Matt Jones
18 Nov 2020

Research Statement

My research develops and tests computational models of human thought. I apply my models to a wide range of domains within Cognitive Psychology and Cognitive Science more broadly, including perception, memory, learning, concept acquisition, knowledge representation, problem solving, and decision making. Typically I study multiple such functions at once, to understand how they interact to produce complex cognition. I formulate models mathematically and implement them in computer simulations or derive predictions from them analytically (i.e., via mathematical proof). I design and conduct behavioral experiments to test model predictions and to decide among competing models. The models I develop draw on tools of advanced mathematics, statistics, machine learning, and sometimes physics. I also study foundational issues of modeling, to help understand what models tell us and how they relate to scientific theory.

The benefits of computational modeling in cognitive science are manifold: They offer a more rigorous understanding of the workings of the mind than is possible from verbal theories. Accordingly, they enable more precise prediction and hence more informative empirical tests. Likewise, cognitive models facilitate optimization of training and education, via simulation of the efficiency and efficacy of alternative training paradigms (materials, timing, etc.). Models can be used as diagnostic tools, in that parameter estimates from fitting a model to a person's behavior or brain activity can serve as dependent measures that reflect separable psychological processes more accurately than can standard summary statistics. Such model-based measures are valuable in studies of individual differences and in clinical applications. Finally, rigorous computational characterizations of human thought can inform machine learning and artificial intelligence, enabling artificial systems to reproduce aspects of human intelligence and creativity that are currently beyond reach.

Since my tenure review in 2014, I have continued my core research program in mathematical modeling of cognition while expanding into several new fields. I began a collaboration in Computational Psychiatry with researchers at CU Denver and NIMH, applying models of learning, decision making, and response time to improve diagnosis and treatment of attentional and emotional disorders. I began a collaboration in Civil Engineering to improve planning and spatial reasoning in skilled construction trades. I started a new line of research applying my modeling approach to Quantum Physics, yielding new insights into the mathematical features of quantum entanglement as well as possible parallels in human decision making. I also continued existing collaborations, applying models and principles of cognitive psychology to classroom education, and applying statistical methods to population genetics.

During this six-year period, my research has produced 20 peer-reviewed publications, with 7 more submitted for publication. I have obtained four new grants from the National Science Foundation (NSF), National Institute of Mental Health (NIMH): one as PI, two as MPI or co-PI, and one as coinvestigator. I am co-PI on a newly funded internal grant from CU's AB Nexus program. I also completed five grants that were active at the time of tenure, one as PI, two as co-PI, and two as coinvestigator. I gave six invited talks (national and international) and made or co-authored 41 conference presentations.

Core research in Cognitive Psychology

My primary theoretical interests lie in the interactions among learning, decision making, and knowledge representation. How does what people learn depend on their background knowledge or how they represent that knowledge? Reciprocally, how does learning about a domain change representations? How do people learn and represent abstract relational concepts? How do people select or seek information in the service of learning? How do the decisions people make depend on their state of knowledge, and how does decision making change with learning? My students, collaborators, and I have addressed these questions in a variety of areas, using various experimental methods and modeling techniques.

Attention in Learning

It is well established that attention affects learning, in that people and nonhuman animals learn more about attended than unattended stimuli. More intriguing is that learning feeds back to impact attention, such that the degree to which different stimuli reliably predict outcomes affects how much attention they receive in new situations. Thus attention itself is learned. Various theories have been proposed for explaining the learning dynamics of attention (Kruschke, 2001; Le Pelley, Mitchell, & Johnson, 2013; Mackintosh, 1975; Pearce & Hall, 1980), but often they are incompatible and none can explain the full range of established empirical phenomena. My senior PhD student Sam Paskewitz and I have developed a new model, CompAct, embodying the principle that stimuli compete for attention at the time of prediction or response selection, and that attention is learned by the same principles of reinforcement learning from prediction error that govern simpler (associative) forms of learning. We have recently published our first theoretical paper on this model, showing how it relates to and unifies previous theories (Paskewitz & Jones, 2020). We have also run a series of experiments that support CompAct's predictions over other models (Paskewitz & Jones, submitted). The model and experimental paradigms we have developed have played a central role in our work in Computational Psychiatry, described below.

Abstract Relational Concepts

My former PhD student Dan Corral and I have spent many years studying how people learn and represent abstract relational concepts. Building on Gentner's (1983) influential structure-mapping theory of analogical reasoning, we have proposed a dual-representation theory of relational concepts, whereby they can be represented either atomically or compositionally. An atomic representation is one that operates directly on constituent objects, such as BIGGER(OBJECT$_1$,OBJECT$_2$). A compositional representation is built of atomic relations operating on shared objects. To use Gentner's classic example, the concept ORBIT could be represented as CAUSE[BIGGER(OBJECT$_1$,OBJECT$_2$) & ATTRACT(OBJECT$_1$,OBJECT$_2$), REVOLVE_AROUND(OBJECT$_2$,OBJECT$_1$)]. Thus, the concept derives its meaning from a second-order causal relation between three primitive relations, all operating on the same objects. Although this distinction is not explicit in structure-mapping theory, it is an essential assumption. We have proposed instead that people have flexibility to represent many concepts in either way, for example collapsing the compositional representation above to an atomic one, ORBIT(OBJECT$_1$,OBJECT$_2$). This hypothesis has implications for learning, memory, and decision making, and we conjecture that it exemplifies the sort of representational flexibility that is essential to human problem solving and has been so elusive in artificial intelligence. In a series of concept-learning experiments in collaboration with Ken Kurtz, we demonstrated that people can shift their representations as we have proposed, to take advantage of the processing efficiency afforded by atomic representations (Corral, Kurtz, & Jones, 2018). Our findings were further corroborated by a longer series of experiments in Dan's dissertation.

Computational Modeling of Reinforcement Learning

Reinforcement learning from prediction error (RL) is one of the greater successes in efforts to link psychological theory to brain function. As such, it offers rich opportunities for applying cognitive models to neurophysiological measures such as EEG and functional MRI. In collaboration with Tor Wager's group, I developed two computational models to explain how RL is biased by prior knowledge (Jepma, Koban, van Doorn, Jones, & Wager, 2018). In their experiment, subjects rated thermal pain administered following two visual cues, which they falsely believed to signal different temperatures. Subjects' ratings were biased by the cues, and surprisingly this bias was not unlearned even after dozens of trials. Fitting the models to the data enabled us to identify two mechanisms contributing to this persistent bias: biased pain perception and biased updating of pain expectations. Correlating the models' dynamics with fMRI signals enabled us to separately localize the brain networks mediating them. Understanding these computational mechanisms and their associated brain networks is an important step in addressing why people persist in their beliefs despite contrary evidence.

In a similar collaboration with Keith Lohse and Matt Miller, I applied an RL model to their single-trial EEG data on reward positivity (RewP), a brain signal that indexes prediction error (Holroyd & Coles, 2002). The model separated the dynamics of reward expectation from those of reward representation (the arithmetic difference between these is the prediction error). This separation in turn enabled us to assess the contribution of RL to learning at different time scales. We found that RL mechanisms drive short-term (trial to trial) adjustments in response probabilities but have no causal role in long-term acquisition of category knowledge. Thus our findings place important boundaries on the contribution of RL to higher-level learning (Lohse, Miller, Daou, Valerius, & Jones, 2020).

RL also plays a role in cognitive control, the executive functions by which stimuli and response strategies are selected in accordance with task goals. In collaboration with Daniel Weissman, I have been studying this connection in the context of the congruency sequence effect (CSE), whereby the performance cost of distracting information is modulated by whether it was congruent or incongruent with task goals on previous trials (Gratton, Coles, & Donchin, 1992). Building on my earlier work linking sequential effects in learning to knowledge representation (Jones, Curran, Mozer, & Wilder, 2013), we conducted a series of experiments demonstrating that the CSE reflects flexible mechanisms for learning statistical dependencies in the environment, rather than simpler mechanisms for managing response conflict as has previously been proposed (Weissman, Grant, & Jones, in press).

Active Learning and Information Selection

Most empirical studies on human learning provide subjects with stimuli and feedback and assess what they can learn. This dominant paradigm neglects a critical aspect of everyday learning, namely that people can choose what information to acquire (what questions to ask, what objects to explore, etc.). The growing body of research on active learning addresses how people select information (Gureckis & Markant, 2012). Of particular interest is sequential information selection, where the choice of each question can be informed by the answers to preceding ones, such as with a scientist performing a series of experiments. Do people plan ahead in such situations, choosing an experiment based on whether it leads to useful followup experiments, or do they choose myopically, performing the most informative experiment at each step? In collaboration with Jonathan Nelson and Björn Meder, I conducted an information-theoretic analysis drawing on the same principles of dynamic control used in more advanced RL models, which identified the conditions under which myopic selection is at odds with long-run efficiency (Nelson, Meder, & Jones, submitted). We then conducted an empirical study that demonstrated both children and adults are strikingly myopic when seeking information sequentially (Meder, Nelson, Jones, & Ruggeri, 2019). These findings place a strong upper bound on the sophistication of people's active learning strategies, and they point to new research questions on how active learning could be improved in both educational and scientific (optimal experiment design) settings.

Foundational Issues

I have long been interested in the relationship between rational and algorithmic models of cognition. Rational models describe the outcomes the mind attempts to optimize ("why"), whereas algorithmic models describe the representations and processes by which calculations are carried out ("how") (Marr, 1982). These two approaches are complementary, and understanding how they relate often yields greater insight than either approach alone. Typically, rational models are considered limiting cases of algorithmic ones: They describe what the mind would do if it had infinite computational resources (Griffiths, Vul, & Sanborn, 2012). Instead, Brad Love and I have shown analytically how many algorithmic models can be treated as limiting cases of rational ones, when the inductive bias (formally, a Bayesian prior distribution) becomes infinitely strong (Parpart, Jones, & Love, 2018; Bobadilla-Suarez, Jones, & Love, submitted). This provides a new way to think about algorithmic models, in terms of the assumptions they embody for the statistical structure of the environment. I recently received a grant from NSF apply this approach to a wide range of models, spanning decision making, RL, and concept learning. In addition to the theoretical integration afforded by our approach, it has the potential to generate new models for both psychology and machine learning.

I have also shown how a rational interpretation can provide useful grounding for models of decision making and response time. My earlier work (Jones & Dzhafarov, 2014) demonstrated that these models are often unfalsifiable because of flexibility in their ancillary (extra-theoretical) assumptions, a line of work that I have continued in recent years (Jones, submitted). On the more positive side, I have shown that the most influential model in this family (the diffusion model; Ratcliff, 1978) can be derived from rational principles of Bayesian inference. This normative grounding eliminates the model's excess flexibility and opens connections to other theoretical frameworks, such as RL to explain how decision making adapts with learning (Jones, 2018).

Computational Psychiatry

Computational methods are becoming increasingly influential in Clinical Psychology and Psychiatry (Huys, Maia, & Frank, 2016), and funding agencies such as NIMH have recently made Computational Psychiatry a priority. Cognitive modeling can contribute to the study of psychopathology in numerous ways, by offering more accurate diagnosis, informing nosology, developing and evaluating new interventions, and predicting and assessing response to existing treatments (Stoddard & Jones, 2019). In collaboration with Joel Stoddard (Psychiatry, CU Denver) and Melissa Brotman's group at NIMH, I have been applying models of learning and decision making to disorders including anxiety, adolescent irritability, and ADHD.

Several of our projects have used variants of the diffusion model, a widely used model of decision making and response time that affords separate measurement of processing efficiency (drift rate), response caution (decision threshold), response bias (starting point), and nondecision time. We have found that the increased response-time variability seen in sustained attention tasks in youth with attention or mood disorders is accounted for by lower drift rates, and surprisingly that this effect is greater in the absence of distractor stimuli (Haller, Stoddard, Pagliaccio, Bui, MacGillivray, Jones, & Brotman, in press). We have developed a model that separates perceptual bias from response bias in adolescents' discrimination of facial cues, and found that these measures correlate differently with anxiety versus irritability (Haller, …, Jones, & Brotman, in preparation). We have combined the diffusion model with a model of spatial attention to assess how an emotionally charged cue (an angry face) biases spatial attention in a subsequent perceptual discrimination task, and how this effect correlates with anxiety (Haller, Stoddard, Jones, & Brotman, in preparation). These results all help to tie the disorders in question to well-studied experimental tasks from Cognitive Psychology and to pinpoint the cognitive processes that are disrupted.

In other projects we have applied models of reinforcement or attention learning to isolate disruptions in learning processes. We developed an RL model of a therapeutic intervention designed to correct biased interpretation of affective social cues (facial expressions), and found that the model can distinguish deficits due to attentional bias, bias in feedback processing, and poor discrimination of facial cues (Stoddard, Haller, Costa, Brotman, & Jones, submitted). Funded by an R21 grant from NIMH, we applied the CompAct model of attention learning described above (Paskewitz & Jones, 2020) to a new associative learning task with facial cues, finding that individuals higher in anxiety can perform worse in this task because they pay less attention to (i.e., avoid) angry faces or because they are slower to shift their attention to predictive social cues. These results show that individuals with seemingly similar difficulties at a behavioral level can suffer quite different disruptions in cognitive processing, and they offer a potential path toward diagnostic methods for individually targeted interventions. We are now pursuing these possibilities in a newly funded grant from CU's AB Nexus program.

Civil Engineering

I have recently begun an NSF-funded project with Paul Goodrum and Matt Hallowell (Civil Engineering) and Tom Yeh (Computer Science) to investigate the impact of augmented reality (AR) and related technologies on skilled workers in the construction trades. This grant is part of NSF's Future of Work at the Human-Technology Frontier program and is an excellent opportunity for computational principles from Cognitive Psychology to have a broad practical impact. Replacing 2d blueprints with AR could dramatically improve worker efficiency and safety, but it also comes with risks of information overload and distraction. Building on cognitive theories of planning (Hayes-Roth & Hayes-Roth Perrault, 1979; Spiegel, Koester, & Schack, 2013), inference under uncertainty (Daw, Niv, & Dayan, 2005), and the impact of trust on planning (La Porta, Lopez De Silanes, Schleifer, & Vishny, 1997; Michaelson & Munakata, 2016), I helped to develop a set of hypotheses regarding how various aspects of an AR system might impact critical outcome measures including planning horizon, ability to anticipate hazards, and need for replanning or rework.

Although some form of AR will likely become common on jobsites in the near future, AR is not currently used by field workers. The pilot experiments and subject-matter expert interviews we have conducted thus far have answered many basic questions about how people interact with such a system during a construction task, and what capabilities they find most useful. We are currently designing the apparatus and task for our main experiment, which will contrast various AR and 2d conditions on the aforementioned measures.

This collaboration has also opened up other lines of work. We recently completed a theoretical paper on statistical measures of workplace safety in which I was able to use my statistics expertise to evaluate the validity and reliability of accident rates as indicators of safety performance (Hallowell, Quashne, Salas, Jones, MacLean, & Quinn, in press).

Physics

Quantum entanglement is arguably the most profound discovery ever made about the nature of reality. Empirical violations of Bell's inequality (Bell, 1964) and related conditions (Clauser, Horne, Shimony, & Holt, 1969; Kochen & Specker, 1967), collectively known as contextuality, demonstrate that physical quantities do not have meaningful values when not being measured, not even in a probabilistic sense. However, the argument behind this conclusion requires measurements to be separated in such a way that they cannot causally influence each other's outcomes. That requirement impedes application of contextuality to more complex systems, both in Physics and beyond. For example, several recent studies have attempted to demonstrate contextuality in human behavior, via particular patterns of correlations between decisions made by an individual that mirror the correlations observed in particle physics (Aerts et al., 2017; Cervantes & Dzhafarov, 2018).

In collaborations with Ehtibar Dzhafarov, I have helped to elucidate how questions of contextuality fundamentally change when measurements can directly influence each other, such as how any decision a person makes can affect their subsequent decisions (Dzhafarov, Kujala, Cervantes, Zhang, & Jones, 2016). My main contribution in this area has been to show how the framework of probabilistic causal models frequently used in Cognitive Science can resolve this problem (Jones, 2019). I proposed a model-based definition of contextuality that accounts for direct influence between measurements, and proved this definition is equivalent to a purely probabilistic definition by Dzhafarov and Kujala (2016) and also to a principle of no-fine-tuning proposed in the Physics literature by Wood and Spekkens (2015) and Cavalcanti (2018). The advantage of the model-based approach is that—like John Bell's original proof—it directly specifies the class of physical models of a system that are ruled out if it is found to be contextual. My approach is beginning to be cited in the Physics literature (e.g., Pearl & Cavalcanti, in press). This is notable because, although there is a long history of formal methods being imported from Physics to Psychology, there are very few examples of the reverse.

References

Aerts, D., Arguëlles, J. A., Beltran, L., Geriente, S., de Bianchi, M. S., Sozzo, S., & Veloz, T. (2017). Spin and wind directions I: Identifying entanglement in nature and cognition. *Foundations of Science*, *23*, 323-335.

Bell, J. S. (1964). On the Einstein–Podolsky–Rosen paradox. *Physics, 1*, 195-200.

Bobadilla-Suarez, S., Jones, M., & Love, B. C. (submitted for publication). Robust priors for regularized regression.

Cavalcanti, E. G. (2018). Classical causal models for Bell and Kochen–Specker inequality violations require fine-tuning. *Physics Review X, 8*, 021018.

Cervantes, V. H., & Dzhafarov, E. N. (2018). Snow Queen is evil and beautiful: Experimental evidence for probabilistic contextuality in human choices. *Decision, 5*, 193-204.

Clauser, J., Horne, M., Shimony, A., & Holt, R. (1969). Proposed experiment to test local hidden-variable theories. *Physical Review Letters, 23,* 880-884.

Corral, D., Kurtz, K. J., & Jones, M. (2018). Learning relational concepts from within- vs. between-category comparisons. *Journal of Experimental Psychology: General, 147*, 1571-1596.

Daw, N., Niv, Y. & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*, 1704-1711.

Dzhafarov, E. N., & Kujala, J. V. (2016). Context-content systems of random variables: The contextuality-by-default theory. *Journal of Mathematical Psychology, 74*, 11-33.

Dzhafarov, E., Kujala, J., Cervantes, V., Zhang, R., & Jones, M. (2016). On contextuality in behavioural data. *Philosophical Transactions of the Royal Society A, 374*, 20150234.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7,* 155-170.

Gratton, G., Coles, M. G. H., & Donchin, E. (1992). Optimizing the use of information: strategic control of activation and responses. *Journal of Experimental Psychology: General, 4*, 480-506.

Griffiths, T. L., Vul, E., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science, 21*, 263-268.

Gureckis, T.M. and Markant, D.B. (2012). A cognitive and computational perspective on self-directed learning. *Perspectives on Psychological Science*, 7, 464-481.

Haller, S. P., Stoddard, J., Jones, M., & Brotman, M. A. (manuscript in preparation). Effects of emotionally charged distractors on spatial attention.

Haller, S. P., Stoddard, J., Kircanski, K., Chen, G., Cárdenas, S. I., MacGillivray, C., Botz-Zapp, C. A., Bui, H., Stavish, C., Jones, M., & Brotman, M. A. (manuscript in preparation). Unique and shared neural mechanisms of face-emotion processing in anxious and irritable youth: A computational modeling approach.

Haller, S. P., Stoddard, J., Pagliaccio, D., Bui, H., MacGillivray, C., Jones, M., & Brotman, M. A. (in press). Computational modeling of attentional impairments in disruptive mood dysregulation and attention deficit/hyperactivity disorder. *Journal of the American Academy of Child and Adolescent Psychiatry*.

Hallowell, M., Quashne, M., Salas, R., Jones, M., MacLean, B., & Quinn, E. (in press). The statistical invalidity of TRIR as a measure of safety performance. *Professional Safety*.

Hayes-Roth, B., & Hayes-Roth Perrault, F. (1979). Cognitive Model of Planning. *Cognitive Science, 3*, 275-310.

Holroyd, C., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review, 109*, 679-709.

Huys, Q. J. M., Maia, T. V., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience, 19*, 404-413.

Jepma, M., Koban, L., van Doorn, J., Jones, M., & Wager, T. D. (2018). Behavioural and neural evidence for self-reinforcing expectancy effects on pain. *Nature Human Behaviour, 2*, 838-855.

Jones, M. (2018). The diffusion model of speeded choice, from a rational perspective. In W. Batchelder et al. (Eds.), *New Handbook of Mathematical Psychology, Vol. 2* (pp. 71-103). Cambridge University Press.

Jones, M. (2019). Relating causal and probabilistic approaches to contextuality. *Philosophical Transactions of the Royal Society A, 377*, 20190133.

Jones, M. (submitted for publication). Flexibility of evidence-accumulation models under practical constraints.

Jones, M., Curran, T., Mozer, M. C., & Wilder, M. H. (2013). Sequential effects in response time reveal learning mechanisms and event representations. *Psychological Review, 120*, 628-666.

Kochen, S., & Specker, E. P. (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics, 17*, 59–87.

Kruschke, J. K. (2001). Toward a Unified Model of Attention in Associative Learning. *Journal of Mathematical Psychology*, *45*(6), 812–863.

La Porta, R., Lopez De Silanes, F., Schleifer, A., & Vishny, R. (1997). Trust in large organizations. *The American Economic Review: Papers and Proceedings*, 333-338.

Le Pelley, M. E., Mitchell, C. J., & Johnson, A. M. (2013). Outcome value influences attentional biases in human associative learning: Dissociable effects of training and instruction. *Journal of Experimental Psychology: Animal Behavior Processes*, *39*(1), 39-55.

Lohse, K. R., Miller, M. W., Daou, M., Valerius, W., & Jones, M. (2020). Dissociating the contributions of reward-prediction errors to trial-level adaptation and long-term learning. *Biological Psychology, 149*, 107775.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*(4), 276-298.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman

Meder, B., Nelson, J. D., Jones, M., & Ruggeri, A. (2019). Stepwise versus globally optimal search in children and adults. *Cognition, 191*, 103965.

Michaelson, L. & Munakata, Y. (2016). Trust matters: Seeing how an adult treats another person influences preschoolers' willingness to delay gratification. *Developmental Science, 19*, 1011-1019.

Nelson, J. D., Meder, B., & Jones, M. (submitted for publication). On the fine line between "heuristic" and "optimal" sequential question strategies.

Parpart, P., Jones, M., & Love, B. C. (2018). Heuristics as Bayesian inference under extreme priors. *Cognitive Psychology, 102*, 127-144.

Paskewitz, S., & Jones, M. (2020). Dissecting EXIT. *Journal of Mathematical Psychology, 97*, 102371.

Paskewitz, S., & Jones, M. (submitted for publication). Predictiveness and reward effects on attention can be explained by a single mechanism.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87*, 532-552.

Pearl, J. C., & Cavalcanti, E. G. (in press). Classical causal models cannot faithfully explain Bell nonlocality or Kochen–Specker contextuality in arbitrary scenarios. *Quantum*.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85*, 59-108.

Spiegel, M. A., Koester, D. & Schack, T. (2013). The functional role of working memory in the (re-)planning and execution of grasping movements. *Journal of Experimental Psychology: Human Perception and Performance*, *39*, 1326-1339.

Stoddard, J., Haller, S. P., Costa, V., Brotman, M. A., & Jones, M. (submitted for publication). Measuring the mechanism of interpretation bias training to treat psychopathology.

Stoddard, J., & Jones, M. (2019). Computational modeling in pediatric mental health. *Journal of the American Academy of Child and Adolescent Psychiatry, 58*, 471-473.

Weissman, D. H., Grant, L. D., & Jones, M. (2020). The congruency sequence effect indexes response-general control. *Journal of Experimental Psychology: Human Perception and Performance, 46*, 1387-1396.

Wood, C. J., Spekkens, R. W. (2015). The lesson of causal discovery algorithms for quantum correlations: Causal explanations of Bell-inequality violations require fine-tuning. *New Journal of Physics, 17*, 033002.