

The Predictive Turn in Neuroscience

Daniel A. Weiskopf

Abstract: Neuroscientists have in recent years turned to building models that aim to generate predictions rather than explanations. This “predictive turn” has swept across domains including law, marketing, and neuropsychiatry. Yet the norms of prediction remain undertheorized relative to those of explanation. I examine two styles of predictive modeling and show how they exemplify the normative dynamics at work in prediction. I propose an account of how predictive models, conceived of as technological devices for aiding decision-making, can come to be adequate for purposes that are defined by both their guiding research questions and their larger social context of application.

Contact information: Department of Philosophy, Georgia State University, PO Box 3994, Atlanta, GA 30302-3984. Email: dweiskopf@gsu.edu

1. The Return of Prediction

Scientific explanation has flourished as a research topic in the past several decades. Its ascendance, however, has come at the expense of attention to prediction (Douglas 2009). Within the philosophy of science, competing theories of explanation and extensive case studies of explanatory practices have proliferated. The neglect of prediction has not been universal; constructive empiricists in the tradition of van Fraassen have championed empirical adequacy as the main desideratum of scientific theorizing, and there have been a handful of other recent exceptions (e.g., Northcott 2017). The asymmetry is nevertheless real and somewhat hard to account for, particularly since many sciences, including celestial mechanics, hydrology, economics, and climatology, have historically been organized around generating and refining predictions (Oreskes 2000).

In recent years demand has increased sharply for the scientific community to produce actionable predictions that can inform policy choices and justify public spending on basic research (Hofman, Sharma, and Watts 2017). Within neuroscience in particular, there is a marked shift underway towards building predictive models. This predictive turn is attested not just in various theoretical manifestos (e.g., Bzdok and Yeo 2017) but also in a host of new experimental, applied, and clinical practices. This trend arguably traces back to the emergence of neurolaw, where promises that neuroscience could inform sentencing decisions were underwritten by studies using brain biomarkers to predict criminal recidivism (Aharoni et al. 2013). Interest in neuroprediction has since spread to practitioners in consumer marketing, psychiatry, and addiction intervention.

The predictive turn can be defined in terms of a commitment to the autonomy of prediction vis-à-vis other organizing scientific principles. Consider this concise formulation:

“Perhaps the biggest benefits of a prediction oriented [approach] within psychology are likely to be realized when psychologists start asking research questions that are naturally amenable to predictive analysis. Doing so requires setting aside, at least some of the time, deeply ingrained preoccupations with identifying the underlying causal mechanisms that are mostly likely to have given rise to some data” (Yarkoni and Westfall 2017, 18). Echoing this point, Bzdok and Ioannadis say that “[s]uch predictive approaches put less emphasis on mechanistic insight into the biological underpinnings of the coherent behavioral phenotype” (2019, 3). Key in this statement of aims is the separation of prediction from explanation: the goals, methods, and products of predictively oriented science are assumed to be distinct and valuable in their own right. Prediction is not pursued primarily for the sake of confirming or deepening explanations, or for elucidating causal structure, but rather for its own sake. This represents not simply a return to the tradition of seeing prediction as a complement to explanation, but a novel call for scientific practices that pursue these distinct goals in parallel.

The predictive turn can be seen as a part of a large-scale reorientation involving both social and technological aspects. Its main drivers include changes in the culture of statistical analysis—primarily a shift to multivariate rather than mass univariate analysis—and straightforwardly technological improvements, such as computational resources for handling large datasets. Innovations in the social structure of scientific fields also contribute, most notably the emergence of multinational consortiums to create and manage platforms for distributed sharing of neural datasets. Finally, as noted above, there is pervasive pressure from government and industry to generate usable scientific results, as well as the novelty factor of adopting tools that have attracted attention and funding in the machine learning community.

Here I analyze several examples of how these new styles of predictive modeling are deployed in neuroscience. Attending to these cases will help to illuminate the larger normative structure of prediction. Starting from the idea that predictive models are tools for guiding decisions and action, I sketch an account of predictive norms that treats them as a species of technology assessment. Key to this account is the interplay between high-level norm schemata that govern prediction and how these norms are situated within practical and investigative contexts. Such contexts both concretize norms and generate dynamic trade-offs among them. The upshot is that norms of prediction, like other forms of technology assessment, involve highly situated interactions among epistemic, practical, and ethical concerns.

2. Bringing Prediction into Focus

Before turning to cases, two core senses of prediction first need to be distinguished.¹ *Logical* prediction is the use of a theory or model to infer some as-yet unobserved state of the world, without restrictions as to whether this state lies in the past, present, or future. Logical predictions frequently center on counterfactuals: were the system in a certain state, or intervened on in a particular way, this would be the result. The sense captures the use of models to draw inferences about potentially observable states and measurements. It also covers what is meant in calling one variable in a model a predictor of another: their relationship is such that the value of one reliably allows inferences to the value of the other. The logical sense of prediction is at issue in the claim that theories and models are tested via the predictions that they make or distinguished from each other on the basis of their consequences.

¹ The terms used here follow Oreskes (2000), but similar distinctions are made elsewhere.

Temporal prediction, by contrast, is the use of a model to infer a future event, event probability, or event probability distribution. In temporal prediction, as Sarewitz and Pielke (1999) note, scientists “use suites of observational data and sophisticated numerical models in an effort to foretell the behavior or evolution of complex phenomena... These predictions of complex phenomena seek to ascribe time, place and characteristics to events” (p. 123). Other researchers agree in this usage: prediction is “the anticipation of a future outcome based on information available in the present” (Poldrack et al. 2018). “Forecasting” is the common term for the art of producing temporal predictions in a particular domain: hurricane paths, earthquakes, disease prognosis, elections, epidemics, sales revenue, advertising success, and so on. Temporal prediction is a species of logical prediction insofar as it involves deriving consequences from a model. But not all models that make logical predictions can make temporal predictions, nor can they necessarily do so with any kind of ease or utility.

A final point of clarification. Contemporary model assessment techniques increasingly rely on prediction, as opposed to measurements of how well models accommodate existing data. Strong fits to past data notoriously risk failing to generalize to new observations (Varoquaux and Poldrack 2019). Assessing a model in terms of its performance on out-of-sample prediction tasks involves procedures such as cross-validation: dividing the data set into two subsets, parameterizing the model on the training subset, and calculating its performance on the test subset. In this context, prediction once again means logical prediction, i.e., generalization to the unobserved. Insofar as temporal prediction is future-oriented it is of course also a form of out-of-sample generalization. However, not all models assessed by their out-of-sample performance generate temporal predictions. If the data used in out-of-sample prediction has already been collected, this practice technically constitutes retrodiction rather than temporal prediction.

The salient feature in the studies examined here is that predictions are the main intended product of the investigation rather than just a means for testing a model. An investigation's intended products can be pinned down by elucidating the research question that it responds to. Questions concerning *why* or *how* something happened or tends to happen are typically explanatory demands for a causal or etiological account. Questions concerning what *will* happen, by contrast, demand answers that link current observables to future conditions. Not only do different kinds of models answer each of these questions, but the norms governing what counts as a satisfactory answer differ in each case. These points can be brought out through some brief case studies.

3. Two Varieties of Predictive Modeling

The first set of cases draws on a cluster of emerging fields that take prediction of large-scale extraneural events as their targets. “Neuromarketing” involves the use of information extracted from brain activity to sharpen predictions concerning the success of a product advertisement, public relations campaign, or other large-scale social phenomenon such as the success of a government anti-smoking initiative or the rise of a pop single. “Consumer neuroscience” and “communication neuroscience,” by contrast, focus on the theoretical study of neural processes involved in consumer behavior divorced from applications in the advertising industry.

These studies exemplify what Berkman and Falk (2013) call a “brain-as-predictor approach”: to treat “neural measures (e.g., activation, structure, connectivity) as independent variables in models that predict longitudinal outcomes as dependent variables” (p. 45). The neural measures in question are typically chosen on the basis of existing theory about the

psychological function of particular features; however, it is also common to uncover predictors in a more theory-free way by looking for patterns in the data. Following the logic of temporal prediction in advertising and polling, inferences in these studies aim to forecast social or population-level events.

In an early study, Falk, Berkman, and Lieberman (2012) scanned 30 participants, all heavy smokers who intended to quit, while they passively watched a series of TV ads for a stop smoking campaign. fMRI measures were collected from a pre-specified ROI in ventromedial PFC hypothesized to play a predictive role. This activation was compared with control ROIs and with self-reported ratings of ad effectiveness. Ad effectiveness itself was measured by increase in the call volume to the quit helpline after each ad ran. Each potential predictor was assessed for how well it captured the relative ordering of the three different ad campaigns screened. Activation in the target ROI was the best predictor of this ranking; notably, it outperformed the traditional survey self-report measures.

Changes in neural activation in a handful of people, then, can predict the outcome of population-level events and interventions happening months later. This makes it tempting to see the main innovation of predictive models as confined solely to increases in accuracy over traditional methods such as surveys or focus groups. Further studies refine this picture, however. Genevsky, Yoon, and Knutson (2017) investigated neural markers that could forecast which online crowdfunded projects would be most successful. Thirty individuals underwent fMRI while viewing images and text from different crowdfunding campaigns. They were then surveyed for their attitudes about the campaigns and whether they personally would fund them. The main results of interest are, first, that activity in nucleus accumbens (NAcc) and medial PFC predicted individual funding choices; second, that NAcc but not mPFC predicted aggregate

crowdfunding performance; and third, that neural measures were better aggregate funding predictors than were survey responses. Neural predictors of aggregate choice can contribute to forecasting accuracy over and above behavioral ones without always accurately predicting individual choices well.

The second set of cases centers on the prediction of individual neural, cognitive, or behavioral outcomes, drawing on the tradition of searching for biomarkers of psychiatric or neurological conditions (Gabrieli, Ghosh, and Whitfield-Gabrieli 2015; Jollans and Whelan 2018; Woo et al. 2017). Existing research has focused on early diagnosis of conditions such as schizophrenia and Alzheimer’s dementia (AD), mood disorders, and alcoholism. Here the inferences point in the opposite direction, from models trained on large population-level samples to individual-level predictions.² The overall research framework of single-subject prediction is translational, driven by the need to find clinically deployable tests that will rapidly, cheaply, and accurately facilitate classifying and treating patients. Ideally, early diagnosis or monitoring can lead to better outcomes.

In a landmark study, Whelan and a multinational team (2014) looked at data from 692 adolescents for predictors of binge drinking from ages 14-16. These included life history, personality traits, and various neural measures. The model they built, at its best, successfully predicted 77% of future binge drinkers and 67% of non-binge drinkers. Many of the chosen behavioral and personality measures, such as sexual activity and low conscientiousness,

² Predictive models of this kind primarily generate group-level predictions, but they are often used as proxies to forecast individuals’ behavior. As Poldrack et al. (2018) note, this raises the “group-to-individual” inference problem; here, though, I will assess these models as they are often used in practice, namely as devices for specifically individual prediction.

predicted both future and current binge drinking well.³ Among the brain measures, the most robust predictors were a combination of structural and functional indices across regions including right middle and precentral gyri, bilateral superior frontal gyrus, and premotor cortex. Notably, good predictors of future binge drinking were not always the same as good classifiers of current binge drinking: e.g., parenchymal volume and grey-to-white matter ratio was a predictor but not a classifier, and vmPFC and left lateral PFC classified current binge drinkers but were poor predictors.

Studies aiming to make such individualized predictions continue to proliferate, and generalizations about their scope and limits can now be ventured. For instance, a recent meta-analysis of 116 studies predicting transition to AD from mild cognitive impairment (MCI) revealed accuracies of 74.5% for fMRI and 76.9% for PET-based methods (Grueso and Viejo-Sobera 2021). In terms of the kind of classifier used, support vector machines were 75.4% accurate and convolutional neural networks were 78.5% accurate. Prediction of progression to AD, however, was always a harder task than sorting of patients according to their current status as normal vs. AD. There are obvious reasons why these tools might be of interest within biomedicine, despite their present limitations.

These sets of models are in many ways complements of each other. They differ in the evidential base on which they are each built (small group samples vs. population level samples), in the targets they aim to predict (socioeconomic events vs. individual outcomes), and in the contexts of their plausible use (corporate or state policymaking vs. clinical guidance). These differences aside, they draw on similar analytic and computational resources to produce their

³ Smoking, however, is so strong that it swamps the remaining predictors and was therefore removed from later analyses.

predictions. As we will see, the practice of normative assessment requires taking all of these factors into account.

4. Completing the Predictive Turn

We can now revisit the question of what makes something a good predictive model. This first needs to be separated from the issue of whether predictive power is an added epistemic virtue of theories and models that already do some explanatory work; that is, whether an explanatory theory that generates novel predictions *thereby* warrants more of our confidence relative to one that just fits existing data. This latter topic has been at issue in debates over predictivism (Douglas and Magnus 2013), but the models surveyed here are not explanatory models that also generate predictions. Rather, they are specifically tailored to answer predictive questions. Models such as these do not become worse at their job simply because they don't also do explanatory work. Instead, their success or failure turns on how closely they mesh with the sorts of decision contexts in which the models will prospectively be applied.

To see this, begin with the fact that prediction questions are generally asked because their answers are meant to be action-guiding. The choice of predicted targets—drawn from domains of social, clinical, and commercial importance—suggests that this consideration is at the forefront of modelers' minds. While some might pursue prediction for reasons of sheer intellectual curiosity, knowing possible futures is notably a way of pre-empting or planning for them. The need for action and the goals of clinicians and policymakers combine to determine the specific ways in which models are assessed. If a model is regarded as providing input to decisions concerning action, how well or poorly it does that should factor into its assessment.

For this reason, predictive modeling provides an exemplary illustration of how a range of values shape representational and modeling practices in science. The question, in short, is what makes neuropredictions “adequate for purpose” in the sense of Wendy Parker (2009): “the model, when used in accordance with specified methodologies, will convey information about the target system that allows model users to infer correct answers to the target questions” (p. 236). This perspective on prediction takes models to be technologies: devices for helping to achieve particular ends. Like other technologies, their assessment turns not just on how well they perform the functions for which they are designed, but also on how well or poorly they work as components of larger sociotechnical networks.

Completing the predictive turn, then, will depend on achieving greater clarity concerning how the settings in which these questions are asked also contribute to generating norms. To begin, consider some candidate dimensions of normative assessment such as accuracy, actionability, generalizability, interpretability, and informational economy. These criteria function more like *norm schemata* than providers of specific guidance to modelers. A norm schema is an abstraction that points towards clusters of ideal properties that a model might have. Each element in these clusters represents a way of making the norm more concrete and specific in its guidance. “Be accurate,” “Be able to potentially guide actions,” and “Use the least amount of informational input possible” are exhortations that commend agreement only insofar as they are platitudes. What is needed to make them substantive is an idea of what it means to comply adequately with *that* norm in *this* situation. Importantly, the investigative context per se does not necessarily give a complete or comprehensive set of clues concerning the right way to flesh out these schemata. Making them concrete requires looking at the wider ends to which the predictive questions are directed, and to the environment in which the technology will be deployed.

Take a norm such as informational economy. Roughly, this enjoins modelers to make use of the least amount of informational input necessary to generate predictions. In the real world, information isn't free. As the expense of taking measurements increases, models generally become less desirable. Within neuromarketing, this manifests as an emphasis on striking a balance among accuracy, cost, and ease of deployment (Hakim and Levy 2019). Similarly, while information-hungry multimodal imaging methods are in some cases the most accurate for predicting future disease states, they require costly equipment that takes special skill to operate, as well as unfamiliar analytic techniques. It is possible, given certain assumptions, to estimate the economic cost and participant burden imposed by using a diagnostic predictive tool (Petroni et al. 2019). While minimizing both of these factors (relative to existing predictive methods) is most desirable, clinical deployment may, for example, prioritize minimizing the burden on patients. What it means to be economical with information depends as well on what the predictions will be used for: the costs of error (e.g., a misdiagnosis) may be sufficiently high that the added expense of gathering more inputs is overall warranted.

Actionability refers to the usefulness of a prediction in formulating and ranking various candidate plans of action. The need for action-relevant information in fields such as marketing is clear from the sheer amount of waste generated by failed products, estimates of which range from 70-95% of all new products, and advertising of unknown efficacy, which consumes tens of billions of dollars per year in the US alone (Spence 2019). Neuropredictions that one advertising campaign rather than another is likely to succeed may influence whether that campaign is funded (all else being equal). But the relationship of prediction to action is highly indirect. In particular, action does not always take the form of intervention on the target. A prediction that assigns to a patient a moderately high risk of developing autism cannot, at the moment, lead to any course of

treatment that will cure the disorder or prevent it from emerging. The same is unfortunately true for many other neurological conditions. Interventions in neuropsychiatry may instead come in the form of situational changes such as increased monitoring of symptoms and provision of support services. How actionable a model is turns on properties that extend outward to the overall decision-making context. As a consequence, a “good” predictive model that provides highly accurate information may still be unsuitable or inadequate for purpose if there is no way that it can usefully shape our choice of actions (Lane, Hunter, and Lawrie 2020).

Interpretability (also called transparency) is another widely discussed norm of prediction (Chirimuuta 2020). Interpretable models are, roughly, those whose operations can be articulated and understood by their human users. Interpretability can also be unpacked in more local ways: e.g., it might be required that models should be interpretable specifically in terms of plausible biological mechanisms. Predictors that use “unbiological” or neurally opaque features would fail by this criterion (Woo et al. 2017, 372). Interpretability, too, is a norm whose application depends on specific contexts of inquiry, and uninterpretable models can sometimes be adequate for purpose. For example, in predicting progression to Alzheimer’s dementia, less interpretable models developed by merging complicated multimodal data sets may be acceptable if no differences in treatment depend on why someone is classified one way rather than another. This can be the case when filtering patients to take part in drug trials (prior to random assignment to conditions), where maximizing the number who will convert to AD is highly desirable. This use would minimize costs to the drug developers while not negatively affecting participants.

By the same token, neuromarketers may have little desire to open the black box of an EEG-based predictor that bears at best an opportunistic relationship to background theory or

ground truths.⁴ Contrast this attitude with that of consumer neuroscientists, who have an interest in understanding the causal basis of the model's predictions (Scholz et al. 2017). Here, predictive models are useful because they provide potential new targets of explanation. In a synergistic setting like this where predictive and explanatory investigations work in tandem, biological interpretability matters far more. These cases indicate that both the degree and kind of interpretability at issue are parameters that need to be set by these different purposes, depending on what is taken to be sufficient to ground the attitude of trust towards a model.

Similar points could be made concerning the other candidate norms above. Collectively, they show that when answering research questions that require building predictive models, we need to take into account a dynamic set of normative criteria. These dynamics show up in two ways. The first is in *norm instantiation*: the descent from norm schemata to particular, concrete instances of norms. The second is in *norm interaction*: the specification of how those norms are ranked in terms of their relative priority to one another. In norm instantiation we move from abstract schemata to more particularized forms of guidance. A norm such as “Be actionable” gets transformed in this process into a specific set of prescriptions for how the model should be assessed as to its action-guiding potential relative to a context of use. What it means for a model to be actionable depends on where and how it is embedded into a sociotechnical nexus—the same model embedded elsewhere may not be positioned to feed into decision making in the same ways, and hence should not be assessed relative to the same criteria. Norm interaction then takes these precisified norms and determines how they are to be weighted, ordered, or traded off against each other. In one context, being transparent might be of paramount importance, while in

⁴ Devices such as this are now being marketed. See, e.g., the iMotions platform: <https://imotions.com/neuromarketing>

others it might be most important to achieve an optimal mix of informational economy and actionability, with transparency being a marginal benefit at best.

In the current literature, predictive models are routinely and rigorously assessed on the usual methodological criteria such as whether they employ correct cross-validation techniques or are trained on an appropriately representative (unbiased) dataset. But these internal conditions do not capture the full range of relevant assessments, since these models are also situated within action-guiding systems. Both of the processes described here (instantiation and interaction) provide potential channels for wider values to shape how models should be assessed. The values in question have to do with whether the predictor is adequate for purpose. Whether or not it is depends in turn on the policy and decision-making procedures within which the predictive model is embedded. This flows directly from the technological perspective on models. New or existing technologies are assessed by measuring them against a set of criteria established through broad consultation with experts and ordinary citizens. A model's adequacy, by the same token, is determined by how well it measures up to the concrete and ordered structure of norms established by its overall situation of use.

5. Conclusion

At the moment, predictive modeling in neuroscience is best conceived of as an exploratory field. Nevertheless, it has had no shortage of critics who have attacked it on practical grounds as well as ethical ones. The approach taken here has been to suspend the attitudes of neuroskepticism and neurohype with an eye not towards judging the field's prospects, but rather towards considering what insights these cases might offer about prediction more broadly. It is too early to say how fruitful the predictive turn will turn out to be and to what research questions

and problems it might best be applied. But the development of ever more sophisticated models provides an opportunity to pay renewed attention to prediction as a distinct scientific goal. This shift is laudable and overdue. Future work may shed greater light on prediction by bringing it, at last, out of explanation's shadow.

References

- Aharoni, Eyal, Gina M. Vincent, Carla L. Harenski, Vince D. Calhoun, Walter Sinnott-Armstrong, Michael S. Gazzaniga, and Kent A. Kiehl. 2013. "Neuroprediction of Future Rearrest." *Proceedings of the National Academy of Sciences* 110 (15): 6223–28. <https://doi.org/10.1073/pnas.1219302110>.
- Berkman, Elliot T., and Emily B. Falk. 2013. "Beyond Brain Mapping: Using Neural Measures to Predict Real-World Outcomes." *Current Directions in Psychological Science* 22 (1): 45–50. <https://doi.org/10.1177/0963721412469394>.
- Bzdok, Danilo, and John P.A. Ioannidis. 2019. "Exploration, Inference, and Prediction in Neuroscience and Biomedicine." *Trends in Neurosciences* 42 (4): 251–62. <https://doi.org/10.1016/j.tins.2019.02.001>.
- Bzdok, Danilo, and B.T. Thomas Yeo. 2017. "Inference in the Age of Big Data: Future Perspectives on Neuroscience." *NeuroImage* 155 (July): 549–64. <https://doi.org/10.1016/j.neuroimage.2017.04.061>.
- Chirimuuta, Mazviita. 2020. "Prediction versus Understanding in Computationally Enhanced Neuroscience." *Synthese*, May. <https://doi.org/10.1007/s11229-020-02713-0>.
- Douglas, Heather. 2009. "Reintroducing Prediction to Explanation." *Philosophy of Science* 76 (4): 444–63. <https://doi.org/10.1086/648111>.
- Douglas, Heather, and P.D. Magnus. 2013. "State of the Field: Why Novel Prediction Matters." *Studies in History and Philosophy of Science Part A* 44 (4): 580–89. <https://doi.org/10.1016/j.shpsa.2013.04.001>.

- Falk, Emily B., Elliot T. Berkman, and Matthew D. Lieberman. 2012. "From Neural Responses to Population Behavior: Neural Focus Group Predicts Population-Level Media Effects." *Psychological Science* 23 (5): 439–45. <https://doi.org/10.1177/0956797611434964>.
- Gabrieli, John D. E., Satrajit S. Ghosh, and Susan Whitfield-Gabrieli. 2015. "Prediction as a Humanitarian and Pragmatic Contribution from Human Cognitive Neuroscience." *Neuron* 85 (1): 11–26. <https://doi.org/10.1016/j.neuron.2014.10.047>.
- Genevsky, Alexander, Carolyn Yoon, and Brian Knutson. 2017. "When Brain Beats Behavior: Neuroforecasting Crowdfunding Outcomes." *The Journal of Neuroscience* 37 (36): 8625–34. <https://doi.org/10.1523/JNEUROSCI.1633-16.2017>.
- Gruoso, Sergio, and Raquel Viejo-Sobera. 2021. "Machine Learning Methods for Predicting Progression from Mild Cognitive Impairment to Alzheimer's Disease Dementia: A Systematic Review." *Alzheimer's Research & Therapy* 13 (1): 162. <https://doi.org/10.1186/s13195-021-00900-w>.
- Hakim, Adam, and Dino J. Levy. 2019. "A Gateway to Consumers' Minds: Achievements, Caveats, and Prospects of Electroencephalography-based Prediction in Neuromarketing." *WIREs Cognitive Science* 10 (2). <https://doi.org/10.1002/wcs.1485>.
- Hofman, Jake M., Amit Sharma, and Duncan J. Watts. 2017. "Prediction and Explanation in Social Systems." *Science* 355 (6324): 486–88. <https://doi.org/10.1126/science.aal3856>.
- Jollans, Lee, and Robert Whelan. 2018. "Neuromarkers for Mental Disorders: Harnessing Population Neuroscience." *Frontiers in Psychiatry* 9 (June): 242. <https://doi.org/10.3389/fpsy.2018.00242>.

- Lane, Natalie M., Stuart A. Hunter, and Stephen M. Lawrie. 2020. "The Benefit of Foresight? An Ethical Evaluation of Predictive Testing for Psychosis in Clinical Practice." *NeuroImage: Clinical* 26: 102228. <https://doi.org/10.1016/j.nicl.2020.102228>.
- Northcott, Robert. 2017. "When Are Purely Predictive Models Best?" *Disputatio* 9 (47): 631–56. <https://doi.org/10.1515/disp-2017-0021>.
- Oreskes, Naomi. 2000. "Why Predict? Historical Perspectives on Prediction in Earth Science." In *Prediction: Science, Decision Making, and the Future of Nature*, edited by Daniel Sarewitz, Roger Pielke Jr, and Radford Byerly Jr, 23–40. Washington, DC: Island Press.
- Parker, Wendy S. 2009. "Confirmation and Adequacy-for-Purpose in Climate Modelling." *Aristotelian Society Supplementary Volume* 83 (1): 233–49. <https://doi.org/10.1111/j.1467-8349.2009.00180.x>.
- Petrone, Paula M., Adrià Casamitjana, Carles Falcon, Miquel Artigues, Grégory Operto, Raffaele Cacciaglia, José Luis Molinuevo, Verónica Vilaplana, Juan Domingo Gispert, and the Alzheimer's Disease Neuroimaging Initiative. 2019. "Prediction of Amyloid Pathology in Cognitively Unimpaired Individuals Using Voxel-Wise Analysis of Longitudinal Structural Brain MRI." *Alzheimer's Research & Therapy* 11 (1): 72. <https://doi.org/10.1186/s13195-019-0526-8>.
- Poldrack, Russell A., John Monahan, Peter B. Imrey, Valerie Reyna, Marcus E. Raichle, David Faigman, and Joshua W. Buckholz. 2018. "Predicting Violent Behavior: What Can Neuroscience Add?" *Trends in Cognitive Sciences* 22 (2): 111–23. <https://doi.org/10.1016/j.tics.2017.11.003>.
- Sarewitz, Daniel, and Roger Pielke. 1999. "Prediction in Science and Policy." *Technology in Society* 21 (2): 121–33. [https://doi.org/10.1016/S0160-791X\(99\)00002-0](https://doi.org/10.1016/S0160-791X(99)00002-0).

- Scholz, Christin, Elisa C. Baek, Matthew Brook O'Donnell, Hyun Suk Kim, Joseph N. Cappella, and Emily B. Falk. 2017. "A Neural Model of Valuation and Information Virality." *Proceedings of the National Academy of Sciences* 114 (11): 2881–86. <https://doi.org/10.1073/pnas.1615259114>.
- Spence, Charles. 2019. "Neuroscience-Inspired Design: From Academic Neuromarketing to Commercially Relevant Research." *Organizational Research Methods* 22 (1): 275–98. <https://doi.org/10.1177/1094428116672003>.
- Varoquaux, Gaël, and Russell A. Poldrack. 2019. "Predictive Models Avoid Excessive Reductionism in Cognitive Neuroimaging." *Current Opinion in Neurobiology* 55 (April): 1–6. <https://doi.org/10.1016/j.conb.2018.11.002>.
- Whelan, Robert, Richard Watts, Catherine A. Orr, Robert R. Althoff, the IMAGEN Consortium, Eric Artiges, Tobias Banaschewski, et al. 2014. "Neuropsychosocial Profiles of Current and Future Adolescent Alcohol Misusers." *Nature* 512 (7513): 185–89. <https://doi.org/10.1038/nature13402>.
- Woo, Choong-Wan, Luke J. Chang, Martin A. Lindquist, and Tor D. Wager. 2017. "Building Better Biomarkers: Brain Models in Translational Neuroimaging." *Nature Neuroscience* 20 (3): 365–77. <https://doi.org/10.1038/nn.4478>.
- Yarkoni, Tal, and Jacob Westfall. 2017. "Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning." *Perspectives on Psychological Science* 12 (6): 1100–1122. <https://doi.org/10.1177/1745691617693393>.