

## PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/198370>

Please be advised that this information was generated on 2019-03-21 and may be subject to change.

# One wouldn't expect an expert bowler to hit only two pins: Hierarchical predictive processing of agent-caused events

Lieke Heil , Johan Kwisthout, Stan van Pelt, Iris van Rooij and Harold Bekkering

Quarterly Journal of Experimental Psychology  
2018, Vol. 71(12) 2643–2654  
© Experimental Psychology Society 2018



Article reuse guidelines:  
sagepub.com/journals-permissions  
DOI: 10.1177/1747021817752102  
qjep.sagepub.com



## Abstract

Evidence is accumulating that our brains process incoming information using top-down predictions. If lower level representations are correctly predicted by higher level representations, this enhances processing. However, if they are incorrectly predicted, additional processing is required at higher levels to “explain away” prediction errors. Here, we explored the potential nature of the models generating such predictions. More specifically, we investigated whether a predictive processing model with a hierarchical structure and causal relations between its levels is able to account for the processing of agent-caused events. In Experiment 1, participants watched animated movies of “experienced” and “novice” bowlers. The results are in line with the idea that prediction errors at a lower level of the hierarchy (i.e., the *outcome* of how many pins fell down) slow down reporting of information at a higher level (i.e., which *agent* was throwing the ball). Experiments 2 and 3 suggest that this effect is specific to situations in which the predictor is causally related to the outcome. Overall, the study supports the idea that a hierarchical predictive processing model can account for the processing of observed action outcomes and that the predictions involved are specific to cases where action outcomes can be predicted based on causal knowledge.

## Keywords

Predictive processing; hierarchy; action outcomes; causality; action observation

Received: 14 July 2017; revised: 20 October 2017; accepted: 27 November 2017

## Introduction

When we observe the actions of another person, we predict what this person is going to do next in order to decide what his or her aim is and to adapt our own response accordingly. To do this, we take the characteristics of this observed person into account. Imagine, for example, that you see a florist reaching for a vase. You could decide to offer help by handing it to him so he can use it to put flowers in. If, on the other hand, you see a small child reaching for the same vase, your prediction of the potential outcome will tell you that this might not be such a good idea, so you can intervene with the child's action and put the vase at an unreachable location. Previous research supports the idea that to interact with the world effectively, we use information from previous experiences to predict a specific visual stimulus (Den Ouden, Friston, Daw, McIntosh, & Stephan, 2009; Summerfield & Koehlin, 2008) or the outcomes of other

people's actions (Aglioti, Cesari, Romani, & Urgesi, 2008). In this article, we explore the idea that predictions about the outcomes of other people's actions arise in a generative model that has a hierarchical structure and consists of causal relations between different levels of this hierarchy.

This idea has its origin in the predictive processing framework. According to this framework, the brain is continuously predicting the input it will receive (Clark, 2013b; Friston, 2005). This means that our brains process

---

Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, The Netherlands

### Corresponding author:

Lieke Heil, Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, PO Box 9104, 6500 HE Nijmegen, The Netherlands.  
Email: l.heil@donders.ru.nl

incoming information not in a purely bottom-up fashion, but by a cascade of predictions from higher level to lower level representations. The top-down neural signal then consists of predicted states of the world and the bottom-up signal consists of the difference between these predictions and the actual input. These differences are called prediction errors. The framework assumes that if lower level representations are correctly predicted by higher level representations, this enhances processing, for example, in terms of speed (O'Reilly et al., 2013). However, when lower level representations are incorrectly predicted, additional processing at higher levels is required to deal with the prediction errors arising at lower levels.

In recent years, empirical evidence has been found to support the idea that the predictive processing framework successfully describes low-level sensory processing (e.g., Bastos et al., 2012; Kok, Jehee, & de Lange, 2012; Phillips, Blenkmann, Hughes, Bekinschtein, & Rowe, 2015; Rao & Ballard, 1999; for a review, see Summerfield & de Lange, 2014). Some researchers have pressed that predictive processing may serve as a general account of brain functioning (e.g., Clark, 2013b; Friston, 2010; Hohwy, 2013), in which case it should also describe processing at higher cognitive levels, including those involved in the processing of agent-caused events. Indeed, it has been suggested that neural responses to biological motion and other people's beliefs and desires are modulated by the predictability of an event, in compliance with the features of predictive processing (Kilner, Friston, & Frith, 2007; Koster-Hale & Saxe, 2013). In accordance with this idea, top-down and bottom-up signals in the brain have been found to be modulated by the probability of an agent-caused event (Van Pelt et al., 2016). Yet, little is known still about the potential cognitive principles that govern the prediction of action outcomes from agent information. In this article, we test whether a hierarchical predictive processing model is, in principle, able to account for the way in which we process the outcomes of other people's actions. We specifically explore a model in which the predictive relations between levels of the generative model are causal.

To study our hypothesis in a naturalistic but controlled setting, we used a behavioural paradigm in which participants viewed animated movies of people playing a bowling game. In these movies, the agent's action (i.e., throwing a bowling ball) caused an outcome (i.e., the pins were knocked down). In Experiment 1, the specific outcome (i.e., the score) could be predicted based on identification of the agent, that is, knowledge about the performance of this agent on previous trials. There were two agents, indicated as the novice and the experienced player, who usually obtained a low and a high score, respectively. In 25% of all trials, however, players obtained scores that were incongruent with their skills. According to the model that we explore here, such an

incongruence results in a prediction error at the level of the hierarchy at which "outcome" is represented. In the current bowling set-up, this prediction error should increase processing at the higher level of the hierarchy at which the "agent" is represented, as this is where the prediction error needs to be "explained away". To investigate the idea of a hierarchically organised model, we asked participants to report after each movie either which agent ("experienced" or "novice") or which outcome ("high score" or "low score") they observed. For each of these two questions, reaction times were compared for questions following predicted versus unpredicted outcomes. Reaction times have previously been found to correlate with the improbability of an event (Bestmann et al., 2008; Den Ouden, Daunizeau, Roiser, Friston, & Stephan, 2010) and are therefore assumed to reflect the prediction error. Given the predictive processing model that we investigate here, we hypothesise that the observation of an unpredicted outcome slows down the response to a question about the agent, as the prediction error arising at the outcome level is to be explained away at the agent level, requiring additional processing at that level and thereby slowing down the reporting of the inferred agent information. In other words, we predicted a longer reaction time for the agent question if it followed an unpredicted rather than a predicted outcome, as a prediction error needs to be explained away at this level. On the other hand, as a prediction error is assumed to be explained away at a level above, not at the level at which the error occurs itself, a prediction error occurring specifically at the outcome level is not expected to influence the reaction times for the outcome question. In this way, measuring reaction times for these questions affords us to test the hypothesis of a hierarchically structured predictive model in which information about agents is processed at a higher level in this model than information about outcomes.

Furthermore, we explored the role of knowledge about the causal structure of the world in the prediction of the outcomes of agent-caused events. Predicting what types of outcomes are likely produced by what types of agents presupposes a generative model that encodes such world knowledge. This knowledge tells us that some agents have a propensity to cause some events more often than others. For instance, a skilled bowler is more likely to hit a strike than a novice bowler. Also, we know that agents may cause certain events but not others. A person may cause an object to fall, but not the sun to shine. This seems to be mainly the result of knowledge about the mechanisms that cause events (Shultz, Fisher, Pratt, & Rulf, 1986). The influence of this type of knowledge also explains why young children are more likely to expect the movement of an object to be caused by an agent rather than by a train (Saxe, Tzelnic, & Carey, 2007), why people only learn to associate a tone

with an air puff if they are aware of the link between the two (Clark & Squire, 1998), and why the remembered speed of a movement is influenced by the effect it seems to have caused in Michotte's (1963) launching effect paradigm (Kerzel, Bekkering, Wohlschläger, & Prinz, 2000).

To investigate the hypothesis that a predictive processing model with causal relations between the levels is able to account for the type of processing that occurs at high levels of the hierarchy where agent-caused events are represented, we compare the results of this first experiment with the results of two follow-up experiments. Whereas in Experiment 1, the outcomes of actions (number of pins thrown over by an agent) could be predicted based on knowledge about the agent (his propensity to throw high or low scores); in Experiment 2, we created a situation in which the outcome could only be predicted from a coloured patch next to the agent. Colours, unlike agents, do not cause outcomes. Therefore, if the predictions made in Experiment 1 were specifically based on causal knowledge of agents causing outcomes, none of the effects predicted for Experiment 1 would be predicted for Experiment 2. In Experiment 3, the predictive cue was the agent's shirt colour. Given that according to world knowledge a shirt colour can be a cue to an agent's identity, we would again expect the same reaction time effects in Experiment 3 as predicted for Experiment 1. On the other hand, if the predictions in all experiments would be merely based on non-causal associations between two events, one would expect similar patterns for all experiments. The set of three experiments together allows us to explore whether or not a hierarchical predictive processing model with causal rather than non-causal relations between the levels can account for the processing of other agents' actions. As such, we hope our results will open the way for future research on this topic.

## Experiment 1

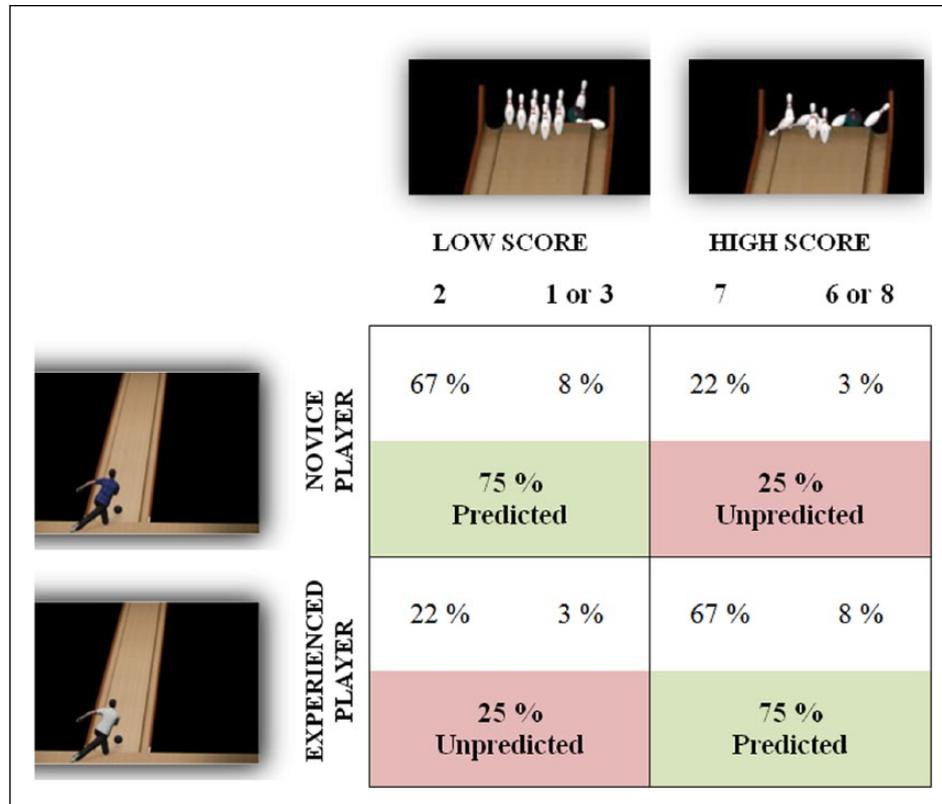
### Methods

**Participants.** A total of 28 healthy, right-handed participants (22 females) aged between 18 and 26 (mean age 22.6) years were recruited for the first experiment. They were paid €10 or received course credits for their participation. The study was approved by the institution's local ethics committee (ECG2012-0910-058) and written informed consent was obtained from each participant.

**Stimuli and design.** A total of 24 animated movies of a bowling game was created using Autodesk's 3ds Max 2014 and MotionBuilder 2014 (<http://www.autodesk.com>). These movies showed a bowling lane and one of two agents, who could be recognised by their clothing. The avatars for the bowling players were selected from Worldviz Vizard

Complete Characters (<http://www.worldviz.com/products/avatars/complete-characters>). In each movie, the agent threw a ball directed at the pins and disappeared at 1,200 ms after the start of the movie, to keep the visual display of the action outcome the same for the two agents. The ball then rolled towards the pins, either a little left or right of the centre and hit 1, 2, 3, 6, 7, or 8 pins. The kinematics of the ball movement only varied in terms of the direction and were not associated with a specific agent or outcome. Each movie lasted 5 s. By keeping the kinematics of the action constant, we were able to investigate predictions that are based on information about the agent, rather than on kinematics. In 75% of all 288 trials, the outcome was as expected based on the agent's skill level. This means that one agent, who was introduced to the participants as the novice player, received a low score (1, 2, or 3) in 108 out of 144 trials, whereas the other agent, who was introduced to the participants as the experienced player, received a high score (6, 7, or 8) in 108 out of 144 trials. More specifically, within the category of low scores, a score of 2 was most frequent (96 out of 144 trials), as was a score of 7 in the category of high scores (Figure 1). The other outcomes (i.e., 1, 3, 6, and 8) were included as fillers that provided variability in scores and thereby made the experiment more realistic. The movies were presented using Presentation software (version 17.2, <http://www.neurobs.com>).

**Procedure.** Participants were seated comfortably in front of the computer on which the experiment was presented. Instructions were presented on the screen and shortly repeated verbally. In the instructions, it was explained to participants that there were two agents, a novice and an experienced player, who usually (but not always) obtained scores that matched their skill level. In addition, participants were told that after each movie, they would be asked to answer one out of two questions, and that they should pay attention to everything they saw, as they would not know which question would be asked afterwards. They were also instructed to answer the question as quickly as possible. After the instructions, the participants performed four practice trials, in which they received information about which agent they would see, before the movie was presented. This allowed them to associate appearance of the agent with his skill level, as they would need this to perform the task. Immediately after each movie of the actual task, participants were asked to answer one out of two questions, which were presented in random order. One of the questions was about the agent ("Did you see the experienced or the novice player?"), whereas the other question was about the outcome ("Was the score high or low?"). The question was presented on the screen for 2 s, with the two answer options presented underneath. Participants answered by using their index fingers to press either the left button or the right button on a button box as quickly



**Figure 1.** Overview of conditions and stimuli in Experiment 1.

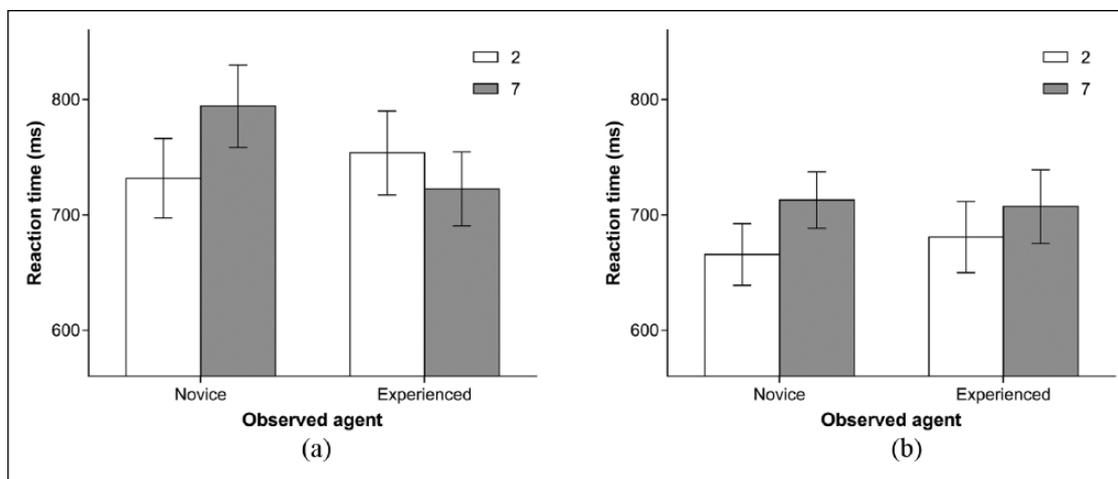
as possible. The order of the answer options was randomised to prevent motor preparation. In the practice trials, participants received feedback on the accuracy of their answer. This was not the case in the actual experiment. After each trial, a fixation cross was presented for a duration that was randomised between 500 and 2,500 ms.

Reaction times to the questions that followed movies in which the outcome was 2 or 7 were analysed using a 2 (agent: novice vs experienced player)  $\times$  2 (outcome: 2 vs 7)  $\times$  2 (question: agent vs outcome) repeated measures analysis of variance (ANOVA). The reaction times to questions that were answered incorrectly were not considered in this analysis. As all responses given after 2,000 ms were labelled incorrect, our final dataset did not include any trials with very long reaction times. The data were checked for trials with very short reaction times (<100 ms), but no such trials were found. Furthermore, reaction times to questions following movies with outcomes 1, 3, 6, and 8 were also not considered in the analysis, as they were included as fillers in the experiments to make the events look as naturalistic as possible and appeared very infrequently (a total of 32 out of 288 trials). Based on our hypothesis that additional processing is required at the agent level as a result of a prediction error, we anticipated that unexpected events (i.e., the novice agent obtaining a high score or the experienced agent a low score) would

bring about longer reaction times than expected events (the novice agent obtaining a low score or the experienced agent a high score). This results in an anticipated interaction effect between agent and outcome, specifically for the agent question.

### Results and discussion

The analysis showed a significant three-way interaction between agent, outcome, and question,  $F(1, 27)=7.41$ ,  $p=.01$ ,  $\eta_p^2=.22$  (Figure 2; see Supplementary Material for a full overview of results). Follow-up analyses showed that for the agent question, there was an interaction between agent and outcome,  $F(1, 27)=18.95$ ,  $p<.001$ ,  $\eta_p^2=.41$ . More specifically, paired-samples  $t$ -tests showed that, for this question, unexpected events resulted in higher reaction times for both the novice player,  $t(1, 27)=-4.00$ ,  $p<.001$ , and the experienced player,  $t(1, 27)=2.60$ ,  $p=.02$ . Overall, participants were on average 46.8 ms slower to respond when the outcome was not as would be predicted from the agent's skills, in line with the hypothesis that unexpected events result in a longer reaction time than expected events as a result of additional processing at the agent level. Such an effect did not appear when participants were asked about the outcome, where there was no interaction between agent and outcome,  $F(1, 27)=0.80$ ,  $p=.38$ ,  $\eta_p^2=.03$ . For this question, there only was a



**Figure 2.** Reaction times (mean  $\pm$  SEM) for (a) the agent question and (b) the outcome question in Experiment 1, separately for bowler expertise and outcome (scores 2 and 7).

significant effect of outcome on reaction time  $F(1, 27) = 19.36$ ,  $p < .001$ ,  $\eta_p^2 = .42$ . More specifically, a score of 7 resulted in a reaction time that was on average 36.9 ms longer than the reaction time following a score of 2.

Additional analyses confirmed that including the questions following movies that showed another outcome (i.e., 1, 3, 6, and 8) did not influence this pattern of results. Furthermore, the average accuracy over all trials was 95.2% (94.6% for the agent question and 95.8% for the outcome question), and the specific pattern of these data was incompatible with the possibility that the reaction time effects were driven by a speed-accuracy trade-off.

The three-way interaction shows that the prediction effects differ between the two questions. As participants were not aware which question they would be asked after each movie, it is unlikely that they made predictions in one situation, but not in the other. Rather, it seems that specifically for the agent question, reaction times were influenced by the violation of the participant's prediction. This is consistent with the idea that prediction errors at the lower level of the causal hierarchy, that is, the outcome level, slow down the reporting of information at a higher level, that is, the agent level.

In this first experiment, predictions about the outcome of another person's action could be based on a causal relation between the agent and the outcome. To investigate the specificity of this type of prediction to this causal relation, we performed a second experiment in which the score could not be predicted based on the agent's skills, but on an arbitrary statistical relation between a coloured box next to the agent and the outcome. If people's predictions indeed crucially depend on causal knowledge about agents causing outcomes, no effect on reaction times of predictability of events should be observed in Experiment 2.

## Experiment 2

### Methods

**Participants.** A total of 28 participants (23 females) aged between 19 and 29 (mean age 23.1) years took part in the second experiment. They were paid €10 or received course credits for their participation. The study was approved by the local ethics committee and written informed consent was obtained from each participant.

**Stimuli and design.** Out of the 24 animated movies from the first experiment, 12 movies with only one of the agents were selected. This means that there were still two ball directions (left and right) and six outcomes (1, 2, 3, 6, 7, and 8). Instead of two different agents, there were now two different coloured boxes (yellow and blue). In each movie, one of these boxes was presented next to the agent. Like the agent, the box was presented from the beginning of the 5-s movie and disappeared after 1,200 ms. The colour of the box correlated with the outcome. If one colour was presented, the outcome was likely to be low, whereas if the other colour was presented, the outcome was likely to be high. The distribution of trials was the same as in the first experiment, with 288 trials in total, of which 75% was as expected based on the colour of the box (Figure 3).

**Procedure.** The testing procedure was similar to the procedure for the previous experiment. Although the instructions were also largely the same, participants were now told that they would see a blue or yellow box that would indicate whether the score is more likely to be high or low. As an example, they were told that one of the colours might indicate that the player would probably get a low score, but that this would not always be the case. Again, participants were explicitly informed about the association

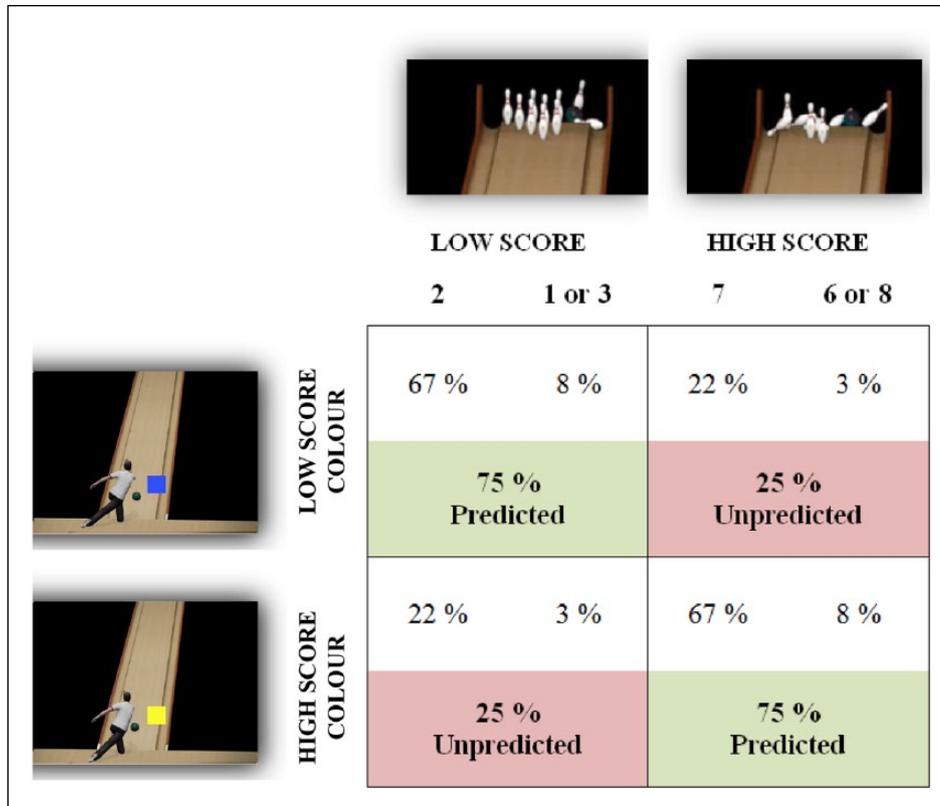


Figure 3. Overview of conditions and stimuli in Experiment 2.

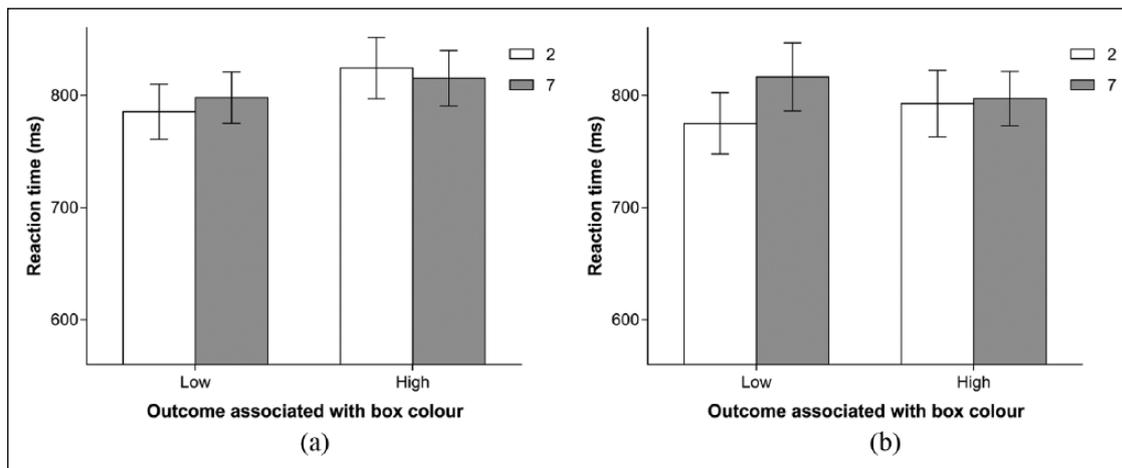


Figure 4. Reaction times (mean ± SEM) for (a) the colour question and (b) the outcome question in Experiment 2, separately for shirt-colour cue (indicative of low or high outcome) and outcome (scores 2 and 7).

between the colour and the outcome during the four practice trials. The questions that followed each movie were about the colour (“Was the box blue or yellow?”) and the outcome (“Was the score high or low?”). As in the first experiment, reaction times to the questions were measured using a button box and analysed using a 2 (colour: blue vs yellow) × 2 (outcome: 2 vs 7) × 2 (question: colour vs outcome) repeated measures ANOVA.

### Results and discussion

Unlike in the previous experiment, no significant three-way interaction between colour, outcome, and question was found,  $F(1, 27) = 0.27, p = .61, \eta_p^2 = .01$  (Figure 4; see Supplementary Material for a full overview of results). There was also no interaction between colour and outcome,  $F(1, 27) = 2.68, p = .11, \eta_p^2 = .09$ . The analysis only

showed an interaction between colour and question,  $F(1, 27)=4.33, p=.05, \eta_p^2=.14$ . To test whether this pattern of results differed from that in Experiment 1, an additional repeated-measures ANOVA was run with the same variables as before and experiment as a between-subjects factor. This analysis showed a significant four-way interaction between colour/agent, outcome, question, and experiment,  $F(1, 54)=4.89, p=.03, \eta_p^2=.08$ , confirming the difference between the experiments.

As in the previous experiment, additional analyses confirmed that including the questions following movies that showed another outcome (i.e., 1, 3, 6, and 8) did not influence this pattern of results. Also, there was no indication that the effects were driven by a speed-accuracy trade-off and the average accuracy over all trials was 96.6% (96.2% for the agent question and 97.0% for the outcome question).

The results do not provide evidence that participants use the correlation between the coloured box and the outcome to predict the outcome. Whereas in Experiment 1 the correlation between an agent's skill and an outcome had a causal interpretation (i.e., an experienced (or novice) player is more likely to cause a high (or low) score outcome), the correlation between colour and outcome in Experiment 2 did not have a natural causal interpretation (i.e., according to our world knowledge, colours in and of themselves have no causal powers to make pins fall down). The difference between these experiments is in line with the idea that predictions during action observation depend on the causal relation between the predictor and the outcome. However, before concluding that the processing of action outcomes indeed crucially depends on causal knowledge about agents causing outcomes, we need to rule out alternative explanations. In Experiment 1, participants were supposed to infer the agent's identity from his shirt colour to answer the agent question. In other words, to distinguish the novice and the experienced agent, participants could not just focus on directly observable cues, as they could for distinguishing the two colours in Experiment 2. To investigate whether this difference between the experiments may account for the findings, we conducted a third experiment in which the questions focused only on directly observable information from the movies.

## Experiment 3

### Methods

**Participants.** A total of 33 participants (31 females) aged between 19 and 28 (mean age 22.2) years took part in this experiment. The sample size was calculated based on the data from Experiment 1, with the assumption that the more implicit causal relation between the colour and the outcome would result in a smaller effect size (it was set at 50% of that of Experiment 1). One participant was excluded from the analyses because the pattern of accuracy for the agent question (i.e., 86.5% correct for the expected outcome versus 15.6% correct for unexpected

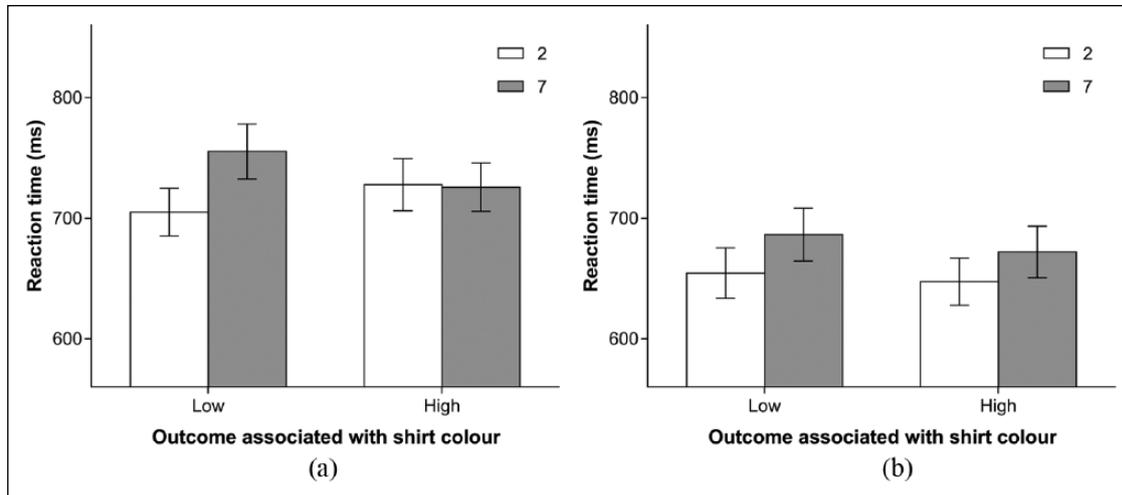
outcome) suggests that she misunderstood the instructions. As in the previous experiments, participants were paid €10 or received course credits for their participation. The study was approved by the local ethics committee and written informed consent was obtained from each participant.

**Stimuli and design.** This experiment was almost exactly the same as Experiment 1 in terms of stimuli and design. We used the same 24 animated movies and there were 288 trials, with a similar distribution as in the previous experiments.

**Procedure.** The procedure used in this experiment was similar to the procedure in Experiment 1. The only difference was in the instructions and the questions that were asked after each movie. As we did not want to instruct participants to think about two different agents with different skill levels, we only told them to pay attention to the colour of the agent's shirt without mentioning that there may be different agents. It was explained to them that the colour of the shirt would indicate whether the score was more likely to be high or low and, thus, that one colour would indicate that the score would probably be low and the other colour that the score would probably be high, although this would not be always the case. Again, the association between the colour and the outcome was mentioned explicitly in the practice trials, but not in the actual experiment. After each movie, participants answered a question about the shirt colour ("Was the shirt purple or white?") or about the outcome ("Was the score high or low?"). Reaction times to the questions were measured using a button box and analysed using a 2 (shirt colour: white vs purple)  $\times$  2 (outcome: low vs high)  $\times$  2 (question: shirt colour vs outcome) repeated measures ANOVA. Trials were excluded from the analysis in the same way as in the previous experiments. There was one trial in which the reaction time was below 100 ms, but as it was only one, it was not excluded from the analysis.

### Results and discussion

Results of the analysis show a three-way interaction between shirt colour, outcome, and question,  $F(1, 31)=4.21, p=.05, \eta_p^2=.12$ , as in Experiment 1 (Figure 5; see Supplementary Material for a full overview of results). Follow-up analyses showed that for the colour question, there was a significant interaction between colour and outcome,  $F(1, 31)=7.07, p=.01, \eta_p^2=.19$ . In paired-samples *t*-tests, unexpected events were found to result in higher reaction times for the colour associated with a low score,  $t(1, 31)=-4.28, p<.001$ , but not for the colour associated with a high score,  $t(1, 31)=0.17, p=.87$ . Overall, for the colour question, the reaction time to a question following an unexpected event was 26.1 ms longer than to a question following an expected event. For the outcome question, there was no interaction between colour and outcome,  $F(1, 31)=0.33, p=.57, \eta_p^2=.01$ . For this question, there



**Figure 5.** Reaction times (mean  $\pm$  SEM) for (a) the colour question and (b) the outcome question in Experiment 3, separately for shirt-colour cue (indicative of low or high outcome) and outcome (scores 2 and 7).

was a significant effect of outcome on reaction time,  $F(1, 31) = 12.29$ ,  $p < .01$ ,  $\eta_p^2 = .28$ . A score of 7 resulted in a reaction time that was on average 28.2 ms longer than the reaction time following a score of 2. This pattern of results resembles that of Experiment 1. A repeated measures ANOVA with the variables from the previous analysis as well as experiment (1 vs 3) as a between-subjects factor indeed showed no four-way interaction between agent/colour, outcome, question, and experiment,  $F(1, 58) = 0.67$ ,  $p = .42$ ,  $\eta_p^2 = .01$ .

Again, additional analyses confirmed that including the questions following movies that showed another outcome (i.e., 1, 3, 6, and 8) did not influence this pattern of results. As in the other experiments, there was no indication that the effects were driven by a speed-accuracy trade-off and the average accuracy over all trials was 96.5% (95.2% for the agent question and 97.7% for the outcome question).

So although participants were not explicitly instructed about different agents with different skill levels, the pattern of results for this experiment is very similar to that of the first experiment. This suggests that the difference between the first two experiments was not simply caused by the fact that one of the questions in Experiment 1 was about the agent, which was not directly observable, whereas in Experiment 2, both questions were about directly observable factors. In Experiment 3, like in Experiment 2, the questions focus on colour and outcome, both of which are directly observable. The pattern of results, however, resembles that of Experiment 1. This is in line with the idea that in both Experiments 1 and 3, colour is used as a cue to the agent's identity and the outcome is then predicted based on this identity. According to this idea, a causal relation between predictor (agent) and

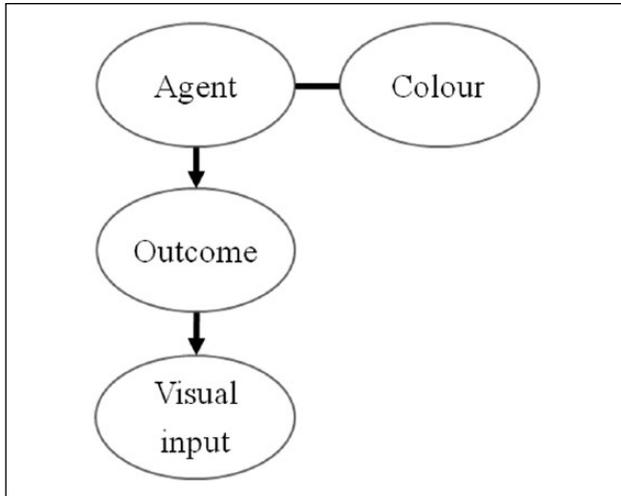
predicted (outcome) is crucial for predicting the outcomes of other people's actions.

To test the idea that predictions enhance processing in terms of speed, we ran an additional analysis in which we compared the overall reaction times for all three experiments, using a one-way ANOVA. This analysis indicated that the average reaction times differed significantly between the experiments,  $F(2, 85) = 5.16$ ,  $p < .01$ ,  $\eta_p^2 = .11$ . Bonferroni-corrected post hoc comparisons revealed that the average reaction time in Experiment 2 ( $M = 798.8$  ms, standard deviation [ $SD$ ] = 124.57) was significantly higher than in Experiment 3 ( $M = 693.81$  ms,  $SD = 108.87$ ),  $p < .01$ , whereas the average reaction times for Experiment 2 were higher than for Experiment 1 ( $M = 715.64$  ms,  $SD = 159.52$ ), although this difference was only marginally significant,  $p = .06$ . For Experiments 1 and 3, reaction times did not differ significantly,  $p = 1$ .

Although the questions answered in the three experiments were different, the length and difficulty of the questions in Experiment 2 cannot account for these findings. Rather, the differences in reaction times between the experiments seem to suggest that predictions that participants made in Experiments 1 and 3 allowed them to respond quickly, whereas the inability of participants to use the relation between colour and outcome to predict the outcome in Experiment 2 resulted in a longer reaction time. This is in line with the idea that predictions speed up cognitive processing.

## General discussion

In a series of experiments, we investigated whether a predictive processing model is, in principle, able to account for the way in which we process the outcomes of other people's actions. More specifically, we explored a model



**Figure 6.** A simplified version of our precise characterisation of hierarchical predictive processing.

in which predictions arise in a generative model that has a hierarchical structure and consists of causal relations between different levels of this hierarchy. The present results support the idea that such a model can account for this type of processing. To further improve the interpretation of the data, we developed a computational characterisation of this hierarchical predictive processing model to assess to what extent the present experiments' qualitative pattern of results is consistent with our theoretical assumptions. A detailed formal description of this characterisation and the associated processes can be found in the Supplementary Material. Here, we will briefly explain the characterisation (for a simplified version, see Figure 6), before outlining the relation with the experimental findings.

Crucially, the model consists of three levels (agent, outcome, and visual input), hierarchically ordered from top to bottom. Predictions are sent from higher to lower levels. This means that the visual input (i.e., the actual configuration of pins falling down) is ultimately predicted based on the observed shirt colour, through predictions at the outcome and agent levels. The relations between the three levels are causal in nature (indicated with arrows), whereas the relationship between colour and agent (indicated with a line) is deterministic. This deterministic relation indicates that the colour of the agent's shirt can be used to identify the agent. This identification then drives a prediction for the most probable outcome (two versus seven pins), which in turn drives predictions about which exact pins will fall down.

Bottom-up input at the lowest level (i.e., a visual stimulus) that is not fully predicted generates prediction errors. These prediction errors represent information about the input that was not already anticipated (Rao & Ballard, 1999) and are instrumental in updating the hypothesis (explaining away the prediction error) at a higher level. In

the predictive processing framework, these prediction errors are weighted, which means that top-down and bottom-up information is balanced by altering the gain on specific prediction error units (Clark, 2013a; Friston & Kiebel, 2009). The prediction error that arises at the level of visual input, and is processed at the outcome level above it, carries a very low weight because of the irreducible uncertainty in this information. This irreducible uncertainty stems from the fact that each outcome is inherently associated with many potential configurations of visual input (e.g., there are many ways in which 2 out of 10 pins can fall down). As a result of this low weight, little additional processing is going on at the outcome level even if some prediction errors arise at the visual level below. On the other hand, the prediction error that is sent to the agent level carries more weight, as the relation between agent and outcome has much more reducible uncertainty and thus allows for updating. Therefore, the discrepancy between the predicted and the perceived outcome (i.e., the number of pins hit) will lead to additional processing at the agent level to explain away the high prediction error. Presumably, this takes time, and as a consequence, when participants answer a question that calls for information from this level, the reaction time increases monotonically with the size of the prediction error at the level below. In this case, the increase in reaction time reflects "explaining away" of the prediction errors by updating the hypothesis about the current situation. More informally, we could say that when participants observe an experienced player hitting only a few pins, their prediction that the outcome is most probably high is violated, leading to a relatively high prediction error, but this prediction error can be explained away based on the knowledge that unpredicted outcomes occur every now and then. After the prediction error is explained away, the information represented at this level can be read and the question can be answered.

With this model in mind, we will now take a closer look at the experimental results. In Experiment 1, we found that prediction errors at the lower (outcome) level indeed slowed down reporting of information at the higher (agent) level. This is consistent with the idea that the generative model created to predict human actions is structured hierarchically, as we assume that prediction errors are sent upwards from lower levels to higher levels, increasing reaction times at this higher level as they are explained away. Given that only responses to the agent question and not the outcome question were influenced, a non-hierarchical representation of the generative model now seems less plausible.

In two follow-up experiments, we investigated the idea that, at high levels of the hierarchy where agent-caused events are represented, the predictive relations between levels of the generative model have a causal representation. In Experiment 1, participants could use their knowledge about the agent to predict the outcome of a bowling action.

A novice player was expected to obtain a low score, whereas an experienced player was expected to obtain a high score. The relation that is learned here could be represented as a causal one: based on general knowledge about the world, we know that agents cause outcomes. So when participants learned which agent usually got which scores in the bowling game, they could use this knowledge to base their prediction about the outcome on. In case this knowledge about agents and the causal structure of the world would not be important, one would expect that changing the agent information to another informative cue would result in the same pattern of reaction times. Therefore, in the second experiment, we used a cue with the same predictive probabilities as the agent. It has previously been suggested that predictions are based on simple associations and the contiguity and contingency between a cue and an outcome determine whether a prediction will be made (Keysers & Gazzola, 2014). Contiguity describes the paired occurrence of the cue and the outcome, whereas contingency implies that one event (i.e., the colour cue or the agent) reliably predicts the other (i.e., the outcome). Thus, the contiguity and contingency between the cue and the outcome were exactly the same in all three experiments. The second experiment differed from the other experiments in terms of the type of relation that could be learned. Here, this relation was not a causal one between an agent and an outcome, but a purely stochastic one between a colour and an outcome. In terms of the model presented in Figure 6, predictions about action outcomes cannot be based on a colour because this colour is not intrinsically linked to an agent, as the representation of colour does not have a causal relation with the representation of the outcome at the level below. This is in agreement with a suggestion that we have made elsewhere (Heil, van Pelt, Kwisthout, van Rooij, & Bekkering, 2014): the degree to which two events are associated can be guided by beliefs about the causal structure of the world. For example, people can only be conditioned to blink their eyes when they hear a tone in case they are aware of the relation between the tone and a puff of air to the eye that caused them to blink (Clark & Squire, 1998). In the same way, Waldmann (2000) found that the blocking effect in conditioning (i.e., the effect that the association of an event A with event Y is prevented if A is presented together with another event B that has previously been associated with event Y) is modulated by whether the participants were led to believe that A and B were either possible causes or possible effects of Y. It seems that awareness of a causal relation between two events influences the way in which these events are processed: if people do not interpret two events as causally related, they are not as easily associated, as is the case when an arbitrary colour correlates with a certain outcome in our bowling setting. We therefore suggest that the lack of a causal interpretation in Experiment 2 results in a pattern of reaction times that does not distinguish between predicted and unpredicted outcomes, even though

participants were made aware that there was a statistical relation between the colour of the box and the outcome.

In Experiment 3, participants could again base their predictions on a causal relation between an agent and an outcome, but this experiment was designed to control for some of the differences between the two previous experiments. We used the same stimuli as in the first experiment, but asked the participants to answer questions about directly observable features, as in the second experiment. Next to reporting the outcome, participants reported the colour of the agent's shirt. Similar to Experiment 1, reaction times to the colour question (which is comparable with the agent question) were again influenced by predictions about the outcome. These findings are in line with the idea that the brain's predictive model of other agent's actions consists of causal relationships between different levels in a hierarchy.

Previous studies investigating the role of causality in cognitive processing already showed that whether or not two events are perceived to be causally related depends on prior knowledge. This knowledge may, for example, concern the causal mechanism (Shultz et al., 1986), the probability that a causal relation exists or the assumed functional form (i.e., deterministic or probabilistic) of this relation (Griffiths, Sobel, Tenenbaum, & Gopnik, 2011). Importantly, this study extends these findings by showing that if events are perceived to be causally related, this allows for the prediction of one event based on the other, which then enhances cognitive processing. It seems that the world model that we use to predict other people's actions and their outcomes revolves around causes and their effects.

This does not mean that it is impossible to make predictions based on arbitrary cues. Previous studies actually showed prediction effects for arbitrarily associated events (e.g., Kok et al., 2012). Conceivably, in the clearly non-arbitrary setting that we created in our experiment, relevant world knowledge is activated, whereas this is not the case for studies in which arbitrary events are associated. For example, in the study by Kok et al. (2012), an auditory cue was the only source of information available to predict the orientation of the visual stimulus. The richness of information in our experiments may have made it more difficult to associate an arbitrary cue with an action outcome. This potential difference in the associations that are learned in arbitrary versus non-arbitrary settings demonstrates the importance of naturalistic, ecologically valid experimental designs in which a causal relation can be inferred. For example, when studying hierarchical predictive processing of agent-caused events, as we do here, it is important to take into account that the structure of the hierarchy may depend on the causal relation between agent and outcome. We suggest that knowing that agents cause outcomes allows us to predict an outcome based on knowledge about the agent.

Importantly, predictive processing is assumed to enhance processing in terms of speed. A comparison between the overall reaction times in the different experiments indeed revealed that in cases in which we found that participants were able to predict the outcomes (i.e., Experiments 1 and 3), processing was speeded up compared with cases in which we found no evidence for these predictions (i.e., Experiment 2). This is in line with the idea that predictions enhance processing in terms of speed and shows the potential benefit of predictive processing as a more general mechanism.

Although our results show that the predictive processing framework is able to account for reaction time effects in the processing of another person's action outcomes, it is also important to consider possible alternative explanations. For instance, in contrast to predictive processing, other accounts may not assign a key role to predictions and prediction errors. An example of such a non-predictive explanation would be an account in which probabilistic inference takes place only after the observation of the events. One could, for instance, hypothesise that the inference of a less probable event requires more processing time than inference of a more probable event, resulting in a longer reaction time for an agent that is not probable given the outcome. Such an account can be substantiated, however, only insofar as there is a plausible mechanism that explains why lower probability events take longer to process. Even if this is granted, such an account would need to make additional assumptions to explain the exact pattern of results. For example, without additional assumptions, the account cannot explain the difference between the two questions, as both need to be inferred and both deal with a similar pattern of low and high probability.

However, it is important to note that many models of cognition may actually be integrated in the predictive processing framework. For example, claims arising from associative theories, such as the associative sequence learning model, do not necessarily oppose those of the predictive processing framework. Press, Heyes, and Kilner (2011) argue that the associative sequence learning model explains how relations are learned, whereas predictive processing explains how learned relations support inferences about other people's actions. In this sense, the two models complement each other because they simply address different questions.

Future research is needed to investigate whether a predictive processing model with a hierarchical structure and causal relations between its levels can also account for different types of cognitive processing. Also, it would be interesting to further disentangle the different processes that may underlie explaining away of the prediction error. In case of hypothesis updating, the probability distribution over the candidate hypotheses is reassessed, whereas the generative model remains stable. On the other hand, revision of the generative model, sensory sampling to obtain more information on the unpredicted event, and active

inference have been proposed as alternative ways to bring prediction and sensory input closer together (Kwisthout, Bekkering, & van Rooij, 2017). For example, if the relation between a certain agent and a certain outcome changes in the course of an experiment (e.g., when the novice bowler gets better), participants need to revise their model.

### Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by a Netherlands Organisation for Scientific Research (NWO)-TOP grant (grant number 407-11-040) to H.B. and I.v.R. from the NWO.

### ORCID iD

Lieke Heil  <https://orcid.org/0000-0002-6375-6126>

### Supplementary Material

Supplementary material is available at [journals.sagepub.com/doi.suppl/10.1177/1747021817752102](https://journals.sagepub.com/doi.suppl/10.1177/1747021817752102)

### References

- Aglioti, S. M., Cesari, P., Romani, M., & Urgesi, C. (2008). Action anticipation and motor resonance in elite basketball players. *Nature Neuroscience*, *11*, 1109–1116.
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, *76*, 695–711.
- Bestmann, S., Harrison, L. M., Blankenburg, F., Mars, R. B., Haggard, P., Friston, K. J., & Rothwell, J. C. (2008). Influence of uncertainty and surprise on human corticospinal excitability during preparation for action. *Current Biology*, *18*, 775–780.
- Clark, A. (2013a). The many faces of precision. *Frontiers in Psychology*, *4*, 270.
- Clark, A. (2013b). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*, 181–204.
- Clark, R. E., & Squire, L. R. (1998). Classical conditioning and brain systems: The role of awareness. *Science*, *280*(5360), 77–81.
- Den Ouden, H. E., Daunizeau, J., Roiser, J., Friston, K. J., & Stephan, K. E. (2010). Striatal prediction error modulates cortical coupling. *Journal of Neuroscience*, *30*, 3210–3219.
- Den Ouden, H. E., Friston, K. J., Daw, N. D., McIntosh, A. R., & Stephan, K. E. (2009). A dual role for prediction error in associative learning. *Cerebral Cortex*, *19*, 1175–1185.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*, 815–836.

- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*, 127–138.
- Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *364*, 1211–1221.
- Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2011). Bayes and blickets: Effects of knowledge on causal induction in children and adults. *Cognitive Science*, *35*, 1407–1455.
- Heil, L., van Pelt, S., Kwisthout, J., van Rooij, I., & Bekkering, H. (2014). Higher-level processes in the formation and application of associations during action understanding. *Behavioral and Brain Sciences*, *37*, 202–203.
- Hohwy, J. (2013). *The predictive mind*. Oxford, UK: Oxford University Press.
- Kerzel, D., Bekkering, H., Wohlschläger, A., & Prinz, W. (2000). Launching the effect: Representations of causal movements are influenced by what they lead to. *The Quarterly Journal of Experimental Psychology*, *53*, 1163–1185.
- Keysers, C., & Gazzola, V. (2014). Hebbian learning and predictive mirror neurons for actions, sensations and emotions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*, 20130175.
- Kilner, J. M., Friston, K. J., & Frith, C. D. (2007). Predictive coding: An account of the mirror neuron system. *Cognitive Processing*, *8*, 159–166.
- Kok, P., Jehee, J. F., & de Lange, F. P. (2012). Less is more: Expectation sharpens representations in the primary visual cortex. *Neuron*, *75*, 265–270.
- Koster-Hale, J., & Saxe, R. (2013). Theory of mind: A neural prediction problem. *Neuron*, *79*, 836–848.
- Kwisthout, J., Bekkering, H., & van Rooij, I. (2017). To be precise, the details don't matter: On predictive processing, precision, and level of detail of predictions. *Brain and Cognition*, *112*, 84–91.
- Michotte, A. (1963). *The perception of causality* (T. R. Miles & E. Miles, Trans.). London, England: Methuen Publishing.
- O'Reilly, J. X., Schüffelgen, U., Cuell, S. F., Behrens, T. E., Mars, R. B., & Rushworth, M. F. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, *110*, E3660–E3669.
- Phillips, H. N., Blenkmann, A., Hughes, L. E., Bekinschtein, T. A., & Rowe, J. B. (2015). Hierarchical organization of frontotemporal networks for the prediction of stimuli across multiple dimensions. *Journal of Neuroscience*, *35*, 9255–9264.
- Press, C., Heyes, C., & Kilner, J. M. (2011). Learning to understand others' actions. *Biology Letters*, *7*, 457–460.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extraclassical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87.
- Saxe, R., Tzelnic, T., & Carey, S. (2007). Knowing who dunnit: Infants identify the causal agent in an unseen causal interaction. *Developmental Psychology*, *43*, 149–158.
- Shultz, T. R., Fisher, G. W., Pratt, C. C., & Rulf, S. (1986). Selection of causal rules. *Child Development*, *57*, 143–152.
- Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*, 745–756.
- Summerfield, C., & Koechlin, E. (2008). A neural representation of prior information during perceptual inference. *Neuron*, *59*, 336–347.
- Van Pelt, S., Heil, L., Kwisthout, J., Ondobaka, S., van Rooij, I., & Bekkering, H. (2016). Beta- and gamma-band activity reflect predictive coding in the processing of causal events. *Social Cognitive and Affective Neuroscience*, *11*, 973–980.
- Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*(1), 53–76.