

Spontaneous Causal Learning While Controlling A Dynamic System

York Hagmayer^{1,*}, Björn Meder², Magda Osman³, Stefan Mangold¹ and David Lagnado⁴

¹Department of Psychology, University of Göttingen, Göttingen, Germany

²Max Planck Institute for Human Development, Berlin, Germany

³School of Biological and Chemical Sciences, Queen Mary University of London, UK

⁴Division of Psychology and Language Sciences, University College London, UK

Abstract: When dealing with a dynamic causal system people may employ a variety of different strategies. One of these strategies is causal learning, that is, learning about the causal structure and parameters of the system acted upon. In two experiments we examined whether people spontaneously induce a causal model when learning to control the state of an outcome value in a dynamic causal system. After the control task, we modified the causal structure of the environment and assessed decision makers' sensitivity to this manipulation. While purely instrumental knowledge does not support inferences given the new modified structure, causal knowledge does. The results showed that most participants learned the structure of the underlying causal system. However, participants acquired surprisingly little knowledge of the system's parameters when the causal processes that governed the system were not perceptually separated (Experiment 1). Knowledge improved considerably once processes were separated and feedback was made more transparent (Experiment 2). These findings indicate that even without instruction, causal learning is a favored strategy for interacting with and controlling a dynamic causal system.

Keywords: Decision making, dynamic systems, causal models, learning, induction.

INTRODUCTION

We encounter complex and dynamic causal systems in both our professional and everyday lives. Indeed we need look no further than ourselves as paradigm examples of such systems. We are constituted by biochemical and neurological mechanisms that are interconnected in a complex manner and operate to maintain a delicate equilibrium. Exercise, diet, and medications are methods by which we affect our own physiological system in order to regulate it, as well as improve it. Likewise, in cases where we interact with the external environment, we often strive to control dynamic systems. For example, workers in industrial companies need to learn to control highly complex systems such as chemical plants and production lines. On a broader scale, politicians, managers, and stockbrokers act on large social systems (e.g., nations, corporations, stock markets), where social agents causally interact in various ways. People generally engage in these systems by making goal directed interventions in order to achieve their desired outcomes. Marketing campaigns, taxation, welfare, and social norms are ways of influencing people's behavior to improve the quality of life and achieve happiness.

In this paper, we are concerned with the kind of knowledge that decision makers acquire when they intervene on and control such complex systems. In particular, we

investigate whether decision makers *spontaneously* acquire causal models, that is, representations that mirror the causal mechanisms that govern the dynamic system acted upon. In other words, we are interested in whether decision makers engage in causal induction during dynamic control tasks, even when they are not asked to do so. First, we will briefly review the literature on control tasks in relation to findings about causal learning. The existing body of evidence on control tasks will be contrasted with research on causal learning, which typically does not examine learning with dynamic systems. Then we will present two experiments that examine whether people spontaneously induce causal models when learning to control a dynamic causal system. We close by discussing the empirical findings and their relevance for the common claim that decision makers fail to acquire causal knowledge in control tasks.

CAUSAL LEARNING WITH DYNAMIC SYSTEMS

Previous research on people's ability to control complex dynamic systems (e.g., simulated industrial power plants, management systems, ecosystems) has largely neglected the role of causal knowledge. This is because theoretical accounts of skill acquisition have focused on explaining the apparent dissociation between accurate control performance and poor reportable knowledge of the structure of the system [1, 2]. Participants' poor reported knowledge is taken as evidence that very little knowledge about the underlying causal structure is acquired. This dissociation between control performance and explicit knowledge is explained in terms of the complexity of the environment. Rather than learning the un-

*Address correspondence to this author at the Department of Psychology, University of Göttingen, Gosslerstrasse 14, 37073 Göttingen, Germany; Tel: +49 (0) 551-398293; Fax: +49 (0) 551-393656; E-mail: yhagmay@gwdg.de

derlying structure, people are supposed to learn to control the system by strengthening their knowledge of the perceptual features of the system and the specific decisions they make (i.e., what actions they chose in a given state of the system) along with the successful outcomes produced [3, 4]. By adding more cases to their knowledge base, they proceed by matching the perceived current state of the system to stored exemplars of successful situation-decision-outcome associations. Because this process relies on generating responses from memory, no inferences are made about the underlying causal relations that govern the system's behavior. While this type of exemplar-based learning is good for controlling the system to a target, it occurs at the expense of gaining knowledge of the system's structural properties. Consequently, people show limited success in transferring control skills to other goals because their knowledge is bound to tasks that have the same perceptual and goal characteristics as the original training task [5, 6].

This view contrasts sharply with current research in causal learning and reasoning. Research on these topics shows that people are able to learn about causal structure through various cues to causality, such as temporal sequence, statistical relations, consequences of interventions, and prior knowledge [7, 8]. In addition, people are not only able to learn about the structure, but also the parameters of causal models [9-11]. Thus, it has been shown that beliefs about causal structure guide the interpretation of covariational data and strongly affect the way people structure the available learning input (for a recent overview see [12]).

Recent research has also emphasized the tight connection between causal beliefs and decision making. The *causal model theory of choice* [13, 14] assumes that people use the available information to induce a causal model of the decision problem and the choice situation. A causal model of the decision problem encompasses the decision maker's knowledge about the structure of the system and their beliefs regarding the causal influences of the available courses of actions. The empirical evidence shows that people indeed use such causal model representations when making simple one-shot decisions [13, 14]. Along similar lines, Hagmayer and Meder [15, 16] investigated whether people spontaneously induce causal models when repeatedly making decisions in order to maximize payoff variables. Their results showed that many participants induced causal models and used them to make decisions. These findings demonstrate that people do not solely learn causal models when asked to do so, but engage in causal induction even when trying to achieve a very different goal. However, in these experiments non-dynamic causal systems were used (i.e., systems that did not change when participants chose not to act). Furthermore, all variables within the system were reset to their initial value on every trial. In control tasks, by contrast, the system usually has some internal momentum and the goal of the decision maker is not to maximize, but to achieve and maintain a certain equilibrium state.

Why Causal Models?

Research on control tasks shows that people can successfully learn to control complex systems without acquiring explicit knowledge about causal structure. One might therefore wonder why decision makers should bother with causal

learning at all. The answer is that causal knowledge is more flexible and adaptive than other forms of knowledge. For example, causal model representations enable decision makers to infer the consequences of novel interventions or to evaluate changes in the structure of a causal system [8, 15, 16]. Crucially, these inferences can be derived from the structure and parameters of the causal model without the necessity of additional learning input. This is a clear advantage over non-causal representations, which require further learning whenever the causal system or the available options change. Causal knowledge is also critical when there are interferences within a dynamic system, such as when an ecosystem is on the edge of collapsing or when an industrial system such as a production line breaks down.

Consider the simple dynamic causal system depicted in Fig. (1a) and assume that the decision maker's task is to maintain the value of the final outcome variable at a certain target level. In total, there are eight causal relations within this causal system (indicated by the arrows in Fig. 1a). First, the decision maker has three options to intervene on the system (termed 'do Alpha', 'do Beta' and 'do Gamma'), whose effects are not known prior to learning. These three actions affect three intermediate variables (*A*, *B*, and *C*) which differentially contribute to the value of the final outcome: variable *A* increases the level by 80 points, *B* by 80 points, and *C* by 120 points. Crucially, outcome variable *A* not only contributes to the value of the final outcome variable, but also activates the intermediate variable *B*. As a consequence, choosing 'do Alpha' raises the level of the outcome variable by 160 points. In addition, the system is dynamic because a decay function reduces the final outcome value by 50% from trial to trial, regardless of the action taken by the decision maker. Thus, even when one does not act upon the system, the value of the target variable changes over time. Formally, the value of the final outcome variable is determined by the following equation:

$$Value_{t_1} = Value_{t_0} \cdot decay + Points | do(X)_{t_0} \quad (1)$$

For example, when the current value of the target variable is 120 and option 'do Alpha' is chosen, the resulting value is $120 \cdot .5 + 160 = 220$ points. Note that Equation 1 can be derived from the causal model depicted in Fig. (1a), but not vice versa. For example, the equation states how the value of the outcome is influenced by a particular intervention, but is silent about the intermediate causal processes (e.g., that 'do Alpha' generates variable *A*, which in turn generates variable *B*).

Assume that the decision maker's task is to maintain the final outcome variable at a target value of 160 points by repeatedly intervening on the causal system. An exemplar-based approach supposes that participants encode which value of the target variable results from a specific action taken in a particular state of the system. Thus the experienced feedback is used to incrementally build up a knowledge base. To decide which action to take next, the current state of the system is matched to previously seen cases, and successful actions are retrieved from memory and repeated. For example, a decision maker may have learned that when the state of the target variable is very low it is useful to choose 'do Alpha', since this action entails the largest increase. Conversely, when the value of the target variable is

than psychological factors may explain the dissociation [18], and many have reported findings in which there is a close correspondence between structural knowledge and control performance [18, 19]. For example, by regularly probing knowledge of structural relations during learning, people show insight into the relations that govern the behavior of the system. In addition, they use this knowledge to control the system in subsequent phases. However, while people may have knowledge of the structure of the relations, the parameters of these relations tend to be only poorly represented [6, 20]. People may know whether a relation is positive or negative, but typically they have very little knowledge about the functional form or the strength of the relation. Crucially, research in this area has yet to examine whether these relations are causal representations of the system, and if so, how they are used to inform decisions to control an outcome to a specified criterion. In particular, if the information available to participants (instructions, feedback resulting from interventions, and background knowledge) is insufficient for inferring a causal model of the underlying system, it is not surprising that participants refrained from doing so. Whenever the underlying causal system is dynamic and consists of probabilistic non-linear mechanisms connecting partially unobservable variables, it might simply be impossible to pin down the structure and parameters of the system because the causal model is underdetermined by the available data. In this case an exemplar-based strategy seems a much better approach, because it can easily adapt to any underlying causal system, as it only represents which action is successful under specific circumstances. However, this does not mean that participants are necessarily reluctant to induce causal models in control tasks. We speculate that participants are inclined to infer the causal mechanisms of a system, because causal knowledge enables them to predict the consequences that would result from a change in the causal system. Thus, decision makers who learn to control a dynamic system should induce a causal model whenever the feedback from the system allows them to do so. We pursued this hypothesis in two experiments.

EXPERIMENT 1

The goal of Experiment 1 was to investigate whether participants engage in spontaneous causal learning when learning to control a dynamic system. To examine the knowledge participants acquired during the control task, we employed a number of tests. In particular, we presented participants with a change in the causal system's structure. Since different kinds of representations of the system entail different reactions to such a change, this modification taps into participants' knowledge. While purely instrumental knowledge only allows for very limited inferences about the changes in effectiveness of previously used interventions, causal knowledge enables the prediction of the consequences resulting from the modification.

Method

Participants and Design

63 students from the University of Göttingen ($n = 31$, 28 females, 3 males mean age = 22.8) and various universities in Berlin ($n = 32$, 18 females, 14 males, mean age = 25.5) participated. Participants from Berlin were paid a small

amount of money (5€); participants in Göttingen could choose between being paid and receiving course credit. Participants were assigned to one of two counterbalanced conditions that only differed in the placement of one particular test question. In one condition, they were immediately queried about the underlying causal structure after the learning phases. In the other condition they were asked about causal structure after completing all other test questions. 27 participants were assigned to the condition in which they were immediately queried about causal structure after the learning phase whereas 36 participants were asked about causal structure after completing all other tasks (the unequal distribution resulted from a miscommunication between the two research locations).

Materials and Procedure

We used a computer-based biological scenario in which participants were instructed to control the level of a certain neurotransmitter in the brain of mice. In order to do so, participants could stimulate the mice's brain with three different types of rays (labeled Alpha-, Beta-, and Gamma-rays). Participants were further told that the radiation might activate different brain areas, which in turn were responsible for the production of a certain amount of the transmitter (cf. Fig. 2). It was also pointed out that the brain areas could be interconnected. No further information was given about the relations between the different types of radiation and brain activation, or between the different brain areas and levels of neurotransmitter.

Next participants were informed that the specific level of the neurotransmitter was essential to the survival of the mice. There were two different kinds of mice, each of which needed a different neurotransmitter level. One of the mice needed a transmitter level of 140 ± 20 points, while the other mouse needed a level of 280 ± 20 . Then participants were asked to maintain the neurotransmitter level on a number of subsequent trials. To maintain the transmitter level, they could stimulate the brain with Alpha-, Beta- or Gamma-rays on each trial (henceforth denoted as 'do Alpha', 'do Beta', and 'do Gamma'). They also had the option to not stimulate the brain ('do nothing').

Fig. (1a) shows the causal system underlying the control task. The level of the neurotransmitter had a decay of 50% on each trial (i.e., $value_{t1} = value_{t0} \cdot .5$), regardless of the action taken. Alpha radiation activated one of the brain areas (area A), which in turn activated a second area (area B). The joint activation of these two areas resulted in a total increase of the transmitter level of 160 points. Beta radiation activated area B and resulted in an increase of 80 points. Finally, Gamma radiation activated the third brain area (area C), thereby increasing the transmitter level by 120 points. (The causal relations from radiation to brain areas on the one hand and brain areas and amount of produced transmitter were counterbalanced across participants.) Participants were never told the causal model's parameters or the functional characteristics of the system. Thus, the underlying causal relations had to be inferred on the basis of the consequences resulting from the interventions. The radiation as well as the subsequent activation of the brain areas were presented through animations (cf. Fig. 2). First, participants saw which brain areas became activated due to the chosen intervention. If

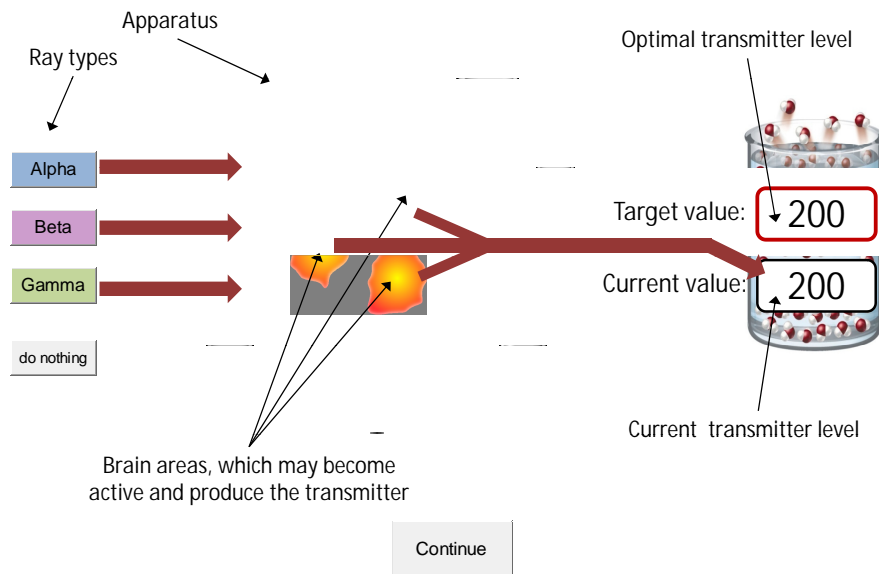


Fig. (2). Figure of the causal system shown to participants.

they chose intervention ‘do Alpha’ they first saw the activation of area *A*, followed by the activation of area *B* (with a delay of 1 s). When they chose ‘do Beta’ or ‘do Gamma’, the corresponding brain areas (*B* and *C*, respectively) became active. Straight afterwards they observed how the transmitter level changed from the previous level to the new level (due to the combined influence of the decay and the intervention). The new level was computed in accordance with the underlying causal model’s parameters (cf. Fig. 1).

Before being presented with the control task there was an introductory phase in which participants could not intervene on the causal system. Rather, they passively observed how the transmitter level decreased when the ‘do nothing’ option was chosen on four consecutive trials. This learning phase was added to ensure that all participants observed the decay of 50% before learning to control the system. Note that this is a necessary requirement to infer the causal model’s parameters. Without knowing the value of the decay, the underlying function cannot be solved as it would contain two unknowns (i.e., the parameter of the decay and the impact of the intervention).

The decay learning phase was followed by two consecutive control task phases in which participants had to maintain a certain value of the transmitter. Each of these phases comprised 20 trials. The desired target levels in the two control task phases were 140 and 280, respectively. The different target levels were chosen to ensure that participants would have to select all of the available options during the control task. For example, once a participant had reached a level of 140, it was best for her to always choose ‘do Beta’, as this would result in a constant level of 160, which lies within the acceptable range of ± 20 . Conversely, to maintain a level of 280 the best strategy was to alternate between ‘do Alpha’ and ‘do Gamma’.

The order of the two control tasks was randomized across participants. On each trial participants could choose among the four options (‘do Alpha’, ‘do Beta’, ‘do Gamma’, and ‘do nothing’). When the task was to maintain a level of 140, the initial starting value was 100 points. When the target level was 280, the initial starting value was 200. On each

trial the consequences of the intervention were computed and feedback was provided as described above (i.e., which brain areas became activated and the resulting value of the transmitter level).

After the 40 control task trials participants received three different tests designed to tap into their knowledge about the causal system: a causal model selection test, two choice tasks, and two predicted value tests. The first test aimed to directly elicit participants’ reportable knowledge of the underlying causal structure. Participants were presented with a graphical representation of two different causal models. Subjects had to choose between a causal model in which the three brain areas were independent from each other (*Independent Causes Model*) and a model in which an activation of area *A* caused an activation of area *B* (*Causal Chain Model*). For both models all three brain areas were designated as causes of the transmitter level. If participants used the experienced feedback from the activation of the brain areas, they should infer that the causal structure corresponds to a causal chain model containing a link $A \rightarrow B$. Since we speculated that this test might have an impact on the other tests, we counterbalanced the position of this question. In one condition participants were asked this question directly after the two control phases, whereas participants in a second condition received this question subsequent to all other tests.

The next two tests aimed to examine how well participants had learned to control the system. For the *choice test* (Test 2a) participants were asked to choose interventions for four different combinations of start values (120 vs. 160 points) and target values (140 vs. 280 points). For example, participants were told that the current value of the transmitter was 120 points and the target value was 280. Then they were asked to choose the intervention that would get them as closely as possible to the desired target value. To ensure that these four decisions were only based on the previously acquired knowledge, no feedback was provided in this phase. The start/target-combinations were designed in a way that there was a clearcut best answer given that participants had learned about the structure and parameters of the causal system. The order of these four questions was randomized.

After the choice test participants were given the *predicted value test* (Test 2b). For this test, participants had to estimate the transmitter level resulting from each of the four options ('do Alpha', 'do Beta', 'do Gamma', and 'do nothing') for two starting values (140 and 280 points, respectively). Hence subjects made eight estimates. Again no feedback was provided. This test was administered to assess people's estimates for the causal impact of each available intervention.

The next tests were designed to investigate whether and how participants would react to changes in the causal underpinnings of the decision problem. Participants were presented with new mice, and were instructed that these were identical to the ones they had encountered during the previous control task, except that the brain area *B* had been surgically removed and could no longer produce the transmitter (cf. Fig. 1b). The removal of variable *B* from the causal system has a number of implications for the effectiveness of the different interventions. As Beta radiation only activates area *B*, this intervention can no longer cause an increase in the transmitter level. Moreover, due to the causal link $A \rightarrow B$, the causal impact of Alpha radiation would also be affected. During the initial control task, Alpha radiation caused an activation of brain area *A*, which in turn activated area *B*. Together these two events generated an increase of +160 points. Due to the removal of area *B*, which produced 80 points, an activation of area *A* with Alpha radiation would now result in an increase of 80 points. This in turn makes Gamma radiation the most effective intervention (+120 points). The crucial point is that the implications of removing variable *B* from the causal system can only be inferred from a causal model representation of the decision context. Thus, if participants engaged in spontaneous causal induction during the control task, they could now capitalize on their causal knowledge to assess the implications of the system's structural modification.

To examine participants' sensitivity to the structural modification of the causal system we employed the same two tasks as before. The crucial difference was that this time the decisions and estimates referred to the modified causal system (i.e., the mice without brain area *B*). Again participants were presented with four different combinations of start values (120 and 160) and target values (140 vs. 280) and were asked to choose the intervention that would get them as close as possible to the desired target value (Test 3a). Finally, we administered the same predicted value test as before, in which learners had to estimate the resulting transmitter value for each possible intervention (Test 3b). Again two different start values (140 and 280, respectively) were used.

Based on our hypothesis that decision makers would use the feedback experienced while learning to control the dynamic system to infer a parameterized causal model of the system, we expected them (i) to prefer the causal chain model over the independent causes model, (ii) change their intervention choices from Test 2a to Test 3a, and (iii) adjust their predicted values for Alpha and Beta radiation subsequent to the removal of variable *B* (Test 2b vs. Test 3b).

Results

Control Task

In order to learn about the structure and the parameters of the causal system it is necessary to experience the consequences of all available interventions. Therefore we first checked which choices were made in the initial control task. Fig. (3a) depicts participants' choices. All participants chose every option at least once; on average each option was chosen roughly 10 times. Thus, in principle participants had sufficient information to learn the structure and parameters of the underlying causal system. For example, imagine that the current value is 200 and a participant chose 'do Gamma'. The outcome of this intervention would be an activation of




Fig. (3). (a) Number of chosen interventions during learning in Experiment 1 (means and SE), (b) distance of actual transmitter level to target level (means and SE). The two blocks (trial 1-20 and trial 20-40) denote the two learning phases with different target values (140 and 280). At the beginning of each block, the target and starting value of the transmitter were reset.

brain area C and the resulting value of the transmitter would be $(200 \cdot .5) + C = 220$. Since the value of the decay was known from the initial observation phase there is only one unknown in the equation. Thus, the equation can be solved for C to determine how much transmitter is produced by brain area C (i.e., $C = 220 - (200 \cdot .5) = 120$). To check whether participants learned to maintain the target level of the transmitter, we computed the distance from the target for each trial. Fig. (3b) depicts the resulting learning curve averaged over participants. It shows that participants learned to control the system, although there was little transfer between the two target levels (see the sharp increase in distance on trial 21 in Fig. 3b). A within-subjects ANOVA comparing the average distance of the first ten trials to the average distance of the last ten trials yielded a clear effect, $F(1,62) = 10.5$, $p < .01$, $MSE = 277.4$. Taken together, these results suggest that participants successfully learned to maintain the target value in the control task.

Causal Model Selection Task

As noted previously, we expected that participants would correctly identify the causal structure of the system. This was in fact the case: 84% of participants correctly chose the causal chain model over the independent causes model, which is significantly higher than 50%, $\chi^2(df = 1, N = 63) = 29.3$, $p < .01$. This finding indicates that participants were sensitive to the structure of the causal system underlying the control task. Also, participants' model choices were not affected by the measurement point, $\chi^2(df = 1, N = 63) = .04$, $p > .80$. We therefore pooled the data for the subsequent analyses.

Intervention Choices

Next we analyzed decision makers intervention choices in Test Phases 2a and 3a, in which they were asked to make four intervention decisions for the original system (comprising all three variable A , B , and C) and the modified causal system in which variable B was removed. Fig. (4) depicts the results.¹ On the left hand side choices are depicted for the causal system participants acted upon in the control task. If participants correctly inferred the consequences of the different interventions, they should prefer intervention 'do Beta' for target level 140 since $(current\ level \cdot decay) + Points/do\ Beta = (120 \cdot .5) + 80 = 140$ and $(160 \cdot .5) + 80 = 160$. For target level 280 'do Alpha' is the best option because the other options result in less points: $(current\ level \cdot decay) + Points/do\ Alpha = (120 \cdot .5) + 160 = 220$ and $(160 \cdot .5) + 160 = 240$. Although not perfect, the most frequent choices conformed to these predictions. On the right hand side of Fig. (4) choices are depicted for the modified causal system in which area B (the area normally activated by Beta radiation) is removed. Based on the modified causal model, participants should now have preferred intervention 'do Alpha' for target level 140 because $Points/do\ Beta_{B_removed} = 0$ and $Points/do\ Alpha_{B_removed} = 80$. For target level 280 'do Gamma' should now be chosen because $Points/do\ Gam-$

$ma_{B_removed} = 120$, that is 'do Gamma' now has a higher impact than 'do Alpha'. Thus, participants should exhibit a differential choice pattern for the original and the modified causal system. In contrast to this prediction, Fig. (4) indicates that participants tended to stick with the choices they made for the system they initially acted upon. To test whether participants changed their preference we conducted a Bowker test for choices with respect to a target level of 140. A Bowker test investigates whether frequencies of different categories (i.e., chosen options) stay the same for two measurements. The test confirmed that a number of participants systematically switched away from choosing 'do Beta', $\chi^2(df = 1, N=126) = 22.1$, $p < .001$. However, not all of them turned to 'do Alpha' as we expected; a substantial number now preferred 'do Gamma'. For a target level of 280 a Bowker test could not be run due to low cell entries. Instead we conducted a McNemar test focusing on the 'do Alpha' and the 'do Gamma' choices. In line with the predictions derived above, participants switched away from choosing 'do Alpha' and now preferred more often 'do Gamma', but still most participants continued with 'do Alpha', $\chi^2(df = 1, N=126) = 17.3$, $p < .001$. Thus, it does not seem as if decision makers used a correctly parameterized causal model to make their choices.

Predicted Values

In Test Phases 2b and 3b participants were asked to estimate the transmitter values resulting from all possible interventions (including the option of 'do nothing'). Estimates were requested for two starting values (140 vs. 280) for both the original and modified causal system. Table 1 depicts the results for the 16 estimates as well as the values entailed by the underlying causal model. For the system participants acted upon in the control phase, they generally underestimated the transmitter level resulting from the interventions. However, the rank order of the estimates corresponds to the order entailed by the underlying causal model.

The crucial analyses concern the comparison of 'do Alpha' and 'do Beta' before and after removing variable B from the causal system. A causal model analysis entails that the ratings for 'do Alpha' and 'do Beta' should decline. The qualitative pattern of estimates conformed to these predictions. Multiple planned comparisons revealed significant differences for 'do Alpha' ratings before and after the modification for both starting values of 140 and 280 (cf. Table 1). The difference for 'do Beta' ratings before and after the modification was only significant for a starting value of 140. The difference for starting value 280 ($p = .055$) missed our error type I corrected criterion of significance. However, the obtained mean estimates also indicate that participants did not decrease their estimates as much as entailed by the underlying causal model. Since the removal of area B completely wipes out the causal impact of 'do Beta' on the transmitter level the predicted values of 'do Beta' should equal the ratings for 'do nothing'. However, only 16 out of 63 participants made this inference at a current value of 140, and only 20 participants at a current value of 280. In line with our predictions the small differences obtained for the estimates of 'do Gamma' and 'do nothing' failed to reach significance for both starting values. Overall, the findings indicate some evidence for participants' sensitivity to the causal model modification, but the results also show that

¹We ran tests to check whether asking participants first about causal structure affected their choices. In order to avoid low cell entries we classified participants' choices as predicted vs. not-predicted for all eight test cases. Then we conducted χ^2 -test for each test case separately comparing the choices in the two counterbalancing conditions. Only one significant difference resulted. Participants more often chose do Alpha after the modification given a starting value of 120 and a target of 140 when they were first asked about causal structure ($\chi^2 = 8.82$, $p < .01$). As this was the only difference, we report the choices for both conditions together in Fig. (4).




Fig. (4). Percentage of participants choosing a certain intervention in the test phase before the causal system was modified and after the modification in Experiment 1. Numbers on the x-axis indicate the specific test case, the first number is the target value (e.g. 140); the second number is the starting value (e.g. 120). Choices predicted on the basis of the correct causal model are indicated by dashed borders.

Table 1. Mean Estimates of Effects Resulting from Interventions in Experiment 1. Numbers in Italics are Predictions Derived from a Correctly Parameterized Causal Model, Numbers in Brackets are Standard Errors

| | Target & Start =140 | | Target & Start = 280 | |
|------------|--------------------------------|------------|-----------------------------|------------|
| | Before | After | Before | After |
| do Alpha | <i>230</i> | <i>150</i> | <i>300</i> | <i>220</i> |
| | 203.95 | 184.41 | 325.10 | 308.92 |
| | (3.35) | (3.92) | (5.22) | (6.66) |
| | $F(1,61) = 15.8; p <.001^{**}$ | | $F(1,61) = 9.59; p <.003^*$ | |
| do Beta | <i>150</i> | <i>70</i> | <i>220</i> | <i>140</i> |
| | 145.03 | 115.79 | 243.56 | 224.78 |
| | (3.65) | (4.94) | (4.06) | (7.78) |
| | $F(1,61) = 20.3; p <.001^{**}$ | | $F(1,61) = 3.82; p =.055$ | |
| do Gamma | <i>190</i> | <i>190</i> | <i>260</i> | <i>260</i> |
| | 161.40 | 155.76 | 268.54 | 270.60 |
| | (3.58) | (3.64) | (4.11) | (3.71) |
| | $F(1,61) = 3.66; p =.060$ | | $F(1,61) < 1; p =.49$ | |
| do Nothing | <i>70</i> | <i>70</i> | <i>140</i> | <i>140</i> |
| | 76.57 | 81.75 | 157.03 | 165.97 |
| | (1.70) | (3.25) | (4.34) | (5.85) |
| | $F(1,61) = 3.22; p =.078$ | | $F(1,61) = 3.24; p =.072$ | |

Note. For all interventions and target values separate ANOVAs were calculated with counterbalancing condition (causal structure test first vs. causal structure test last) as a between-participants variable and modification (before vs. after) as a within-participants variable. To correct for accumulation of type 1 error the significance level was set to $\alpha = .00645$. It turned out that neither the counterbalancing condition nor the interaction of the two factors ever reached significance (all $F_s < 2.1, p > .15$). Therefore only the main effects of the modification of the causal system are reported in Table 1. Results marked with an asterisk are significant.

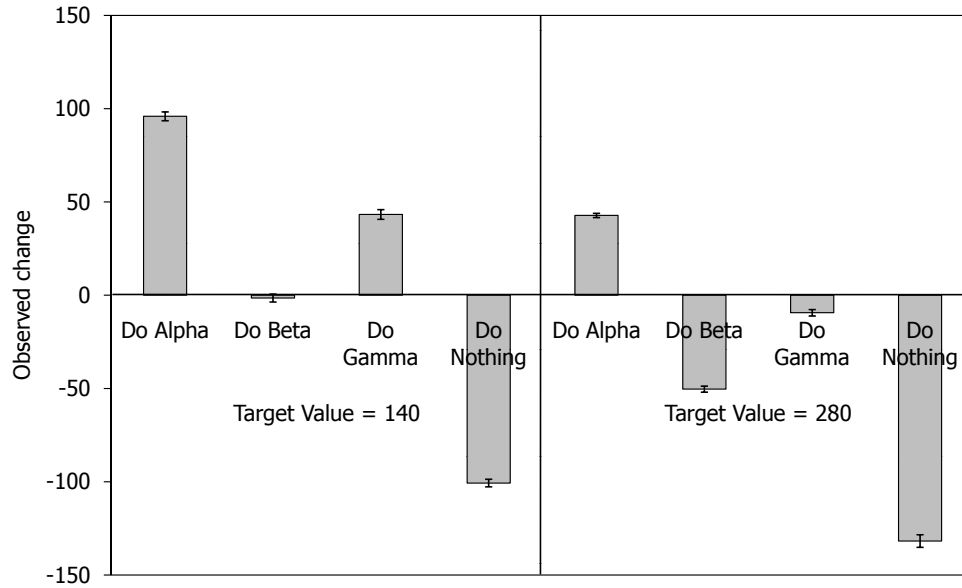


Fig. (5). Observed changes in transmitter level in Experiment 1 (means and SE). Change is the difference of the transmitter value resulting from the intervention and its starting value prior to the intervention.

decision makers did not fully grasp the implications of removing variable *B* from the causal system.

Modeling Decision Making Strategies

The findings suggest that participants did not successfully learn about the causal model and its parameters in the initial control task. For a more detailed analysis we modeled alternative decision making strategies and analyzed how well participants' judgments conformed to these strategies. The first strategy is representative for models assuming that decision makers learn the relations among options and outcome value. The second strategy is exemplar-based. The other three strategies model different types of causal model approaches.

Mean Change Heuristic

The first strategy is inspired by the *natural means heuristic* proposed by Hertwig and Pleskac [21] as a strategy for making repeated decisions under uncertainty. The natural means heuristic assumes that participants cope with uncertainty by encoding the outcomes resulting from the different options. This strategy is sensitive to the predicted value of the different options, because outcomes depend both on the probability and the value of the outcome. In a dynamic system, however, the actual outcome also depends on the initial state of the system. Therefore the assumption underlying the relative change heuristic in the context of a control task is that participants do not learn the absolute impact of the different interventions, but rather encode the *relative* changes resulting from their actions.² In particular, this model assumes that participants do not disentangle the influence of the interventions from the decay, but merely encode the average change resulting from a chosen option. For example, if the current value of the transmitter is 200 and a participant

chose 'do Alpha', the resulting value of the transmitter would be $(200 \cdot .5) + \text{Points} \mid \text{do A} = (200 \cdot .5) + 160 = 260$. Thus, the *relative* change resulting from the chosen action would be +60 points. In summary, the relative change heuristic is agnostic about causal structure and only focuses on the relative impact of the interventions on the state of the outcome variable (i.e., the transmitter value).

To model this strategy, we computed the relative changes participants observed for each type of intervention (i.e., do Alpha, do Beta, do Gamma, and do nothing) during the control task. Fig. (5) depicts the experienced mean changes separately for target 140 and target 280. Note that differences between the two target levels are due to differences in starting values. For example, if the current value of the transmitter is 150 choosing 'do Beta' results only in a small relative change $((150 \cdot .5) + 80 = 155)$. For high starting values the observed relative change may even be negative, since the decay overrides the positive impact of the intervention (e.g., if one chooses 'do Beta' at a level of 240 the resulting value is $(240 \cdot .5) + 80 = 200$). As can be seen from Fig. (5), the observed changes do not correspond to the actual parameters of the underlying causal model. However, their order reflects the order of impacts of the available interventions ('do Alpha' > 'do Gamma' > 'do Beta' > 'do nothing').

The mean changes observed by the individual participants were then used as predictors for their intervention choices (Test phases 2a/b) and predicted value ratings (Test Phases 3a/b). To predict the choices in Test phase 2a, the mean averages for each option were added to the starting value and the difference to the target was computed. The option resulting in the minimal difference to the target was assumed to provide the best choice for a participant encoding only the relative impact of the interventions. Consider a situation in which the current value of the transmitter is 120 and the target value is 140. Further assume that the mean changes observed by the participant were +100 for 'do Al-

²Note that the mean change heuristic is identical to the natural means heuristic if the system is reset to the same starting value on each trial.

pha', +20 for 'do Beta', '+70 for 'do Gamma', and -75 for 'do nothing'. The mean change heuristic then predicts that the decision maker would choose 'do Beta', since $120 + 20$ comes closest to the desired target value.

To capture the implications of the removal of area *B* we assumed that decision makers simply subtracted the experienced mean change of 'do Beta' from the experienced mean change for interventions 'do Beta' and 'do Alpha'. Thus, participants following this encoding strategy should assume that the intervention 'do Beta' is not any longer effective after removing *B* from the causal system (Fig. 1b). Note that subtracting the mean change of 'do Beta' from 'do Alpha' does not imply sensitivity to causal structure. It merely assumes that participants were sensitive to the fact that a choice of 'do Alpha' was followed by an activation of areas *A* and *B*.

Exemplar-Based Strategy

The second strategy is an implementation of the generalized context model [22]. The generalized context model is one of the dominant exemplar models in categorization, but it can also be used to model repeated decision making. We chose this model because it has been widely used and proved highly successful in the area of judgment and decision making [23]. The exemplar-based strategy assumes that participants encode each trial they encounter. More specifically, for each trial the starting value of the transmitter, the action chosen, the activated areas and the resulting change in the outcome value are assumed to be stored. To predict the consequences of an option on a new trial, the starting value of this particular trial is compared to the starting values of all previous trials in which the same option had been chosen. Previous outcome changes are weighted depending on the similarity of each previous starting value to the starting value of the current trial. Trials resembling the current trial are weighted more. Identical trials receive the maximum weight. By integrating the weighted changes of outcomes from previous trials an estimate of the outcome value for the considered option for the current trial is derived. Estimates for all options are compared to the target value to identify the best option.

To model this strategy, we first computed the change participants observed for each type of intervention (i.e., do Alpha, do Beta, do Gamma, and do nothing) on each trial, $change_{j, do(x)}$. Then we computed the absolute distance of each starting value of each test trial to the starting values of all trials seen before [22]. For current trial *i* and previous trials $j_1 \dots j_n$ the distance was defined as

$$d_{ij} = |value_{i,t_0} - value_{j,t_0}| \quad (2)$$

Next, we transformed absolute distances into psychological distances reflecting perceived similarities. To do so we used the exponential transformation proposed in the generalized context model [22]:

$$\eta_{ij} = e^{-d_{ij}} \quad (3)$$

These similarities were then used to weight the observed outcome change on each trial on which the same option was chosen. By summing up over the weighted observed outcome changes and adding the current starting value the outcome value for each option was predicted. The following formula was used:

$$value_{i,do_x,t_1} = \sum_{j=1}^{j=n} \eta_{ij} \cdot change_{j,do_x} / \sum_{j=1}^{j=n} \eta_{ij} + value_{i,t_0} \quad (4)$$

For example, suppose a participant chose 'do Beta' twice, once with a starting value of 100 points resulting in a change of +30 points and once with a starting value of 150 resulting in a change of -15 points. Now the participant is confronted with a starting value of 140 and tries to predict the resulting value using these two previous trials. The distances of the new starting value to the previous starting values are 40 and 10 points respectively, therefore similarities are $4.25 \cdot 10^{-18}$ and $4.54 \cdot 10^{-5}$. Thus the predicted value is $(4.25 \cdot 10^{-18} \cdot 30 + 4.54 \cdot 10^{-5} \cdot (-15)) / (4.25 \cdot 10^{-18} + 4.54 \cdot 10^{-5}) + 140 = 125$. As the example shows, closer previous trials exhibit a much stronger influence on the prediction. Predicted values for each option were then compared to the target value. The option whose value comes closest to the desired target value is chosen.

To capture the implications of the removal of area *B* we assumed that decision makers would use the knowledge they had about cases in which only area *B* was activated before. These were the cases in which the option 'do Beta' was chosen. Therefore the removal of *B* can be accounted for by subtracting the predicted value of option 'do Beta' from the predicted value of 'do Alpha' and 'do Beta' for each particular test case. Thus,

$$value_{do_Alpha|B_{removed},t_1} = value_{do_Alpha,t_1} - value_{do_Beta,t_1} \quad (5)$$

Imagine that the value of 'do Alpha' has to be predicted for a starting value of 160 points after area *B* had been removed. Then the predicted values for both 'do Alpha' and 'do Beta' are computed. Let's assume that the resulting values are +80 for 'do Alpha' and -1 for 'do Beta'. Hence the predicted value of 'do Alpha' given a removal of area *B* would be $160 + 80 + 1 = 241$. Note that this subtraction does not imply that participants are sensitive to causal structure. Like for the mean change strategy it merely entails that participants are sensitive to the fact that 'do Alpha' is followed by an activation of areas *A* and *B*.

Normative Causal Model Strategy

The third strategy assumes that participants use the information about the decay and the feedback on the interventions to infer the structure and the parameters of the underlying causal system. Hence, it presupposes that the actual parameters of the causal system are inferred. Choices and estimates are then assumed to be based on the correctly parameterized causal model. Predictions derived from this normative causal model have been presented in the previous sections (cf. Table 1). We decided to include this strategy as a benchmark, although the analyses presented above indicate that participants in general did not follow this strategy.

Subjective Causal Model Strategy

The fourth strategy assumes that participants engaged in causal learning and correctly inferred the underlying causal structure, but failed to derive correct parameter estimates. Thus, it may be the case that participants engaged in *causal* decision making, but based their choices and predictions on their *subjective* estimates of the causal model's parameters. These subjective parameter estimates may deviate from the

actual parameters since extracting the causal model's parameters required participants to disentangle the impact of the different interventions from the influence of the decay.

To model this strategy, we used participants' estimates for the consequences of the interventions to derive their individual parameter estimates of the causal model. We first reconstructed participants' assumptions about the decay by dividing their individual predicted value estimates for the consequences of 'do nothing' before the causal structure modification by the respective starting value (140 or 280) and then taking the mean of these two estimates. Assume a participant estimated that 'do nothing' at current transmitter values of 140 and 280, respectively, would result in values of 90 and 160. The subjective estimates for the decay would then be $(90/140 + 160/280) / 2 = .61$.

Next we calculated the supposed influence for the other interventions from the predicted value estimates given for the original causal model. To do so, we multiplied the starting values by the participants' subjective estimate of the decay and then subtracted the result from the participants' subjective estimate for the resulting transmitter value for this intervention. Assume a participant predicts that the resulting transmitter level would be 220 for 'do Alpha' at a current value of 140. Her subjective estimate of the decay was .6. The subjective estimate for the impact of 'do Alpha' would then be $220 - (140 \cdot .6) = 136$. We again averaged the results for each intervention across the two starting values (140 and 280). Note that the resulting parameter estimates were based on the ratings participants gave before the causal system was modified. In order to predict the parameters of the modified causal system, the parameter for 'do Beta' was subtracted from the derived estimates of 'do Alpha' and 'do Beta'. Hence, the parameter values decreased to zero for 'do Beta' and remained unchanged for 'do Gamma' and 'do nothing'. These individual parameter estimates (including the decay) were then used to predict the individual choices in Test Phases 2a and 3a (i.e., the intervention choices made for the original and modified causal system) and the predicted values for each intervention after area B had been removed (Test Phase 3b).

Subjective Causal Model Without Decay

A fifth possible strategy is that participants induced a causal model and estimated its parameters, but assumed that no decay would be present once an intervention was taken. That is, decision makers may have assumed that stimulating the brain by means of radiation eliminates the transmitter's decay. Causal Bayes net theories in fact make the assumption that certain ('strong') interventions eliminate all other causal influences on the variable intervened on [24]. Another reason for not assuming decay when acting upon the system might have been that the impact of the decay was not directly observable during the control task. The decay was only explicitly observed when the 'do nothing' option was chosen, otherwise the transmitter level was merely observed to change up- or downward. The assumption of no decay given interventions entails different parameters for the causal model. To calculate these parameters we again used participants' ratings of predicted values in Test Phase 2b (i.e., the predicted value estimates for the unmodified causal system). The only difference to the previous model is that the decay is ignored when the impact of the three

interventions is calculated. Otherwise the calculation is identical to the previous strategy. The same holds for the derivation of the predictions.

Evaluating the Decision Models

To evaluate the different strategies we compared the models' predictions with participants' individual choices and ratings. For the choice task, the number of correctly predicted interventions was computed separately for each participant.³ The number of correct predictions ranged from zero (none of the choices was predicted) to four (every choice made was predicted by the model). The resulting distributions for all strategies are depicted in Fig. (6). For example, the normative causal model correctly predicted all four choices made by a participant with respect to the system observed during learning (i.e., before modification) for 15 participants, it predicted 3 out of 4 choices for another 34 participants, 2 out of 4 for ten participants, and only 1 out of 4 for 4 participants. There were no participants for which the normative causal model failed completely. In Fig. (6) also the prospective results from a random choice model are depicted. Random choices were modeled by a binomial distribution with $n = 4$ choices and a probability of $p = .25$ to predict the actual choice by chance. If participants chose randomly, then none of a participant's choices should be correctly predicted for 20 participants, only one choice for 27 participants, two choices for 13 participants and 3 out of 4 choices for 3 participants. Given random choices no complete matches are expected (see Fig. (6) black lines and diamonds).

As Fig. (6a) shows, all strategies predicted participants' choices before the modification of the causal system better than the random choice model. The resulting distributions all deviated significantly from the random binomial distribution ($\chi^2_{\text{mean_change}} = 1963.0, p < .001$; $\chi^2_{\text{exemplar}} = 1179.0, p < .001$; $\chi^2_{\text{causal_normative}} = 1250.9, p < .001$; $\chi^2_{\text{causal_subjective+decay}} = 1239.7, p < .001$; $\chi^2_{\text{causal_subjective+no_decay}} = 1260.2, p < .001$). Hence, all strategies can predict the decisions made by the participants with respect to the system they intervened upon in the initial control task. A different picture arises for the intervention choices made regarding the modified causal system in which variable B was removed. As shown in Fig. (6b) only the mean change heuristic, the exemplar-based strategy and the subjective causal model without decay were able to predict participants' choices, although the predictions from the other strategies also deviated from the random choice model ($\chi^2_{\text{mean_change}} = 179.4, p < .001$; $\chi^2_{\text{exemplar}} = 123.3, p < .001$; $\chi^2_{\text{causal_normative}} = 17.0, p = .002$; $\chi^2_{\text{causal_subjective+decay}} = 61.2, p < .001$; $\chi^2_{\text{causal_subjective+no_decay}} = 190.6, p < .001$). Thus, only strategies not explicitly taking the system's momentum (i.e., the decay of the outcome value) into account were able to predict participants' choices before and after the causal modification. Moreover, the three strategies had almost the same predictive power (cf. Fig. 6b).

³We decided to include all participants in the evaluation of the five models. Although it would be interesting to see whether participants who correctly identified the underlying causal model used different strategies, we decided not to split up participants into subgroups. The reasons are that (i) participants had a probability of 50% to pick the true model in the model selection task by chance, and (ii) all other conceivable indicators of causal assumptions would have to be derived from the variables that are used in modeling.

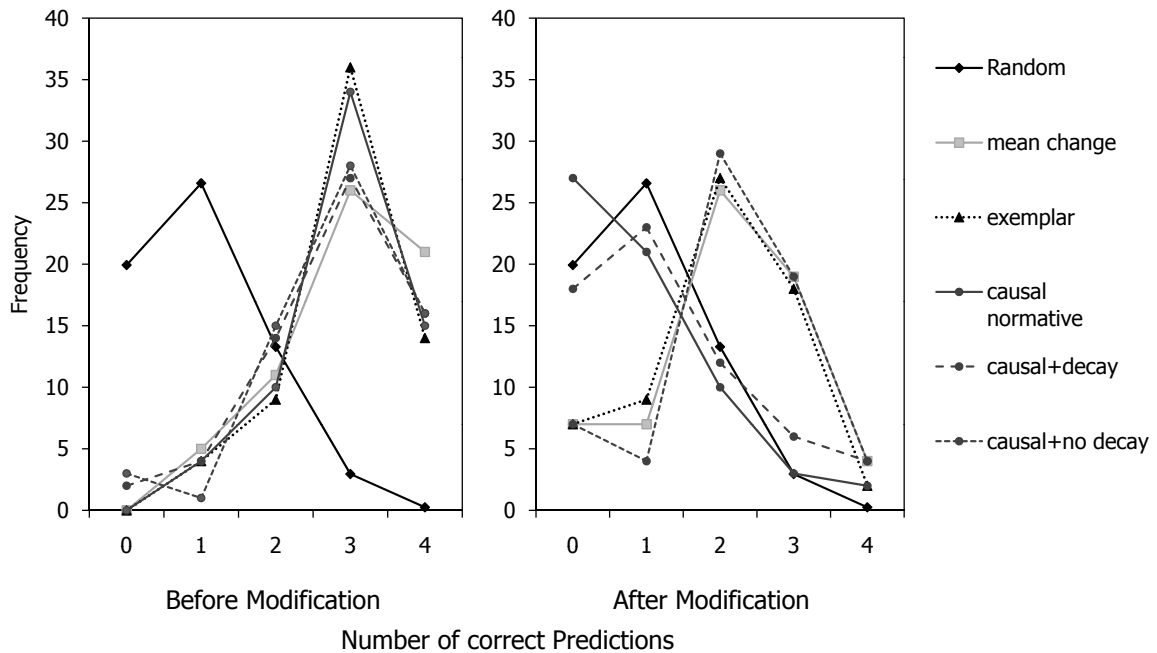


Fig. (6). Results of strategies modeled in Experiment 1: Mean change heuristic, normative causal model, subjective causal model with decay, subjective causal model without decay. Depicted is the distribution of the number of correctly predicted choices for each participant. As a benchmark the predictions of a random choice model are depicted (binomial distribution with $n=4$ and $p=.25$). See text for explanations.

Next we compared the models' predictions with participants' estimates of the predicted values for the modified causal system (i.e., the ratings given in Test Phase 3b). Recall that we calculated the parameters of the causal model strategies from participants' ratings in Test phase 2b. To assess the fit between participants' actual ratings of the transmitter level and predictions we calculated the normalized root mean square errors (NRMSE)⁴. Lower values of NRMSE indicate a better fit between the strategies' predictions and participants' predicted value estimates. It turned out that the subjective causal model without decay fitted slightly better than the mean change heuristic and the exemplar-based strategy, $\text{NRMSE}_{\text{mean_change}} = .285$, $\text{NRMSE}_{\text{exemplar}} = .234$, $\text{NRMSE}_{\text{causal_normative}} = .261$, $\text{NRMSE}_{\text{causal_subjective+decay}} = .242$, $\text{NRMSE}_{\text{causal_subjective+no_decay}} = .213$. Within-participants t-tests⁵ confirmed that the causal model without decay predicted participants ratings better than the mean change heuristic ($t(62) = 10.4$, $p < .01$), the exemplar model ($t(62) = 3.04$, $p < .01$), and the normative causal model ($t(62) = 2.57$, $p < .05$). Only a marginal difference resulted for the causal model with decay ($t(62) = 1.71$, $p < .10$).

Summary

The findings of this first experiment provide mixed evidence for the induction of causal models in the context of a control task. On the one hand, the results of the causal model selection task indicate that a majority inferred the true underlying causal structure. On the other hand, we found only minimal evidence that participants inferred the actual parameters of the system. Neither their choices of interven-

tions for the modified causal system nor their predicted values for these interventions conformed to the true causal model and its parameters. Even when we reconstructed participants' parameter estimates from their ratings on an individual basis answers could not be predicted (see results for the subjective causal model with decay). However, one strategy predicted participants' choices and ratings quite well: A causal model strategy which assumes that interventions eliminate decay.

EXPERIMENT 2

A possible explanation for the findings of Experiment 1 is that participants assumed that no decay was present whenever they intervened on the system. This idea receives some support from the literature on causal models. For example, a number of recent studies have examined how people reason about interventions that deterministically fix the state of a variable [9-11, 14, 24] (see [25] for a formal analysis based on causal Bayes net theory). These studies show that people tend to assume that interventions override the influence of the target variable's natural causes. Thus, one may speculate that decision makers in Experiment 1 assumed that the decay is only present when no intervention is attempted, but is eliminated once an intervention is made. An immediate consequence of this assumption would be that they did not attempt to separate the causal impact of the interventions from the influence of the decay in order to derive the causal model's parameters. The modeling results also provide some support for this hypothesis. The subjective causal model without decay fitted participants' choices and estimates quite well.

If this hypothesis is correct, then there should be a substantial difference when participants are provided with in-

⁴This measure normalizes the deviations on the range of values.

⁵The distributions of NRMSEs were one-peaked and not obviously skewed. Therefore t-tests are a robust way to test for differences.




Fig. (7). (a) Number of chosen interventions during learning in Experiment 2 (means and SE), (b) distance of actual transmitter level to target level (means and SE). The two blocks (trial 1-20 and trial 20-40) denote the two learning phases with the different target values (140 and 280). At the beginning of each block, the target and starting value were reset.

formation that the decay is present even when they intervene on the causal system. In this case they should infer the parameters of the underlying causal system more accurately and should derive better predictions for the modified causal system. Therefore we replicated Experiment 1, but now on each trial showed participants how the transmitter level decreased due to the decay, before the positive impact of the chosen intervention was added.

Method

Participants and Design

61 students from different universities in Berlin participated (32 females, 29 males, mean age = 25.0). They were paid €5 for participating. Again they were randomly assigned to one of the two counterbalanced conditions that only differed in the placement of the causal model selection task. 30 participants were asked about the underlying causal structure directly after the control task, 31 were asked after all other test phases.

Materials and Procedure

The materials and procedure were the same as in Experiment 1. The only difference was that on each trial during the control task participants could observe the impact of the decay. After making their choice, participants first observed the activation of the brain areas resulting from the chosen intervention. Then they saw how the value of the transmitter decreased by 50%. Finally, they observed how the transmitter level raised to its final level (assuming they had intervened on the system, if they had chosen 'do nothing' the level remained at the state caused by the decay). For example, when the current state was +100 and participants chose 'do Alpha' they would first observe how brain areas A and B

became successively activated. Next they would observe how the transmitter level decreased to 50 points (due to the decay); then they would observe it rose to 170 points (due to the intervention). The whole sequence was presented as one animation without any interruptions.

Results

Control Task

The interventions chosen by the participants are depicted in Fig. (7a). The choice pattern was very similar to the first experiment. Participants preferred to choose intervention 'do Beta' when the target level was 140 and interventions 'do Alpha' and 'do Gamma' when that target value was 280. These choices are in accordance with all of the discussed theoretical accounts, including decision making based on a correctly parameterized causal model. Fig. (7b) depicts the learning curve for the participants (i.e., the distance between the target level and the achieved level of the transmitter), which indicates that they successfully learned to control the system and to maintain the target level of the transmitter. A within-participants ANOVA comparing the average difference in the first ten trials to the average distance in the last ten trials confirmed participants' improvement, $F(1,60) = 16.8, p < .001, MSE = 302.5$.

Causal Model Selection Task

As with the previous experiment a majority of participants (80%) correctly preferred the causal chain model over the independent causes model. This percentage again deviated significantly from random choices, $\chi^2(df = 1, N = 61) = 22.4, p < .01$. A 2×2 Chi-Square test revealed no difference between the two counterbalancing conditions, $\chi^2(df = 1, N =$

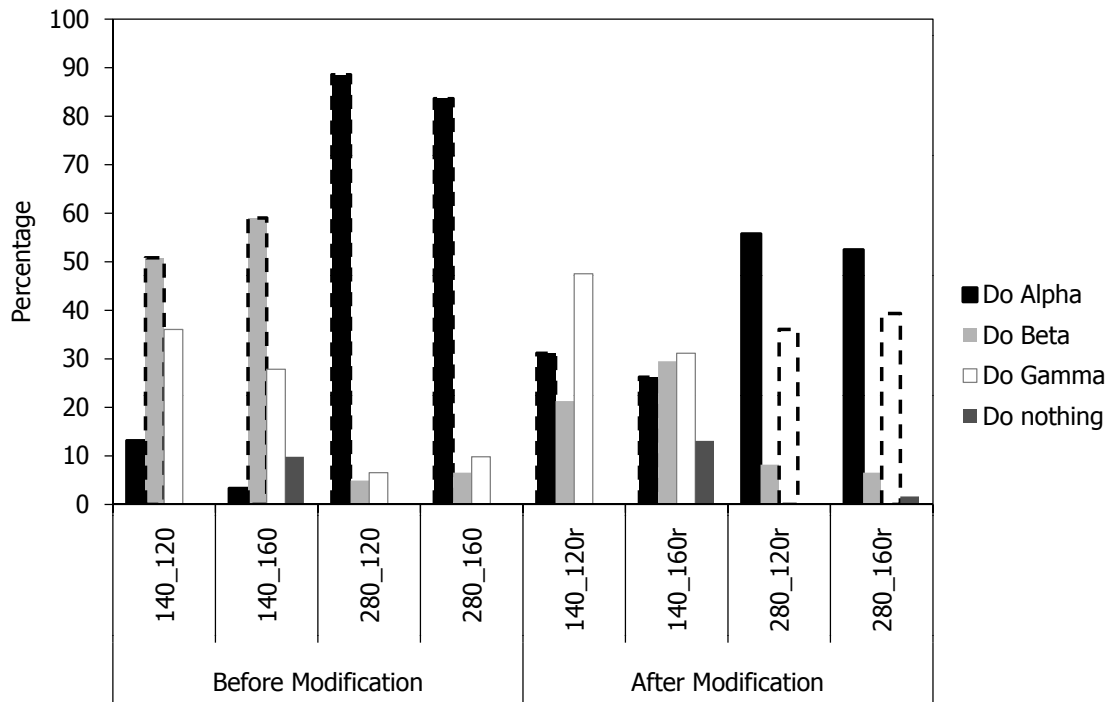


Fig. (8). Percentage of participants choosing a certain intervention in the test phase before the causal system was modified and after the modification (Experiment 2). Numbers on the x-axis indicate the specific test case, the first number is the target value (e.g. 140); the second number is the starting value (e.g. 120). Choices predicted on the basis of the correct causal model are indicated by dashed borders.

61) = .50, $p > .50$. We therefore pooled the data for all subsequent analyses.

Intervention Choices⁶

Fig. (8) depicts participants' choices for the system they acted upon during the control task (left hand side) as well the interventions chosen for the modified causal system (right hand side). For the original system, participants preferred to choose 'do Beta' given a target of 140 and 'do Alpha' given a target of 280. Thus, choices again conformed to the predictions derived from the correctly parameterized causal model. However, the obtained choice pattern differs considerably for the modified causal system. After the removal of area *B* participants chose 'do Alpha' and 'do Gamma' more often for a target of 140, and 'do Gamma' more often for a target of 280. Most of these changes were in accordance with the predictions derived from the correctly parameterized causal model. A Bowker test for target level 140 confirmed a significant difference between the choices made before and after the modification, $\chi^2(df = 6, N = 122) = 42.8, p < .001$. For a target level of 280 we again used a McNemar test for choices of 'do Alpha' and 'do Gamma' due to low cell entries. This test again yielded a significant result, $\chi^2(df = 1, N = 122) = 33.1, p < .001$, indicating that participants switched from 'do Alpha' to 'do Gamma'. Thus, in contrast to Experiment 1 we found clear indications that participants acquired substantial knowledge of the causal structure and its parameters and used this knowledge to adapt their choices when variable *B* was removed from the system.

Predicted Values

Participants' estimates of the transmitter level resulting from the interventions are shown in Table 2. Numbers on the left show the estimates for the original causal system; numbers on the right represent participants' estimates for the modified causal system. Participants significantly reduced the predicted values of 'do Alpha' and 'do Beta' for the modified causal system while the predicted values for 'do Gamma' and 'do nothing' remained about the same. This inference pattern was obtained for both starting values (see Table 2 for the results of the statistical analyses). Thus the qualitative pattern conformed to the predictions derived from a correctly parameterized causal model. Although the quantitative estimates still deviated from the values derived from a causal model analysis, much larger differences between the original and modified causal structure were obtained than in Experiment 1 (cf. Table 1). A closer inspection of the data revealed that 21 (target 140) and 19 (target 280) out of the 61 participants lowered their estimates for 'do Beta' to the amount implied by the parameters of the causal model ($\Delta_{\text{before-after}} = -80 \pm 20$). On the other hand, 15 (target 140) and 22 (target 280) participants did not lower their estimates ($\Delta_{\text{before-after}} \geq 0$). The rest of the participants lowered their estimates to a different degree.

Evaluation of Strategy Models

Participants' intervention choices and predicted value estimates were again modeled using the strategies defined in Experiment 1 (mean change heuristic, exemplar-based strategy, normative causal model, subjective causal model with decay, and subjective causal model without decay). Based on the predictions derived for each individual participant the

⁶We again tested whether the order of the test questions affected participants' choices. There was no difference between the two counterbalancing conditions, all $\chi^2 < 1, p > .33$.

Table 2. Mean Estimates of Effects Resulting from Interventions in Experiment 2. Numbers in Italics are Predictions Derived from a Correctly Parameterized Causal Model, Numbers in Brackets are Standard Errors

| | Target & Start =140 | | Target & Start = 280 | |
|------------|----------------------------------|------------|----------------------------------|------------|
| | Before | After | Before | After |
| do Alpha | <i>230</i> | <i>150</i> | <i>300</i> | <i>220</i> |
| | 199.8 | 155.6 | 315.2 | 274.8 |
| | (4.82) | (4.98) | (5.18) | (8.57) |
| | $F(1, 59) = 66.5; p < .001^{**}$ | | $F(1, 59) = 32.4; p < .001^{**}$ | |
| do Beta | <i>150</i> | <i>70</i> | <i>220</i> | <i>140</i> |
| | 150.9 | 111.0 | 248.4 | 215.1 |
| | (4.37) | (5.07) | (6.03) | (8.51) |
| | $F(1, 59) = 53.8; p < .001^{**}$ | | $F(1, 59) = 14.3; p < .001^{**}$ | |
| do Gamma | <i>190</i> | <i>190</i> | <i>260</i> | <i>260</i> |
| | 163.3 | 156.4 | 269.1 | 266.1 |
| | (4.17) | (3.92) | (4.77) | (5.29) |
| | $F(1, 59) = 4.33; p = .04$ | | $F(1, 59) < 1; p = .58$ | |
| do Nothing | <i>70</i> | <i>70</i> | <i>140</i> | <i>140</i> |
| | 74.3 | 76.2 | 155.2 | 156.6 |
| | (1.81) | (2.22) | (4.69) | (4.63) |
| | $F(1, 59) = 1.79; p = .19$ | | $F(1, 59) < 1; p = .44$ | |

Note. For all interventions and target values separate ANOVAs were calculated with counterbalancing condition (causal structure test first vs. causal structure test last) as a between-participants variable and modification (before vs. after) as a within-participants variable. To correct for accumulation of type 1 error the significance level was set to $\alpha = .00645$. It turned out that neither the counterbalancing condition nor the interaction of the two factors reached significance. Therefore only the main effects of the modification of the causal system are reported in Table 2. Results marked with an asterisk are significant.

number of choices matching to the actual choices was counted. The resulting distribution of matches was again compared to a binomial distribution modeling random choices (binom($n = 4, p = .25$)). Fig. (9) shows the distributions of the strategies and the random choice model. For the original system, all strategies achieved a better fit than the random choice model ($\chi^2_{\text{mean_change}} = 1565.4, p < .001; \chi^2_{\text{exemplar}} = 1846.2, p < .001; \chi^2_{\text{causal_normative}} = 1252.3, p < .001; \chi^2_{\text{causal_subjective+decay}} = 1675.4, p < .001; \chi^2_{\text{causal_subjective+no_decay}} = 371.5, p < .001$). This was also true for the modified causal system: $\chi^2_{\text{mean_change}} = 56.0, p < .001; \chi^2_{\text{exemplar}} = 106.0, p < .001; \chi^2_{\text{causal_normative}} = 142.5, p < .001, \chi^2_{\text{causal_subjective+decay}} = 280.9, p < .001; \chi^2_{\text{causal_subjective+no_decay}} = 113.1, p < .001$. Looking at the distribution and the χ^2 -values, the subjective causal model with decay strategy seems to have done best. It predicted at least 3 out of the 4 choices for 16 of the participants; the subjective causal model without decay did the same for 8 participants, the mean change heuristic for 7, and the exemplar based strategy for 4. For the remaining participants the models either did not yield differential predictions or were not able to account for at least 3 out of the 4 predictions. These results indicate that participants probably used a variety of strategies. However, most seem to have preferred a causal strategy.

In a second step, we also evaluated the models' predictions for the predicted value estimates. We again computed normalized root mean squared errors (NRMSE). The results mirror the findings obtained for the intervention choices, $\text{NRMSE}_{\text{mean_change}} = .393, \text{NRMSE}_{\text{exemplar}} = .304, \text{NRMSE}_{\text{causal_normative}} = .229, \text{NRMSE}_{\text{causal_subjective+decay}} = .214, \text{NRMSE}_{\text{causal_subjective+no_decay}} = .261$. The subjective causal model with decay strategy fitted participants' ratings significantly better than the mean change heuristic ($t(60) = 5.85, p < .01$), the exemplar-based strategy ($t(60) = 3.25, p < .01$), and the subjective causal model without decay ($t(60) = 2.80, p < .01$). There was no difference to the normative causal strategy ($t(60) = 1.18, p = .24$).

Summary

Participants were confronted with the same causal system in both Experiment 1 and 2. The only difference was that in Experiment 2 information about the causal impact of the decay was provided when decision makers intervened on the system. As before, the majority of participants was able to correctly infer the structure of the underlying causal system. In contrast to the results of Experiment 1 we now found that participants systematically adapted their choices and estimates when the structure of the causal system was modified. In accordance with the predictions of a causal model ap-

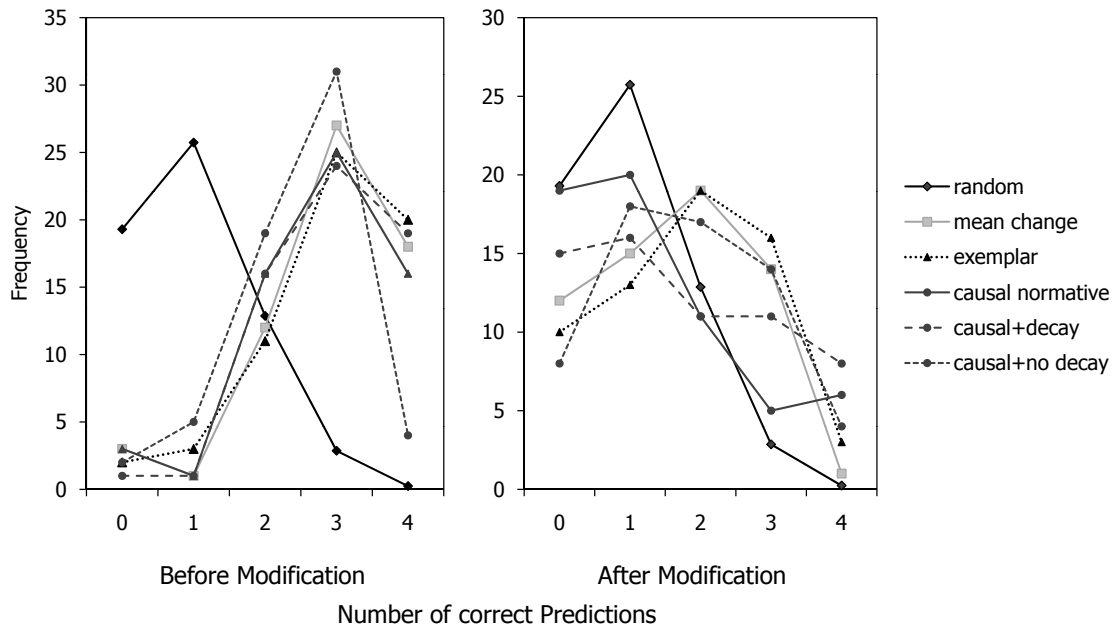


Fig. (9). Results of strategies modeled in Experiment 2: Mean change heuristic, normative causal model, subjective causal model with decay, subjective causal model without decay. Depicted is the distribution of the number of correctly predicted choices for each participant. As a benchmark the predictions of a random choice model are depicted (binomial distribution with $n=4$ and $p=.25$).

proach, decision makers switched their choices away from the previously preferred option towards the option that leads to the best outcome for the modified system. In addition, participants revised their estimates of predicted values and took into account the fact that the removal of area *B* reduced the impact of options ‘do Beta’ and ‘do Alpha’. The modeling results confirmed that the subjective causal model with decay fitted participants’ answers best. These findings show that participants learned about the causal structure and the parameters of the actual system they were confronted with.

GENERAL DISCUSSION

Most research on the control of dynamic causal systems claims that decision makers acquire only limited knowledge about the structure and the parameters of the system [1, 2]. By contrast, studies on causal learning and reasoning suggest that people have the capacity to learn about causal relations when sufficient information is available (e.g., temporal delays, statistical relations, interventions) [7, 8]. Based on the latter findings, we hypothesized that decision makers would induce causal models during control tasks, as long as the environment and experienced learning input provides sufficient information. In particular, we proposed that people would induce causal models even when they were not explicitly asked to do so, and were instead required to achieve and maintain a specific goal state.

In the two experiments reported in this paper, we presented participants with a control task that in principle allowed them to infer the underlying causal system. Although participants were only instructed to achieve and maintain a certain goal state, which they successfully did, in both experiments a clear majority of participants acquired substantial knowledge about the underlying causal structure. By

contrast, more precise knowledge about the system’s parameters depended on the transparency of the underlying causal processes. When the causal impact of the decay and the chosen intervention were not perceptually separated (Experiment 1), only a few participants inferred the causal system’s parameters. Given that they nevertheless correctly selected the causal chain model, the most parsimonious explanation for this finding is that many participants assumed that their interventions would eliminate the decay. Strictly speaking, this assumption is in accordance with the experienced feedback. In the first experiment they never saw a separate decay process when they intervened upon the causal system, but only the final value of the outcome variable resulting from both causal processes. This interpretation also receives support from other research on causal learning showing that people in general seem to understand interventions as causal variables that render the variable intervened upon independent of all other causal influences [9, 10, 11, 25, 26]. However, when the causal impact of the decay was explicitly presented to the participants (Exp. 2), the differential contributions of the chosen interventions and the decay became transparent. In consequence, more participants differentiated between the decay and the causal impact of the intermediate cause variables. Thus, it seems that participants inferred causal models in both experiments, but made different assumptions about the interplay of interventions and the decay due to differences in the transparency of the causal processes. We would like to emphasize that in our view the distinction is not a qualitative one between causal and non-causal learning, but rather a quantitative one in terms of how elaborate the acquired causal representations are.

The current findings challenge alternative theoretical models. Models that assume that decision makers only acquire situation-decision-outcome associations (e.g., the mean

change heuristic) or base new decisions on memory retrieval (e.g., the exemplar based strategy) cannot account for our results. These approaches cannot explain how people made correct causal model judgments about the intermediate variables, nor why they changed their choices and estimates systematically after the structure of the system had been modified. The modeling results showed that these alternative models can account for participants' choices from the first experiment, but not from the second. Moreover, participants' estimates of the values resulting after the modification were best predicted by a causal model strategy in both experiments (without decay in Exp. 1 and with decay in Exp. 2). Thus, although some participants probably employed alternative strategies, the induction of causal models seems to be the strategy favored by the majority of participants.

Consistent with the findings presented here is a set of recent studies showing that people tend to induce causal model representations when repeatedly intervening on a static causal system in order to maximize a specific payoff [15, 16]. For example, using a non-dynamic causal system Meder and Hagmayer showed that participants spontaneously induced causal models from probabilistic data and made appropriate inferences for both novel options and the removal of variables [16]. The experiments were designed to pit causal model theory against instrumental learning models, which entails that decision makers only acquire action-outcome contingencies. The results showed that decision makers induced a causal model representation of the choice task and acquired causal knowledge that went beyond mere action-outcome contingencies.

Conditions Promoting Causal Learning During Control Tasks

The present experiments indicate that decision makers have the capacity to engage in causal learning during control tasks. What are the conditions that enable and support such learning? Generally, it has been argued that the lack of evidence for causal learning found in previous studies results from methodological rather than psychological factors [18]. This argument also receives empirical support. For example, when decision makers are probed about their structural knowledge during the control task, they show improved insight into the underlying causal structure and capitalize on this knowledge for regulating the system [18]. Thus, in contrast to the traditional view on control learning, it seems that decision makers can go beyond solely procedural skills and acquire explicit causal knowledge.

Another factor that might promote causal learning is the kind of system acted upon. The simulated environments used to study control tasks often comprise several continuous input variables that are connected to several continuous outputs. Typically, inputs and outputs are linked by complex, nonlinear functions. In such situations, the precise form of the functional relations may simply be underdetermined by the observable data. In this case, knowledge of previously successful actions with respect to certain circumstances allows for better decisions. Unfortunately, this strategy may further constrain the observed data. For example, once a learner has achieved a specific equilibrium and knows how to maintain it, she will probably refrain from exploring other actions that may lead her away from the goal state [27]. As a

consequence, the chosen inputs and observed outputs are likely to stay in a certain range of values, thereby making it even more difficult to reveal the underlying functions. A particular kind of underdetermination of causal relations by data is also found in our studies. Although decision makers knew that a decay took place when they did not intervene on the system, the feedback they received in Experiment 1 did not allow them to infer whether the decay still maintained its influence once they acted upon the system. As a consequence, they learned little about the causal parameters. By contrast, when participants were provided with information about the impact of the decay while learning to control the system (Experiment 2), they acquired more causal knowledge.

Finally, it is probably crucial that the number of causal hypotheses is constrained by cues to causality, such as background knowledge about causal mechanisms, and the time course of events. The more knowledge a person can bring to bear, the easier it is to pinpoint the underlying causal system [8, 28, 29]. In addition, the feedback provided to participants is critical for causal induction. For example, if in our studies the status of the intermediate variables subsequent to an intervention had not been observable, it would have been impossible to infer the structure of the system and to evaluate any structural changes. In such a situation causal learning would be restricted to mere action-outcome contingencies. Imagine you turn the key of your car and it does not start. Without further information about the state of the intermediate causal variables (e.g., fuel, battery, ignition) you will be unable to identify the problem.

When these prerequisites are met, other conditions may further enhance causal induction. First, free exploration of the causal system without any specific goal has been shown to improve structure learning [20]. While in control tasks the acquisition of knowledge about the system conflicts with keeping the system in a particular state, free exploration allows the system's behavior to be scrutinized under a wide range of circumstances. The empirical evidence shows that participants who are instructed with a specific goal (i.e., maintaining a certain equilibrium state) perform worse than participants who are allowed to freely explore the dynamic system [20]. Second, motivation may also affect causal learning. Personal involvement and high stakes create the motivation to scrutinize information more carefully [30]. Third, personality factors like "need for cognition" and intelligence contribute to successfully controlling a system and learning about its structure [31, 32]. Fourth, expertise can both hamper and promote causal induction. Expertise is in general connected to elaborate background knowledge, which constrains possible causal hypotheses, so that causal learning becomes easier. On the other hand, experts have a huge knowledge base of previously encountered cases, which often allows for exemplar-based decision making. Research on naturalistic decision making has shown that experts tend to use their exemplar knowledge when a given case resembles previously encountered situations, but rely more on causal induction and mental simulations when they cannot retrieve an appropriate case from memory [33]. Finally, knowing in advance about possible changes of the causal system may create additional motivation to engage in causal learning. As pointed out before, causal representa-

tions, but not other forms of knowledge, allow people to quickly react to changes in the causal system, including the arrival of new options [15, 16].

CONCLUSION

The findings reported in this paper indicate that people spontaneously induce causal models when they make judgments and decisions about a dynamic system. They do this even when simpler learning and decision making strategies would have sufficed for the control task assigned to them. We take this as evidence that people prefer to induce causal models so long as the situation allows for causal learning. Only when the task structure does not allow them to engage in causal learning will they resort to simpler strategies. Hence, we do not claim that when people control a dynamic system they will always succeed in learning the underlying causal model, but we suspect that they will often try.

REFERENCES

- [1] Berry D, Broadbent DE. On the relationship between task performance and associated verbalizable knowledge. *Q J Exp Psychol* 1984; 36: 209-31.
- [2] Berry D, Broadbent DE. The combination of implicit and explicit knowledge in task control. *Psychol Res* 1987; 49: 7-15.
- [3] Dienes Z, Fahey R. The role of implicit memory in controlling a dynamic system. *Q J Exp Psychol* 1988; 51: 593-614.
- [4] Dienes Z, Fahey R. Role of specific instances in controlling a dynamic system. *J Exp Psychol Learn Mem Cogn* 1995; 21: 848-62.
- [5] Berry D. The role of action in implicit learning. *Q J Exp Psychol* 1991; 43: 881-906.
- [6] Burns BD, Vollmeyer R. Goal specificity effects on hypothesis testing in problem solving. *Q J Exp Psychol* 2002; 55: 241-61.
- [7] Einhorn HJ, Hogarth RM. Judging probable cause. *Psychol Bull* 1986; 99: 3-19.
- [8] Lagnado DA, Waldmann MR, Hagmayer Y, Sloman SA. Beyond covariation: Cues to causal structure. In: Gopnik A, Schulz L, Eds. *Causal learning: Psychology, philosophy, and computation*. Oxford: Oxford University Press 2007; pp. 154-72.
- [9] Meder B, Hagmayer Y, Waldmann MR. Inferring interventional predictions from observational learning data. *Psychon Bull Rev* 2008; 15:75-80.
- [10] Meder B, Hagmayer, Y, Waldmann, M R. The role of learning data in causal reasoning about observations and interventions. *Mem Cogn* 2009; 37: 249-64.
- [11] Waldmann MR, Hagmayer Y. Seeing versus doing: two modes of accessing causal knowledge. *J Exp Psychol Learn Mem Cogn* 2005; 31: 216-27.
- [12] Waldmann MR, Hagmayer, Y, Blaisdell AP. Beyond the information given: Causal models in learning and reasoning. *Curr Dis Psychol Sci* 2006; 15: 307-11.
- [13] Sloman SA, Hagmayer Y. The causal logic of choice. *Trends Cogn Sci* 2006; 10: 407-12.
- [14] Hagmayer Y, Sloman, SA. Decision makers conceive of their choice as intervention. *J Exp Psychol Gen* 2009; 138: 22-38.
- [15] Hagmayer Y, Meder B. Causal learning through repeated decision making. *Proceedings of the 30th Annual Conference of the Cognitive Science Society*; Love C, McRae K, Sloutsky VM, Eds. 2008: 179-184.
- [16] Meder B, Hagmayer Y. Causal induction enables adaptive decision making. *Proceedings of the 31th Annual Conference of the Cognitive Science Society*, Taatgen NA, van Rijn H, Eds, 2009; 1651-56.
- [17] Lagnado D, Sloman, SA. Time as a guide to cause. *J Exp Psychol Learn Mem Cogn* 2006; 32: 451-60
- [18] Osman M. Evidence for positive transfer and negative transfer/Anti-learning of problem solving skills. *J Exp Psychol Gen* 2008; 137: 97-115.
- [19] Sanderson PM. Verbalizable knowledge and skilled task performance: Association, dissociation, and mental models. *J Exp Psychol Learn Mem Cogn* 1989; 15: 729-47.
- [20] Vollmeyer R, Burns BD, Holyoak KJ. The impact of goal specificity and systematicity of strategies on the acquisition of problem structure. *Cogn Sci* 1996; 20: 75-100.
- [21] Hertwig R, Pleskac TJ. The game of life: How small samples render choice simpler. In: Chater N, Oaksford M, Eds. *The probabilistic mind. Prospects for Bayesian Cognitive Science*. Oxford, UK: Oxford University Press 2008; pp. 209-36.
- [22] Nosofsky RM. Similarity, attention and the identification-categorization relationship. *J Exp Psychol Gen* 1986; 115: 39-57.
- [23] Juslin P, Olsson H, Olsson AC. Exemplar effects in categorization and multiple cue judgment. *J Exp Psychol Gen* 2003; 132: 133-56.
- [24] Pearl J. *Causality*. Cambridge, MA: Cambridge University Press 2000.
- [25] Sloman SA, Lagnado DA. Do we “do”? *Cogn Sci* 2005; 29: 5-39.
- [26] Lagnado D, Sloman SA. The advantage of timely intervention. *J Exp Psychol Learn Mem Cogn* 2004; 30: 856-76.
- [27] Fugelsang JA, Thompson VA. A dual-process model of belief and evidence interactions in causal reasoning. *Mem Cogn* 2003; 31: 800-15.
- [28] Griffiths TL, Tenenbaum JB. Structure and strength in causal induction. *Cogn Psychol* 2005; 51: 334-84.
- [29] Griffiths TL, Baraff ER, Tenenbaum JB. Using physical theories to infer hidden causal structure. *Proceedings of the 26th Annual Conference of the Cognitive Science Society*; Korbus K, Gentner D, Regier T, Eds 2004. 500-505.
- [30] Petty RE, Cacioppo JT. The elaboration likelihood model of persuasion. In: Berkowitz L, Ed. *Advances in experimental social psychology*, 29. New York: Academic Press 1986; pp. 123-205.
- [31] Putz-Osterloh W. *Problemlöseprozesse und Intelligenztestleistung*. Bern: Huber 1981.
- [32] Suess H-M. Intelligenz und komplexes Problemlösen – Perspektiven für eine Kooperation zwischen differentiell-psychometrischer und kognitionspsychologischer Forschung. *Psychol Rundsch* 1999; 50: 220-8.
- [33] Klein G. *Sources of Power*. Cambridge, MA: MIT Press 1998.

Received: August 30, 2009

Revised: December 15, 2009

Accepted: January 07, 2010

© Hagmayer et al.; Licensee Bentham Open.

This is an open access article licensed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/3.0/>) which permits unrestricted, non-commercial use, distribution and reproduction in any medium, provided the work is properly cited.