* <https://doi.org/10.1007/s13162-021-00216-w>.
- Guala, Francesco. 2005. *Methodology of Experimental Economics*. Cambridge: Cambridge University Press.
- Guala, Francesco. 2007. "How to Do Things with Experimental Economics." In *Do Economists Make Markets? On the Performativity of Economics*, ed. Donald MacKenzie, Fabian Muniesa, and Lucia Siu, 128-62. Princeton, NJ: Princeton University Press.
- Hurwicz, Leonid. 1973. "The Design of Mechanisms for Resource Allocation." *The American Economic Review* 63 (2): 1-30.
- Kominers Scott D., Alexander Teytelboym, and Vincent P. Crawford. 2017. "An invitation to market design." *Oxford Review of Economic Policy* 33 (4): 541-71.
- Leonard, Robert. 2010. *Von Neumann, Morgenstern, and the Creation of Game Theory*. Cambridge: Cambridge University Press.

MacKenzie, Donald. 2006a. *An engine, not a camera. How financial models shape markets*. Cambridge: MIT Press.

MacKenzie, Donald. 2006b. "Is economics performative? Option theory and the construction of derivatives markets." *Journal of the History of Economic Thought* 28 (1): 29-55.

Mäki, Uskali. 2013. "Performativity: Saving Austin from MacKenzie." In *EPSA11 Perspectives and Foundational Problems in Philosophy of Science*, ed. Vassilios Karakostas, and Dennis Dieks, 443-53. Dordrecht: Springer.

Parker, Wendy S. 2009. "Confirmation and Adequacy-for-Purpose in Climate Modelling." *Proceedings of the Aristotelian Society, Supplementary Volumes* 83:233-49.

Parker, Wendy S. 2020. "Model Evaluation: An Adequacy-for-Purpose View." *Philosophy of Science* 87:457-77.

Roth, Alvin E. 1991. "Game Theory as a Part of Empirical Economics." *The Economic Journal* 101 (404): 107-14.

Roth, Alvin E. 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica* 70 (4): 1341-78.

Rubinstein, Mark. 1985. "Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978." *Journal of Finance* 40:455-80.

van Basseches, Philippe, Lucie White, Donal Khosrowi, and Mathias Frisch. 2021. "Three Ways in Which Pandemic Models May Perform a Pandemic." *Erasmus Journal for Philosophy and Economics* 14 (1): 110-27.

Wilson, Robert B. 2021. "Strategic Analysis of Auctions." *Econometrica* 89 (2): 555-61.

Winsberg, Eric, and Stephanie Harvard. 2022. "Purposes and duties in scientific modelling."
Journal of Epidemiology and Community Health doi:10.1136/jech-2021-217666.