Concrete Models and Holistic Modelling[*]

Wei Fang[ϕ]

Department of Philosophy, University of Sydney

Abstract: This paper proposes a holistic approach to the model-world relationship, suggesting that the model-world relationship be viewed as an *overall structural fit* where one organized whole (the model) fits another organized whole (the target). This approach is largely motivated by the implausibility of Michael Weisberg's weighted feature-matching account of the model-world relationship, where a set-theoretic conception of the structures of models is assumed. To show the failure of Weisberg's account and the plausibility of my approach, a concrete model, i.e. the San Francisco Bay model, is discussed.

---

[*] Draft paper, please do not quote without permission.

[ϕ] Address: University of Sydney, NSW 2006, Australia. Email: wfan6702@uni.sydney.edu.au.

1. Introduction

One philosophical interest in the philosophy of modelling focuses on the problem of the model-world relationship, also known as the representation problem. Among many approaches to this problem, the similarity account has attracted much attention recently. Ronald Giere (1988, 1999a, 1999b, 2004, 2010), Peter Godfrey-Smith (2006) and Michael Weisberg (2012, 2013) have made the most substantial contributions.

The core of this account, first developed by Giere, is a view of the model-world relationship:

> The appropriate relationship, I suggest, is *similarity*. Hypotheses, then, claim a *similarity* between models and real systems. But since anything is similar to anything else in some respects and to some degree, claims of similarity are vacuous without at least an implicit specification of relevant *respects and degrees*. The general form of a theoretical hypothesis is thus: Such-and-such identifiable real system is similar to a designated model in indicated respects and degrees. (Giere 1988, 81; author's emphasis)

However, critics point out that this account is only schematic since it falls short of specifying the relevant *respects* and *degrees* (Suárez 2003). Moreover, Giere argues that a philosophical account of scientific representation should also take into consideration factors such as the *roles* played by scientists, and the *intentions* those scientists have when modelling (Giere 2004, 2010). Given these considerations, Weisberg develops a more sophisticated similarity account, called the *weighted feature-matching* account

(2012, 2013). The basic idea of his account comes from psychologist Amos Tversky's *contrast* account of similarity, which states that the similarity of objects *a* and *b* depends on the features they share and the features they do not. In light of this, Weisberg proposes his own account:

$$S\,(m, t) = \frac{\theta f(M_a \cap T_a) + \rho f(M_m \cap T_m)}{\theta f(M_a \cap T_a) + \rho f(M_m \cap T_m) + \alpha f(M_a - T_a) + \beta f(M_m - T_m) + \gamma f(T_a - M_a) + \delta f(T_m - M_m)} \qquad (1)$$

f(x) refers to the weighting function, $\alpha, \beta, \gamma, \delta, \theta,$ and ρ denote weighting terms (parameters), subscripts *a* and *m* stand for attributes and mechanisms,[1] and *M* denotes the model and *T* the target. ($M_a \cap T_a$) stands for attributes shared by the model and the target, ($M_a - T_a$) attributes that the model has while the target does not, and ($T_a - M_a$) attributes that the target has while the model does not. The same story goes for mechanisms *m*. Attributes and mechanisms as a whole are called *features* of the model and the target.

An interpretation for this equation is needed. First, there must be a feature set Δ, and the set of features of the model and the set of features of the target are defined as sets of features in Δ. The elements of Δ are determined by a combination of context, conceptualization of the target, and the theoretical goals of the scientist. Besides, the

---

[1] Properties and patterns of systems are termed attributes, and the underlying mechanisms generating these properties and patterns are termed mechanisms (Weisberg 2013, 145).

contents of Δ may change through time as science develops, which in turn might result in a reevaluation of the established model-world relationship (*Ibid.*, 149).

Second, consider the values of weighting parameters $\alpha, \beta, \gamma, \delta, \theta$, and ρ. On Weisberg's account, different kinds of modelling require different weighting parameters. For example, if what interests us is the *minimal modelling* which concerns merely the mechanism responsible for bringing about the phenomenon of interest, the goal of this modelling is written as:[2]

$$\frac{|M_m \cap T_m|}{|M_m \cap T_m| + |M_a - T_a| + |M_m - T_m|} \rightarrow 1 \tag{2}$$

Finally, consider the weighting function *f(x)*, telling us the relative importance of each feature in the set Δ. Weisberg says scientists in most cases have in their mind some subset of the features in Δ, which they regard as especially important. Hence some features are weighted more heavily, while others are equally weighted. Besides, the background theory determines which features in Δ should be weighted more heavily. If the background theory is not rich enough, deciding which should be weighted more heavily is partly an empirical problem.

Having presented an outline of Weisberg's account, I will now argue that this account fails to capture the relationship between concrete models and their targets. To illustrate this

---

[2] Weisberg also describes three other kinds of modelling requiring different weighting parameters: hyperaccurate, how-possibly and mechanistic modelling (2013, 150-52).

shortcoming (Sec. 3), I will first describe the San Francisco Bay model (Sec. 2). Sec. 4 will propose a holistic alternative to Weisberg's account, suggesting that the model-world relationship be viewed as an *overall structural fit* where one organized whole fits another organized whole. Sec. 5 will examine a case where the organization of the whole can be treated as simply another feature.

2. The San Francisco Bay Model

John Reber worried about the fragility of the water supply in the San Francisco Bay area in the 1950s. To solve this problem, he proposed an ambitious proposal, namely, to dam up the Bay. Carrying out this plan would not only supply San Francisco with unlimited drinking water but also entirely change the area's transportation, industrial, military and recreation landscape (Weisberg 2013, 1). However, his critics worried that Reber's plan would only achieve its aims at the cost of destroying commercial fisheries, rendering the South Bay a brackish cesspool, creating problems for the ports of Oakland, Stockton, and Sacramento, and so on (Jackson and Peterson 1977; Cf. Weisberg 2013, 1).

To settle this dispute, the Army Corps of Engineers was charged with investigating the overall influence of the Reber plan by building a massive hydraulic scale model of the Bay (Weisberg 2013, 1-2). Once the model was built, it was adjusted to accurately reproduce several measurements of the parameters such as tide, salinity, and velocities actually recorded in the Bay (for details see Army Corps of Engineers 1963). After the adjustment, it was time to verify the model:

> Agreement between model and prototype for the verification survey of 21-22 September 1956, and for other field surveys, was excellent. Tidal elevations, ranges and phases observed in the prototype were accurately reproduced in the model. Good reproduction of current velocities in the vertical, as well as in the cross section, was obtained at each of the 11 control stations in deep water and at 85 supplementary stations. The salinity verification tests for the verification survey demonstrated that for a fresh-water inflow into the Bay system […], fluctuation of salinity with tidal action at the control points in the model was in agreement with the prototype (Huggins and Schultz 1967, 11).

After the verification, modellers were in a good position to assess the Reber plan through the model built. The investigation showed that it would considerably reduce water-surface areas, reduce the velocities of currents in most of South San Francisco Bay, reduce the tidal discharge through the Golden Gate during the tidal cycle, and so forth (Huggins and Schultz 1973, 19). Given these disastrous consequences, the Army Corps then denounced Reber's plan (Weisberg 2013, 9).

3. How Could Weisberg's Account Shed Light on the Bay Model?

I have argued elsewhere that Weisberg's account cannot shed light on mathematical models due to its atomistic conception of features and its assumption of the set-theoretic approach to model structures (citation anonymized). I find that the same charges can be raised in the case of concrete models.

Consider the first charge: Weisberg’s account is committed to an atomic conception of features. The key of Weisberg’s account is the claim that the similarity of objects *a* and *b* depends on the features they share and the features they do not share. Let us take a closer look at the equation (1). The numerator invites us to weight features shared, and the denominator asks us to weight all features involved (including three feature subsets: features shared, features possessed by the model but not the target, and features possessed by the target but not the model). Each feature is weighted independently and only once, with it falling into one of the three feature subsets. The numerator is the weighted sum of features shared, the denominator is the weighted sum of features shared and unshared, and the similarity measure is the ratio of the numerator to the denominator.

However, features in the Bay model are not atomistic and independent of each other. As Huggins and Schultz put it explicitly, “Among the problems to be considered were the conservation of water […]; […] the tides, currents and salinity of the Bay as they affect other problems […]. None of these problems can be studied separately, for each affects the others” (1973, 12). The reason why none of these problems can be studied separately is because factors involved in these problems cannot be studied separately.

Consider, for instance, the relationship between two key features in the model: tide and salinity. Salinity levels vary along an estuary depending on the mixing of freshwater and saltwater at a site. An estuary “is the transition between a river and a sea. There are two main drivers: the river that discharges fresh water into the estuary and the sea that fills the estuary with salty water, on the rhythm of the tide” (Savenije 2005, Preface ix).

To illustrate this “rhythm of the tide”, consider the effect of the spring-neap tidal cycle on the vertical salinity structure of the James, York and Rappahannock Rivers, Virginia, U.S.A.:

> Analysis of salinity data from the lower York and Rappahannock Rivers (Virginia, U.S.A.) for 1974 revealed that both of these estuaries oscillated between conditions of considerable vertical salinity stratification and homogeneity on a cycle that was closely correlated with the spring-neap tidal cycle, i.e. homogeneity was most highly developed about 4 days after sufficiently high spring tides while stratification was most highly developed during the intervening period. (Haas 1977, 485)

This short report shows not only that characteristics of salinity (such as stratification and homogeneity) are influenced by characteristics of the tide, but also that there is a phase connection (or synchronization) between tidal cycle and salinity oscillations. The former is a causal relationship while the latter is a temporal relationship. The phase connection among features was also emphasized by the Army Corps when verifying the Bay model, saying “These gages were installed in the prototype and placed in operation several months in advance of the date selected to collect the primary tidal current and salinity data required for model verification, since *it was essential to obtain all data simultaneously for a given tide over at least one complete tidal cycle of 24.8 hours*” (1963, 50; my emphasis). Moreover, the same story goes for tide and tidal currents (for details see Army Corps 1963, 20).

In short, features in a model bear not only causal relationships, but also temporal relationships to one another. This implies that, when verifying the model, features of the model causally interact with each other in producing certain outputs (e.g. predictions, effects, phenomena, etc.), rather than that they individually or separately produce outputs. So although outputs of key features in the Bay model can be identified and measured separately, they are not produced separately.

It is important to note that the causal interaction among features may lead to a different kind of interaction, i.e. a "similarity interaction",[3] wherein features interact with one another in producing the similarity value. That is, one feature's contribution to the similarity value depends on other feature(s)' contribution to that value.[4] The difference between causal and similarity interaction is that the latter is a statistical relationship among measured features, and can be viewed as a reflection of the former when coupled with an assumption that there might be such an underlying causal structure.[5] For example, a similarity interaction is shown by the verification of salinity in the Bay model, where the measurement of salinity (as a measurement of one feature's contribution to the similarity

---

[3] I thank X for suggesting this term for me.

[4] This point can be best illustrated with the curve fitting example: when computing the fit of a straight line $y=ax+b$ to a cloud of points, a and b will depend on each other to produce the best fit (I thank X for giving me this example).

[5] This assumption is important because there are cases where the fact that there is similarity interaction cannot guarantee that there is also causal interaction, because some randomly generated data set may also show interaction among features. In other words, causal interaction can lead to similarity interaction and the reverse is not true (I thank Y for letting me know this). I will discuss this assumption, called "precondition" later, in Sec. 4.

value from Weisberg's perspective) depended on other features in the way in which other features were kept constant: "salinity phenomena in the model were in agreement with those of the prototype *for similar conditions of tide, ocean salinity, and fresh-water inflow*" (*Ibid*., 54; my emphasis).

The way that similarity interaction reflects causal interaction, when coupled with the assumption mentioned above, can be expressed as follows: if what is under verification is a causal structure to which modellers do not have direct access (so the structure cannot be a feature in Weisberg's formula), then the coherent behavior of features (i.e. their similarity interactions such as phase connections) is a way of verifying, or at least indicating, the causal interactions in the underlying causal structure.[6] That is the reason why it was so essential to obtain all data simultaneously within a complete tidal cycle for the Bay model, and why all other features must be kept constant when verifying salinity (or other features).

Given features' causal interactions in the model and their similarity interactions when measuring them, it seems that assessing the relationship between a model and its target cannot be simply achieved in the way suggested by Weisberg's equation, for features' contribution to the similarity relationship is not *additive* but *interactive*. That is, to assess the relationship between a model and its target, one cannot measure each feature's contribution independently and then add them together.

4. Set-Theoretic or Non-Set-Theoretic? A Holistic Alternative

---

[6] I thank X for bringing this point to my attention.

Now we arrive at the problem of why Weisberg's account is deeply committed to an atomistic conception of features. As I have argued elsewhere, this problem ultimately comes down to Weisberg's understanding of the structure of models (citation anonymized). Weisberg says models are *interpreted structures* (2013, 15), so concrete models are interpreted concrete structures. At first glance, I have no quarrel with this understanding. On closer inspection, however, it can be shown that Weisberg's account on the model-world relationship assumes a set-theoretic approach to the structure of models.[7] This is because Weisberg's similarity measure can be derived from the *Jaccard similarity coefficient* between two sets, a coefficient assuming a set-theoretic conception of objects (citation anonymized).

The key to the set-theoretic approach to structures is its assumption that elements of objects (i.e. models and targets) are independent of each other, just as elements of a set are independent of each other. In other words, it construes both the model and the target as a set of independent elements, the similarity between which consists in the ratio of the number of elements shared to the number of all elements (citation anonymized). However, as discussed in Sec. 3, features are not independent. More importantly, their causal interactions may result in a similarity interaction among features.

This similarity interaction undermines Weisberg's account, for it cannot properly capture the dependence relationship of features' contribution to the overall similarity

---

[7] Note that Weisberg *explicitly* objects to the set-theoretic approach to models (2013, 137-42). However, I think it is compatible to claim that someone *implicitly* assumes what someone explicitly rejects.

measure between a model and a target. Nonetheless, there is still a way to save the very intuitive notion of similarity, by abandoning the set-theoretic conception of structures. That is, if the structure of a model is viewed as an *organized whole* in which each component of the whole is interconnected to other component(s) (directly or indirectly) in such a way that they interact with one another in producing certain phenomena of interest (i.e. outputs). Under such an understanding, therefore, assessing the relationship between a model and its target cannot be simply achieved by assessing each individual feature's relationship and then adding them together. Nor can this be done by assessing each connection among two or more features and then adding them together, even if connections (causal or non-causal) are also interpreted as features. On the other hand, however, the notion of similarity can be minimally preserved by claiming that assessing the similarity or *fit* (I will use fit hereafter) between a model and a target amounts to assessing the *overall structural fit* between the model and its target.

Generally speaking, structural fit means the structure of the model fits the structure of the target *as an organized whole*. That said, nevertheless, it should be stressed that there is no univocal meaning for the term "structural fit" that could encompass all circumstances, nor can a single equation or formula capture all situations. This is largely due to the heterogeneity of modelling practice and its multifarious goals. On the other hand, however, instructive points can still be asserted. In what follows I will elaborate some basics regarding the conception of "structural fit".

Structural fit in mathematical modelling means different things than in concrete modelling. For example, in a very simple case of curve fitting where a straight line $y=ax+b$

is fitted to a cloud of points, features *a* and *b* will interact with each other to produce the best fit. That is, what fits the cloud of points is the overall structure, not the additive sum of each individual feature. As I have argued elsewhere, in more complicated mathematical modelling such as the *maximum likelihood estimation*, the fit is usually achieved through comparing the predicted data set derived from the model *as a whole* to the observed data set derived from the target system (citation anonymized). Individual features of the model simply disappear, and causally related features, as constituting a whole, that co-occur in the data set are what really matters.

In the case of concrete modelling, admittedly, the claim that assessing the fit between a model and a target amounts to assessing the overall structural fit seems to be less apparent. On closer examination, however, the same claim still holds. Let us go back to the verification of the Bay model. At first glance, it seems the verification of the model was achieved by independently verifying the output (i.e. data sets) of each individual feature, as the report showed (see Sec. 2 for the verification report). That is, it seems that by verifying that each feature in the model fits its counterpart in the target, scientists made the judgment that the model fits the target system.

Underlying this seemingly plausible reasoning, however, there remains the problem of why we are allowed to confirm the verification of the model by means of only verifying several outputs of individual features. Or, to put it slightly differently, in terms of what does the fit of features guarantee the judgment about the fit of the model to the target? I take it that it is more than the fit of individual features themselves that makes sense of the reasoning that the model fits the target. There must be a precondition for this reasoning (remember the “assumption” made in the last section). After all, there are many cases in which the fit of features does not guarantee the fit of the model itself to the target. For instance, a drawing of Tom’s face may accurately capture all features of his face, e.g., nose, eyes, mouth, etc., but still falls short of fitting his face, because of the wrong organization of these features, e.g., putting the mouth in between the eyes and nose (Weisberg would argue that the organization could be a feature. I will discuss this point in Sec. 5.).

So if the fit of features is insufficient to vindicate the fit of a model to its target, what could provide this vindication? My claim is, contrary to Weisberg, that it is the *overall structural fit* of the model to the target system that warrants the fit judgment about the model and its target. In other words, the fit of individual features can only succeed in supporting the fit of the model to the target by the precondition that these features can be organized into the whole (i.e. the assumption that there is such an underlying causal structure), not the other way around.

To understand this “holistic reasoning”, let me articulate the specifics involved step by step. We first build a concrete model, i.e. a concrete structure, wherein features are interconnected with one other in such a way that they have the potential to interactively produce certain phenomena of interest (i.e. outputs). Before verifying the model, we need to adjust key features to make sure the model works very well. Note that any adjustment will not simply be the adjustment of individual features but also of their interconnections, resulting in the adjustment of the overall structure of the model. Finally, we verify the model by comparing the outputs of the model to the outputs of the target. As with mathematical models, this verification is also usually made via comparing data sets, as shown in the Bay model. Note that though these outputs can be identified, derived and measured independently, it is causally connected features that interact in producing them. In other words, although you verify each feature separately, the support provided by a single feature is not confined to that feature of the model, but confirms all aspects of the model that are involved in generating that output.

Thus understood, therefore, the gist of verifying a concrete model such as the Bay model can be captured as follows. The verification of each feature, as a component of a whole, is simply the verification of one aspect of the structure. So the verification of different features is the verification of the same structure from different perspectives. Thus, if the model is an organized whole, then the more features that are independently verified the more likely it is that the model resembles the reality. On the other hand, if what is under verification is not an organized whole but an aggregation of independent items, then the verification of each lends no credence to other parts of the aggregated whole—because these items are not causally linked, the verification of each item is only the verification of that item itself.

In sum, the relationship between a concrete model and its target is a holistic matter wherein an organized whole fits (to a certain degree) or fails to fit another organized whole. Though it seems at first blush that the verification of the whole results from the sum of the verification of each component, the real picture is just the reverse: the whole is always in place and the component can gather force in supporting the verification of the whole only when it can be organized into the whole.

5. Organization and Features

As mentioned above, Weisberg would argue that the organization could be a feature, so a drawing of Tom’s face capturing accurately not only his nose, mouth, eyes but also their organization can be a good model of Tom’s face. A holistic account agrees that organization could be a feature, but disagrees with the way that organization is treated in Weisberg’s similarity measure. Intuitively, we may say that a drawing of one person’s face is a good model if it has the right features: such as a nose, a mouth, eyes, and the organization of all of these. So it seems that if you get each individual feature right, then you get the whole model right. That is, features *additively* contribute to the goodness of the model.

This intuitive way of understanding scientific modelling, however, obscures the fact that features may interact in producing the fit of a model, as shown in Sec. 4. To reiterate this point and to draw a connection to our current discussion, consider another ordinary example.[8] Suppose Anne’s face is an ideal one which scientists want to model. Anne has an ideal nose, which is straight, in contrast to a non-ideal nose, which might be bumped or concave. She also has an ideal nostril, which is round, in contrast to a non-ideal one, which might be triangular or square. Scientist A draws a face for Anne that has a round nostril and a concave nose, while scientist B draws a face that has a triangular nostril and a bumped nose. Drawing A has an ideal feature (the round nostril), but neither feature of drawing B is ideal. Now we ask which drawing better fits Anne’s face. It is likely that we

---

[8] I thank X for giving me this nice example.

will say that B is better because our contemporaries' taste tells us that there is no face so ugly as one with a round nostril and a concave nose, though a round nostril itself is ideal. Hence we see a case wherein the nostril and nose interact to produce the fit of a model to a target.

This discussion leads to a more general question: what are features? In Weisberg's account, a model can *more or less* fit a target, but features are either shared or not. Yet as Wendy Parker points out, "relevant similarities often seem to occur at the level of individual features, not just at the level of the model" (2015, 273). This is because features themselves can be objects such that they more or less fit each other.[9] Weisberg may argue that this problem can be fixed by the assumption that a feature can be redescribed as a set of sub-features, so the similarity between two features can be measured as the result of the similarity between their sub-features. However, I see this treatment as a non-starter, for the similarity between sub-features may also be a matter of degree such that it should be measured as the result of the similarity between their sub-sub-features, and between their sub-sub-sub-features, and so on.

On the other hand, a holistic account does not encounter this problem: if a feature is an object, then it can be viewed as an organized whole. So the relationship between a feature in a model and a feature in a target also consists in their structural fit. Take a minimal model for instance. Most minimal models primarily attempt to represent repeatable patterns of behavior largely insensitive to underlying microscopic details (Batterman 2002, 27). Suppose we are interested in the buckling behavior of struts, and write a

---

[9] I thank X for bringing this to my attention.

phenomenological formula, called Euler’s formula, to characterize it (see Batterman 2002 for details). It seems the pattern of behavior is the only feature involved in this case, i.e., a dependence relationship among several parameters. So assessing the fit between the model and the target comes down to assessing the fit between the feature in the model and the feature in the target. For this, a holistic account can easily come through: the relationship is an overall structural fit, wherein a dependence relationship as a feature fits another dependence relationship.

6. Conclusion

This paper has shown that the assumption of a set-theoretic approach to structures makes Weisberg’s account fail to shed light on the San Francisco Bay model. Alternatively, a holistic approach to models, viewing the model-world relationship as an overall structural fit, fares better not only in capturing the Bay model, but more generally in making sense of modelling practice.

References

Army Corps of Engineers. 1963. *Technical Report on Barriers: A Part of the Comprehensive Survey of San Francisco Bay and Tributaries, California*. Appendix H, Volume 1: Hydraulic Model Studies. San Francisco: Army Corps of Engineers.

Batterman, Robert. 2002. “Asymptotics and the Role of Minimal Models.” *British Journal for the Philosophy of Science* 53 (1): 21-38.

Giere, Ronald N. 1988. *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.

Giere, Ronald N. 1999a. *Science without Laws*. Chicago: University of Chicago Press.

Giere, Ronald N. 1999b. “Using Models to Represent Reality.” In *Model-Based Reasoning in Scientific Discovery*, ed. Lorenzo Magnani, Nancy J. Nersessian, and Paul Thagard, 41-57. Springer Science & Business Media.

Giere, Ronald N. 2004. “How Models Are Used to Represent Reality.” *Philosophy of Science* 71 (5): 742-752.

Giere, Ronald N. 2010. “An Agent-Based Conception of Models and Scientific Representation.” *Synthese* 172 (2): 269–281.

Godfrey-Smith, Peter. 2006. “The Strategy of Model-based Science.” *Biology and Philosophy* 21 (5): 725-740.

Haas, Leonard W. 1977. “The Effect of the Spring-Neap Tidal Cycle on the Vertical Salinity Structure of the James, York and Rappahannock Rivers, Virginia, U.S.A.” *Estuarine and Coastal Marine Science* 5:485-496.

Huggins, Eugene. M., and Edward A. Schultz. 1967. “San Francisco Bay in A Warehouse.” *Journal of the IEST* 10 (5): 9-16.

Huggins, Eugene M., and Edward A. Schultz. 1973. “The San Francisco Bay and the Delta Model.” *California Engineer* 51 (3): 11-23.

Jackson, W. Turrentine, and Alan M. Peterson. 1977. *The Sacramento-San Joaquin Delta: The Evolution and Implementation of Water Policy*. Davis: California Water Resource Center, University of California.

Parker, Wendy. 2015. “Getting (even more) serious about similarity.” *Biology and Philosophy* 30 (2): 267-276.

Savenije, Hubert H. G. 2005. “*Salinity and Tides in Alluvial Estuaries*.” Elsevier Science.

Suárez, Mauricio. 2003. “Scientific Representation: against Similarity and Isomorphism.” *International Studies in the Philosophy of Science* 17 (3): 225-244.

Weisberg, Michael. 2012. “Getting Serious about Similarity.” *Philosophy of Science* 79 (5): 785-794.

Weisberg, Michael. 2013. *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.