

Causal Emergence and Real Patterns

Joe Dewhurst

April 3, 2020

In several recent publications, Erik Hoel and colleagues have proposed a new model of causal emergence based on an information theoretic measure of causation. In this paper I will first introduce their measure, which they call ‘effective information’, and describe how they use it to argue for causal emergence. In brief, the idea is that certain kinds of complex system are structured such that an intervention characterised at the macro-level will be more informative than one characterised at the micro-level, and that this constitutes a form of causal emergence. Having introduced Hoel’s proposal, I will then assess the extent to which it is genuinely ‘causal’ and/or ‘emergent’, and argue that its interventionist approach to causation supports only an epistemic form of emergence. Finally I will suggest that the best way to make sense of Hoel’s proposal is in terms of Ladyman & Ross’ information theoretic gloss on Dennettian ‘real patterns’, which can clarify the sense in which emergence can be both causal and epistemic.

1 Effective information and causal emergence

The notion of ‘effective information’ was first introduced by Tononi & Sporns (2003) in the context of integrated information theory (IIT), a formal approach to the study of consciousness developed by Tononi.¹ In this paper they define effective information (EI) as “a quantity capturing all *causal* interactions that *can* occur between two parts of a system” (Tononi & Sporns 2003, emphasis in original). It is formalised as the mutual information (MI) between two complementary partitions (A and B) of a system (S), where A has been intervened on such that its outputs are maximally entropic (using an independent noise source), to give a measure, $EI(A \rightarrow B)$, of the effect of this intervention on B. They explicitly state that this measure should be interpreted causally:

¹For the latest version of this theory, see Oizumi, Albantakis, & Tononi (2014). I will not say anything more about consciousness in this paper, but it is important to keep this theoretical background in mind, as it plausibly provides some of the motivation for wanting a robust account of causal emergence (i.e., if consciousness is going to be understood as a non-epiphenomenal emergent phenomenon). Furthermore, it has been suggested that IIT should be understood as supporting a hylomorphic (structure-first) ontology (see Owen 2019), in which case many of the usual arguments against causal emergence might be misdirected. For the purposes of this paper, however, I will assess Hoel’s proposal on its own terms, without discussing IIT or hylomorphism directly.

Since A is substituted by independent noise sources, the entropy that B shares with A is due to *causal effects* of A on B [...] Thus, if the connections between A and B are strong and specialized, different outputs from A will produce different firing patterns in B, and EI(A \rightarrow B) will be high. On the other hand, if the connections between A and B are such that different outputs from A produce scant effects, or if the effect is always the same, then EI(A \rightarrow B) will be low or zero. (Tononi & Sporns 2003, emphasis added)

EI measures the effect that an intervention on one part of a system has on another part of that system, and is thus intended as an information theoretic measure of the degree to which two parts of a system are causally integrated.² It is similar in this sense to the information theoretic analyses of causation given by Pearl (1988, 2000) and Spirtes, Glymour, & Scheines (1993), which provided some of the inspiration for Woodward’s (2003) interventionist account of causation – a connection that Hoel draws explicitly, and that I will return to in the next section. For now I will move on to Hoel’s application of EI to the question of causal emergence.

Hoel (2017; see also Hoel *et al* 2013 and Albantakis *et al* 2019) demonstrates how adjusting the granularity of our description of a system can vary the quantity of EI that holds between one state of the system and another, such that there is potentially greater EI (and thus ‘more’ causation) when the system is analysed at the macro-level. He first defines a variant formalisation of EI, where a uniform intervention is applied across all states of the system, and then the results of this intervention are observed.³ The more certain the effects of an intervention, the greater the EI of the system. This can be most clearly demonstrated in terms of transition probability matrices (TPMs), which show the probabilities of one state transitioning to some other state. Consider three different 4-state TPMs (adapted from Hoel 2017):

$$\begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix} \text{ TPM 1}$$

$$\begin{bmatrix} 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \text{ TPM 2}$$

²A very similar account of causation is given by Griffiths *et al* (2015), who define an information theoretic measure of ‘causal specificity’ that is intended to adjudicate between different possible sources of phenotypic variation in biology (e.g. genetic vs. environmental).

³Hoel *et al* (2016) present another measure of causal emergence, ϕ^{Max} , which draws more on the integrated information theory mentioned above. It imposes stronger requirements than EI, but continues to rely on coarse-graining, which means that much of what I have to say here will still apply. In future work I would like to explore the differences between EI and ϕ^{Max} in greater detail.

$$\begin{bmatrix} 1/4 & 1/4 & 1/4 & 1/4 \\ 1/4 & 1/4 & 1/4 & 1/4 \\ 1/4 & 1/4 & 1/4 & 1/4 \\ 1/4 & 1/4 & 1/4 & 1/4 \end{bmatrix} \text{TPM 3}$$

TPM 1 has high (in fact, maximal) EI: if we know which initial state it is in, then we also know, with certainty, which state it will transition to (here EI = 1). The EI of TPM 2 and TPM 3 is lower: for TPM 2 only states 3 and 4 give us certain predictions, while states 1 and 2 are less predictive (the EI of the whole system = 0.5), whereas TPM 3 is entirely random (EI = 0). Taking EI as a measure of causation, we can say that the states of TPM 1 are more causally integrated than those of either TPM 2 or TPM 3 (the latter of which is not at all causally integrated). This should come as no surprise, as it is fairly intuitive that a stronger causal connection should be more informative about the future states of a system.

What does this have to do with causal emergence? Again, Hoel gives a clear demonstration by comparing the EI of two more TPMs, where this time one is a macroscale description of the other:

$$\begin{bmatrix} 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 0 \\ 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 0 \\ 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 0 \\ 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 0 \\ 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 0 \\ 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 0 \\ 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \text{TPM 4}$$

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \text{TPM 5}$$

In the first, 8-state TPM, the EI is relatively low (EI = 0.18), as only one state will transition with high probability, while each of the other seven states will transition with equal probability to any of the same seven. This means that if we apply an arbitrary intervention to the system, we would have (on average) relatively little knowledge of which state it will transition to next. Compare this with the second, 2-state TPM, where the seven unpredictable (micro) states of the first system have been grouped together into a single, now highly predictable, (macro) state. Under this new (macroscale) description, the EI of the system has risen to 1, as any arbitrary intervention will give us perfect knowledge of which state it will transition to next. If, as Hoel claims, EI is an adequate measure of causation, then there is ‘more’ causation in the macroscale description of this system (TPM 5) than in the microscale description (TPM 4), demonstrating a novel kind of causal emergence. Causation emerges under the macroscale description because we can more reliably predict the outcomes of our macroscale interventions, or to put it another way, the kinds of system that can be coarse-grained in the way described above are those that (according to Hoel) exhibit causal emergence.

The potential applications of this approach should be apparent. In particular, if mental states can be treated as macroscale coarse-grains of microscale neural structures, then this could offer a way of formally demonstrating that the macroscale possesses novel causal powers, and thus blocking reductionism or epiphenomenalism of the kind threatened by causal exclusion arguments (cf. Kim 2000, 2005). However, the sense in which EI measures causation, and the kind of emergence that this measure supports, is not yet clear. I will address each point in turn.

2 Is it really causal?

Hoel's commitment to an interventionist account of causation is relatively clear. His method for measuring EI is inspired by Pearl (2000), and the sense in which it tracks causal structure is very much interventionist in flavour (cf. Woodward 2003). Essentially, what EI measures is the extent to which an arbitrary intervention is predictive of the future state of a system, such that a highly structured (or determinate) system will have high EI, and a very chaotic (or indeterminate) system will have low EI.⁴ Another way of putting this is to say that we can intervene more reliably on a system with high EI, and it is in this sense that EI is a measure of causation.

When we coarse-grain over a system we are focusing our attention on particular aspects of its structure, by grouping together states whose (somewhat chaotic) outcomes we are happy to treat as equivalent (for whatever reason). In doing so we raise the EI of the system, because now our interventions will more reliably bring about the state that we are interested in – a macro-state that is a coarse-grained aggregate of many different fine-grained micro-states (compare TPM 5 with TPM 4). It is only in this interventionist sense that there is 'more' causation in the macroscale system than in its microscale equivalent, because we are able to more reliably predict the outcome of an arbitrary intervention (relative to our coarse-grained grouping of the micro-states).

Those who tend to get excited about causal emergence are typically interested in something like the emergence of genuinely novel causal *powers*, such that (for example) mental states could have causal powers over and above those of their physical supervenience base (see e.g. Robb 2019). It is not obvious that interventionism gets us anything like causal powers in this sense, rather than an account of causal *explanation*.⁵ Woodward's interventionist account is dis-

⁴List & Pivato (2015) have argued for something like the opposite of Hoel's causal emergence: the emergence of macroscale *indeterminacy* from a determinate microscale system. Similarly to Hoel, though, they advocate an ontic interpretation of this kind of emergence, despite its apparently epistemic appearance.

⁵See e.g. Reutlinger (2012) for some criticism of interventionism as a fundamental account of causation, and Andersen (2017) for an attempt to provide an "underlying metaphysics" for interventionist causation. Andersen's account stresses the importance of patterns and

tinctively pragmatic in flavour, insofar as it is concerned more with the role of causation in scientific methodology than with metaphysical questions about the ontology of causation (cf. Woodward 2015a). So it does not seem like the interventionist approach is well-equipped to tell us whether there are genuinely novel causal *powers* at the macro-level of some system, rather than merely novel causal *explanations*.⁶

Even if we were to grant that interventionist causation is genuinely causal, in something like the powers sense, it would still be unclear whether the macroscale interventions described by Hoel actually give rise to any ‘more’ causation than there already is at the microscale.⁷ One potential concern here has to do with so-called ‘fat-handedness’ objections (cf. Baumgartner & Gebharder 2016). If every macroscale structure supervenes upon a microscale structure,⁸ then any intervention on the macroscale will at the same time also be an intervention on the microscale, making it hard (or maybe even impossible) to distinguish causation at different scales. In order to measure EI at the macro-level, we need to intervene at the macro-level, but doing so will necessarily also involve a micro-level intervention. So even if we can intervene more reliably (at least relative to macro-level outcomes) by targeting the macro-level, it doesn’t seem like there are really any novel causal processes taking place at this level, because corresponding causal processes will always be taking place at the micro-level. This is not to say that there cannot be macro-level causation in some sense (cf. Shapiro & Sober 2007; Woodward 2015b), but rather that if EI were measuring an increase in causal power at the macro-level, then there would also need to be some commensurate change at the micro-level, which after all has just undergone the very same intervention (in the sense that any intervention must simultaneously target both macro- and micro-level).

None of this is to say that Hoel’s argument for causal emergence doesn’t track ‘causation’ of some kind, but rather that it is unlikely that this kind of causal emergence is going to satisfy anyone looking for something stronger than mere

information for interventionist causation, and as such seems like a natural fit with Hoel’s proposal for causal emergence.

⁶Hoel does seem aware of this distinction, writing that his account “does not contradict other theories of emergence, such as proposals that truly novel laws or properties may come into being at higher scales” (2017: 13), so it is possible that he would accept that he has only so far offered an account of emergent causal *explanations*.

⁷Of course, this assumes that we are taking micro-causation as our starting point, but insofar as Hoel’s account is an account of causal *emergence*, it seems fair to ask whether it can demonstrate the emergence of macro-causation over-and-above micro-causation. If it is just intended as an account of causation *at the macro-level*, rather than an account of causal emergence *per se*, then it would be much less exciting or controversial than it first appears.

⁸Hoel explicitly states that “The set of all possible causal [macro-]models, $\{\mathbf{S}\}$, is entirely fixed by the base S_m [the micro-structure]” (2017: 5). He goes on to note that “In technical terms this known as supervenience: given the lowest scale of any system (the base), all the subsequent macro causal models of that system are fixed” (*ibid.*). So his commitment to the supervenience of the (supposedly causally emergent) macro-structure on some underlying micro-structure should not be in question.

explanatory emergence. Either interventionism (and thus EI) tracks something like causal explanation rather than causal power, which is unlikely to satisfy many traditional emergentists, or else it is likely to suffer from something like fat-handedness objections. It is possible that the latter could be blocked (see e.g. Kästner & Anderson 2018 for recent discussion), but this is very much an ongoing debate, and one that Hoel will have to address as well. I will now turn to the question of emergence itself, and consider in what sense the causes measured by EI can really be described as ‘emergent’.

3 Is it really emergent?

In the literature on emergence, it is typical to distinguish between ontological (or strong) emergence and epistemic (or weak) emergence (see e.g. Silberstein & McGeever 1999; Chalmers 2006; Wilson 2015). This distinction is not entirely uncontroversial, and in particular there might be reason to further distinguish between strong and weak varieties of ontological emergence, etc., but for my purposes the ontological/epistemic distinction will suffice. Ontological emergence, in the sense that matters here, concerns the emergence of genuinely novel properties at some non-fundamental level, while epistemic emergence concerns the emergence of greater explanatory or predictive power at some higher level of description, relative to our epistemic capacities.

Hoel’s emergence might initially seem to be clearly epistemic, as it simply gives us a more informative description of coarse-grained processes that could also be described in fine-grained terms at the lower level, without introducing any genuinely novel properties or powers. Hoel himself seems to recognise this, writing that “a macroscale description of a system (a map) can be more *informative* than a fully detailed microscale description of the system (the territory)” (2017: 187, emphasis added). Understood in this way, Hoel’s claim is just that the higher level description can be more informative (in the sense measured by EI), not that it introduces anything that isn’t already present at the lower level.

However, as discussed in the previous section, it seems like Hoel is also committed to the emergence of novel causal powers at the macro-level, writing that causal emergence “is when the macro beats the micro in terms of efficacy, informativeness, or *power of its causal relationships*” (2017, emphasis added). Ontological emergence would certainly seem to offer a better characterisation of the emergence of genuinely novel causal powers. To say that novel causal powers can emerge at a higher level of organisation seems to be saying something metaphysically robust, and to entail the addition of something genuinely new to our ontology, not just the addition of explanatory or predictive power. But given Hoel’s interventionist approach to causation, it doesn’t seem that he can achieve strong emergence of this kind. Any intervention at the macro-scale is also going to involve an intervention at the micro-scale that it supervenes on, and so any higher-level ‘causation’ is also going to be accompanied by a change

in lower-level causal dynamics. Hoel's causal emergence just doesn't seem to be ontologically robust, and in this sense it is unclear whether we would should really say that it is 'causal' (at least in the sense of causal powers).

Bracketing for a moment the question of whether it is genuinely causal, Hoel's emergence could still be interesting even if it were merely epistemic. Even if he has only provided a formal approach to measuring the sense in which a higher level description of a system can be more informative, Hoel will already have made a valuable contribution (cf. Beckers & Halpern 2019 for another approach to this kind of question). By measuring the strength of an intervention in terms of effective information, he gives a helpful analysis of the circumstances under which a coarse-grained description of a system might be more (epistemically) beneficial than a fine-grained one, and consequently an analysis of the kinds of system where a 'special science' approach is most appropriate.

This epistemically emergent nature of complex systems is captured well by the idea of what Dennett calls 'real patterns', i.e. higher level descriptions of complex systems that are better at tracking relevant and informative regularities than lower level descriptions (Dennett 1991). Ross (2000; cf. Ladyman & Ross 2007) provides an information-theoretic presentation of this approach that will help to clarify how Hoel's proposed emergence can be both 'causal' (in the relevant interventionist sense) and nonetheless 'weak' (in the sense that it does not involve the emergence of any novel powers). According to Ross' definition, a pattern is real if and only if:

- (i) it is projectible under at least one physically possible perspective;
and
- (ii) it encodes information about at least one structure of events or entities S where that encoding is more efficient, in information-theoretic terms, than the bit-map encoding of S, and where for at least one of the physically possible perspectives under which the pattern is projectible, there exists an aspect of S that cannot be tracked unless the encoding is recovered from the perspective in question.

(Ladyman & Ross 2007: 226; adapted from Ross 2000)

This neatly captures the exact sense in which Hoel wants to say that a causal structure can emerge at a higher level. This can happen when there is a higher level pattern that encodes information about the whole structure (at every level) in a more efficient manner than the bit-map encoding (which simply performs a one-to-one mapping between microstructural states of the system and states of our description). Furthermore, this pattern is projectible, in the sense that it allows us to make predictions about future states of the system, and insofar as these predictions could not have been made without appealing to that pattern, we could say that they are (weakly or epistemically) emergent.⁹ Macroscale patterns of this kind are *real*, according to Ladyman & Ross, because the question

⁹Ladyman & Ross (2007: 193) express sympathy with the kind of weak emergentism defended

of which patterns encode information more efficiently is to be determined objectively. Andersen also stresses the objective nature of real patterns, although she concedes that “pattern ontologies” might nonetheless be perspectival in the sense that “different interests or goals lead to genuinely different ways of carving up the world into patterns” (2017: 13 of preprint). Hoel’s approach to causal emergence can be seen as giving us something like a formal method for answering this question, by calculating which description of a system gives the greatest EI (Andersen appeals for just a method in the conclusion of her 2017 paper). Nonetheless, it would be a mistake to conclude that this means there is greater causal *power* at the macroscale, as nothing has changed other than our way of describing the system.

In order to illustrate this more clearly, consider what is required for a causally emergent real pattern to occur. There must be a microstructurally complex system that defies concise description at the lower level, but which exhibits stable (and thus more easily predictable) regularities at the higher (or macrostructural) level. These regularities constitute ‘real patterns’ in the Dennettian sense, as they provide a more efficient encoding of certain features of the system’s structure. This can be considered a form of (weak) causal emergence insofar as it facilitates predictions and interventions that we could not otherwise have made. However, it is always possible (at least in principle) that the same predictions and interventions could have been made from the lower level, provided that we were able to keep track of every state and process of this complex system in a computationally tractable manner.¹⁰ Another way of putting this is to say that there must always be some physical microstructure responsible for implementing the macrostructural patterns that enable higher level predictions and interventions.¹¹ Furthermore, any intervention at the higher level is going to simultaneously intervene on the lower level structure, and any subsequent changes at the higher level will also supervene on changes at the lower level (see fn.7 for Hoel’s explicit commitment to supervenience). So even if we agree with Hoel that certain higher level patterns ought to be described as causally emergent, we have no reason to think that this emergence is strong in any ontologically robust sense.

Weak causal emergence of this kind might nonetheless be sufficient for everything that we could want from an account of higher level causation. I take it that this is what Hoel has in mind when he talks about ‘macro beating mi-

by Batterman (2002), but indicate that they prefer to avoid talking of ‘emergence’ altogether, instead using the label ‘scale relativity of ontology’ to refer to the kinds of macroscale patterns discussed here. Ross & Spurrett (2004) defend a similar kind of view about the circumstances under which macroscale explanations might be superior to microscale.

¹⁰Just as the “Laplacian super-physicists” introduced by Dennett (1981) are able to.

¹¹Given their hylomorphist sympathies, some defenders of IIT might deny this, but doing so would not convince anyone interested in the *emergence* of novel causes from a physical microstructure, rather than the autonomy (or even priority) of higher level causal structures. The latter would be an interesting result in its own right, but this seems importantly distinct from the emergence debate.

cro’, or there being greater ‘power’ in the causal relationships at higher levels of description – not that anything ontologically novel has emerged, but rather that this level of description gives us a better grip on the ‘real’ causal dynamics of the system. Perhaps this is in fact all there is to causation, as Andersen (2017) argues, in which case we can see Hoel and colleagues as providing some of the mathematical tools necessary for the analysis of genuine higher level causal powers (at least in the weak, interventionist sense). This could even get us ontological emergence in a deflationary sense: if there is nothing more to causation than information-theoretic dynamics, then weak emergence may suffice, but it will still not be the kind of causal emergence that would satisfy those with a more traditional idea of what causation is. I am not interested here in taking sides in these more fundamental debates about the nature of causation, but rather just want to make it clear what exactly Hoel’s proposal can offer.

Conclusion

I have presented Hoel’s proposal for a new model of causal emergence, based on a formal measure of ‘effective information’, and argued that it is causal only in the somewhat deflationary interventionist sense, and thus emergent only in a weak or epistemic sense. Anyone hoping for a demonstration of the strong emergence of ontologically novel causal powers should therefore not be too excited about this new information-theoretic proposal. Nonetheless, Hoel has provided a valuable analysis of the circumstances under which a macroscale model is more informative about the underlying causal dynamics of a system, at least relative to our epistemic interests. The Dennettian ‘real patterns’ framework offers a useful way of thinking about this analysis, as it allows us to see how Hoel’s emergence could be both ‘causal’ (in that it helps us to identify real causal structure) and at the same time ‘epistemic’ (in that it does not involve the emergence of genuinely novel causal powers). This could even allow for something approaching a kind of ‘weak ontological’ emergence, but only if one were willing to accept as fundamental the somewhat deflationary account of causation offered by Hoel and the interventionists. My aim here was not to intervene in metaphysical debates about the nature of causation, but rather just to clarify what kind of causal emergence is offered by Hoel’s proposal, so that those who might be looking for something stronger at least know what they are getting.

Acknowledgements

I would like to thank Matteo Grasso, Cecily Whitely, David Silverman, Jonny Lee, Krzysztof Dolega, Natesh Ganesh, and audience members at the Second Joint MCMP-Hannover Workshop “Philosophy of Science” and the 2019 Joint Session for their helpful comments, suggestions, and encouragement at various stages of this project.

References

- Albantakis, L., Marshall, W., Hoel, E., & Tononi, G. 2019. "What Caused What? A Quantitative Account of Actual Causation Using Dynamical Causal Networks." *Entropy*, 21/5: 459.
- Andersen, H.K. 2017. "Patterns, Information, and Causation." *The Journal of Philosophy*, 114/11: 592-622.
- Baumgartner, M., & Gebharter, A. 2016. "Constitutive relevance, mutual manipulability, and fat-handedness." *British Journal for the Philosophy of Science*, 67: 731-56.
- Batterman, R.W. 2002. *The Devil in the Details*. Oxford: OUP.
- Beckers, S., & Halpern, J.Y. 2019. "Abstracting Causal Models." In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*.
- Chalmers, D. 2006. "Strong and Weak Emergence." In Clayton & Davies (eds.), *The Re-emergence of Emergence*. Oxford: OUP.
- Dennett, D. 1981. "True Believers." In Haugeland (ed.), *Mind Design*. Cambridge, MA: MIT Press.
- Dennett, D. 1991. "Real Patterns." *The Journal of Philosophy*, 88/1: 27-51.
- Griffiths, P., Pocheville, A., Calcott, B., Stotz, K., Kim, H., & Knight, R. 2015. "Measuring Causal Specificity." *Philosophy of Science*, 82: 529-55.
- Hoel, E. 2017. "When the Map is Better Than the Territory." *Entropy*, 19/188: doi:10.3390/e19050188
- Hoel, E., Albantakis, L., & Tononi, G. 2013. "Quantifying causal emergence shows that macro can beat micro." *Proceedings of the National Academy of Science*, 110: 19790-5.
- Hoel, E., Albantakis, L., Marshall, W., & Tononi, G. 2016. "Can the macro beat the micro? Integrated information across spatiotemporal scales." *Neuroscience of Consciousness*, 1: doi:10.1093/nc/niw012
- Kästner, L. & Anderson, L. M. 2018. "Intervening into mechanisms: Prospects and challenges." *Philosophy Compass*, 13:e12546.
- Kim, J. 2000. *Mind in a Physical World*. Cambridge, MA: MIT Press.

- Kim, J. 2005. *Physicalism or Something Near Enough*. Princeton: Princeton University Press.
- Ladyman, J. & Ross, D. 2007. *Everything Must Go*. Oxford: OUP.
- List, C. & Pivato, M. 2015. “Emergent Chance.” *Philosophical Review*, 124/1: 119-52.
- Oizumi, M., Albantakis, L., & Tononi, G. 2014. “From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0.” *PLoS Computational Biology*, 10/5: e1003588.
- Owen, M. 2019. “Exploring Common Ground between Integrated Information Theory and Aristotelian Metaphysics.” *Journal of Consciousness Studies*, 26/1-2: 163-87.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Elsevier.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge: CUP.
- Reutlinger, A. 2012. “Getting Rid of Interventions.” *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 18/4: 787-95.
- Robb, D. 2019. “Emergent mental causation.” In Gibb, Hendry, & Lancaster (eds.), *The Routledge Handbook of Emergence*. Abingdon: Routledge.
- Ross, D. 2000. “Rainforest realism: A Dennettian theory of existence.” In Ross, D., Brook, A., & Thompson, D. (eds.), *Dennett’s Philosophy: A Comprehensive Assessment*. Cambridge, MA: MIT Press.
- Ross, D. & Spurrett, D. 2004. “What to say to a skeptical metaphysician: A defense manual for cognitive and behavioral scientists.” *Behavioral and Brain Sciences*, 27: 603-47.
- Shapiro, L. & Sober, E. 2007. “Epiphenomenalism—The Do’s and Don’t’s.” In Wolters & Machamer (eds.), *Thinking About Causes*. Pittsburgh: University of Pittsburgh Press.
- Silberstein, M. & McGeever, J. 1999. “The Search for Ontological Emergence.” *The Philosophical Quarterly*, 49/195: 201-14.
- Spirtes, P., Glymour, C., & Scheines, R. 1993. *Causation, Prediction, and Search*. Springer.

Tononi, G. & Sporns, O. 2003. "Measuring information integration." *BMC Neuroscience*, 4: 31.

Wilson, J. 2015. "Metaphysical Emergence: Weak and Strong." In Bigaj & Wuthrich (eds.), *Metaphysics in Contemporary Physics*. Poznan Studies in the Philosophy of the Sciences and the Humanities.

Woodward, J. 2003. *Making Things Happen*. Oxford: OUP.

Woodward, J. 2015a. "Methodology, ontology, and interventionism." *Synthese*, 192: 3577-99.

Woodward, J. 2015b. "Interventionism and Causal Exclusion." *Philosophy and Phenomenological Research*, XCI/2: 303-47.