

Causal Induction Enables Adaptive Decision Making

Björn Meder^{1,2} (meder@mpib-berlin.mpg.de)

York Hagmayer² (york.hagmayer@bio.uni-goettingen.de)

¹Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

²Department of Psychology, University of Göttingen, Gosslerstr. 14, 37073 Göttingen, Germany

Abstract

The present paper examines the interplay between causal reasoning and decision making. We use a repeated decision making paradigm to investigate how people adapt their choice behavior when being confronted with changes in the decision environment. We argue that people are sensitive to the causal texture of a decision problem and adjust their choice behavior in accordance with their causal beliefs. In the first study we examine how people adapt their decision making behavior when new options whose consequences have not been observed yet become available. In the second study the causal system underlying the decision problem is modified to investigate how prior experiences with the choice task affect decision making. The results show that decision makers' choice behavior is strongly contingent on their causal beliefs and that they exploit their causal knowledge to assess the consequences of changes in the decision problem situation. A high consistency between hypotheses about causal structure, expected values, and actual choices was observed.

Keywords: Decision making; Causal reasoning; Learning

Introduction

How do decision makers assess the implications of changes in the decision environment? And how do they adapt their choice behavior to such changes? Previous research on adaptive decision making (e.g., Payne, Bettman, & Johnson, 1993) has primarily focused on strategy selection. Going beyond this research, we here focus on adaptivity in terms of flexibly responding to changes of the causal underpinnings of the decision problem. Our main hypothesis is that decision making is contingent on causal considerations and that decision makers exploit their causal knowledge to adapt to the respective structure of the choice situation.

For example, cancer is often treated by chemotherapy. However, the potential benefits of this treatment strongly depend on a number of factors, such as the constitution of the patient. When the patient's liver is working properly, chemotherapy is often the most promising treatment. However, when a patient suffers from liver dysfunction it may be necessary to switch to a different treatment, such as radiation therapy. Thus, the potential benefits of the available courses of action strongly depend on the specific properties of the underlying causal system. Therefore, knowledge of the causal system enables the decision maker to determine which of the different options is most promising under the prevailing circumstances. Importantly, it does so without requiring any further learning experience.

Surprisingly, most theories of decision making still neglect the importance of causal considerations. For example,

likelihood \times value theories (e.g., expected utility theory) distinguish between options, possible outcomes, and the associated uncertainties, but causal learning and causal reasoning are not addressed by these theories. Rather, these accounts implicitly assume that likelihood estimates correctly mirror causal relations, although observable statistical ("evidential") relations may not necessarily reflect underlying causal processes (Hagmayer & Sloman, 2009).

Other researchers, however, have emphasized the tight connection between causal reasoning and decision making. For example, Sloman and Hagmayer (2006) have argued that people tend to construct a mental *causal model* (Waldmann, Hagmayer, & Blaisdell, 2006) of the choice situation. A causal model of the decision problem encompasses the causal influences between options, outcome events, and payoffs. It therefore enables decision makers to mentally simulate the consequences of the available courses of actions. In a number of studies Hagmayer and Sloman (2009) demonstrated that peoples' decisions are contingent on their causal beliefs when making simple one-shot decisions in hypothetical scenarios. These studies also showed that people spontaneously activate their causal beliefs before making a choice.

Repeated Decision Making: Becoming Adapted and Being Adaptive

In the present paper, we focus on decision situations in which people repeatedly face a binary choice task with the goal of maximizing their payoff. Since no prior information concerning the payoff distributions is provided, the outcomes of the available options have to be assessed from the experienced feedback. Erev and Barron (2005) referred to this setting as *minimal information paradigm*. In such studies decision makers are presented with two options (e.g., two buttons *A* and *B*) between which participants can choose. Each choice leads to a certain payoff drawn from an unknown distribution. To describe the learning process during such repeated decision making, Barron and Erev (2003) put forward the *value assessment model*, a reinforcement learning model that estimates the payoff distributions and expected values (EV) of the available options from feedback. A similar approach underlies instrumental learning theories, which describe how an organisms' choice behavior is shaped through reinforcement.

A characteristic feature of such reinforcement models is that they entail that choice behavior gradually changes in accordance with the experienced contingencies between

actions and outcomes. Thus, the decision maker becomes adapted to the decision problem. The underlying learning mechanisms also entail that agents can adapt to changes in the choice situation. For example, imagine the payoff structure is altered after a number of learning trials. Because such a change results in different feedback, the decision maker can adapt to the new situation. However, adaptation is rather slow since re-learning is influenced by the previously acquired knowledge (i.e., expected value estimates, associative weights). Moreover, new learning experiences are essential – no feedback, no adaptation.

However, knowledge gained from previous experiences with a particular choice situation can also enable the decision maker to flexibly respond to changes in the decision context. This is the case when (i) the choices pertained to a causal system, and (ii) the observed consequences allowed it to infer the structure of the underlying causal system. Then, people may acquire knowledge that reflects the causal influences in the environment. For example, endocrine therapy has proved to be effective for treating breast cancer. The benefit of this therapy is based on the fact that some types of cancer cells possess hormone receptors which are responsible for the nutrition of the cells which, in turn, affects tumor growth. Thus, blocking these receptors by a drug stops tumor growth. However, when the tumor cells do not possess these receptors, the therapy is ineffective. In this case, knowing that one of the variables in the causal chain is missing allows the agent to directly infer that taking this action is not any longer sensible. As envisioned by Tolman and Brunswik (1935), representations that mirror the causal texture of the environment are highly adaptive.

To induce causal structure, people can capitalize on different cues, such as prior knowledge, spatio-temporal contiguity, and statistical information (Lagnado, Waldmann, Hagmayer, & Sloman, 2007). The *causal model theory of choice* (Sloman & Hagmayer, 2006; Hagmayer & Sloman, 2009) extends this idea to decision making by assuming that people tend to induce a causal model of the decision problem and the choice situation. A causal model of a decision problem comprises the causal influences between actions, outcomes, and payoffs. Figure 1 shows some simple examples of causal models. These models detail the relations between the available options (L , W), the intermediate outcome variables (A , B) that are causally affected by the available courses of action, and the payoffs associated with the occurrence of the outcome variables. A characteristic feature of causal models is that they represent only causal relations, but not the statistical contingencies that are generated by them (cf. Sloman, 2005). For example, the relation among endocrine therapy, the blockade of receptors, and tumor growth can be described by a causal chain like the one shown in Figure 1 left hand side.

Goals and Hypotheses

In a previous study (Hagmayer & Meder, 2008) we examined how pre-existing causal beliefs about a domain guide repeated decision making. To do so, we manipulated

decision makers' beliefs about the causal structure underlying the decision problem while keeping the observed consequences constant. The results showed that, depending on peoples' initial causal beliefs, identical learning experiences can lead to very different conclusions. These beliefs, in turn, strongly affected participants' reactions to changes of the underlying causal system.

In the present set of studies we did not provide participants with an initial causal hypothesis. In fact, we did not even point out that causal knowledge might be helpful. Participants were not informed that the causal system they initially acted upon might change. Also, instead of presenting participants with deterministic causal systems like we did in our previous studies we here examine decision making with probabilistic causal systems. Thus, the available courses of actions only probabilistically generated their outcomes. Finally, we used two different modifications of the decision problem to assess whether participants are able to adapt to changes. In the first experiment we surprisingly provided participants with a novel option, whose consequences had never been observed before but could be inferred from a causal model representation of the decision situation. In the second experiment participants were unexpectedly informed that a causal variable had been removed from the system. Again, causal knowledge would allow them to spontaneously and accurately adapt to the new situation.

Based on previous research into causal learning and causal decision making we expected participants to acquire a causal model representation of the decision problem (Hagmayer & Meder, 2008; Hagmayer & Sloman, 2009). Following up on research on causal learning, which has shown that people can exploit causal knowledge to infer the outcomes of hypothetical interventions from causal model representations (Meder, Hagmayer, & Waldmann, 2008), we expected participants to be able to assess the consequences of the novel option in Experiment 1. We also expected them to capitalize on their causal model representation when the causal system underlying the decision problem is altered (Experiment 2).

Experiment 1

The goal of the first study was to examine whether people would adapt to the causal texture of a repeated decision making task and spontaneously induce a causal model representation. In particular, the question was how people would react to changes of the decision problem, and whether they would be capable to adapt their choice behavior accordingly. The experiment consisted of two repeated decision making phases, with the second phase being the test phase in which unexpectedly an additional option was introduced. In both phases participants' task was to maximize the value of a certain payoff variable. However, feedback about the outcomes of the decisions made was only provided in the initial decision making phase, but not in the test phase.

Figure 1 shows the two causal structures used in this study, *causal chain* (CH) and *common cause* (CC), and the associated feedback structures (the relations between options, variables, and payoff were counterbalanced across

participants). In the *Causal Chain condition*, option *L* influenced variable *B* only by way of *A* ($Do L \rightarrow A \rightarrow B$), whereas in the two common cause conditions *Do L* independently affected *A* and *B* ($A \leftarrow Do L \rightarrow B$). We employed two different common cause conditions. In one condition (*Common Cause 1*), the available options (*L*, *W*) had identical expected values as in the causal chain condition. This manipulation, however, requires different probabilistic relations than the ones in the chain condition (e.g., $P(B | Do L)_{Chain} > P(B | Do L)_{CC1}$). We therefore designed a second common cause condition (*Common Cause 2*) in which all causal relations had the same strengths as in the chain condition. As a consequence, the expected value for *Do L* was higher than in the chain condition. However, the rank order of the options' expected values was the same as in the other two conditions (i.e., $EV(Do L) > EV(Do W)$).

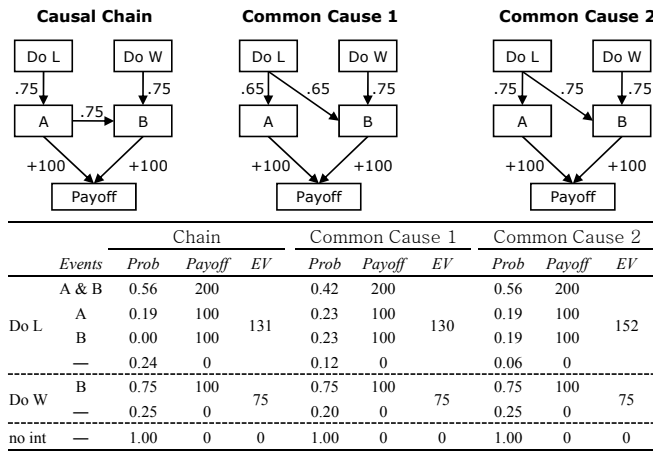


Figure 1. Causal structures and feedback of Experiment 1.

Participants and Design Sixty University of Göttingen undergraduates took part for course credit or were paid 7€. They were randomly assigned to one of the three conditions (*Causal Chain*, *Common Cause 1*, *Common Cause 2*).

Materials and Procedure We used a biological scenario according to which certain bacteria produce a vaccine against diseases. Participants were told that the production of the vaccine is regulated by two genes, *A* and *B*, which are inactive by default. Then they were instructed to try to produce as much vaccine as possible by activating these genes through applying two “trigger substances”, *L* and *W*. Participants were not informed how exactly the trigger substances relate to the activation of the genes, but it was pointed out that the two genes may also be causally interrelated.

The initial *Repeated Decision Making (RDM) Phase* consisted of 100 decision trials in which participants were requested to maximize the amount of produced vaccine by repeatedly choosing one of the three options (*Do L*, *Do W*, no intervention). Each of the 100 decision trials referred to bacteria whose genes were inactive prior to any intervention. After making a decision, participants first observed which genes became active and then received information on the payoff (amount of produced vaccine). The temporal order conformed to the underlying causal structure. Thus, in

the chain condition participants observed first that gene *A* became activated and then, with a delay of 1s, that gene *B* also became active. In the common cause condition both genes turned on simultaneously. As option *L* had a much higher expected value than option *W* (cf. Figure 1) we expected participants to prefer option *L* regardless of condition.

After making 100 decisions the *Test Phase* began, which consisted of 10 additional decision trials. In this phase a novel option was introduced. Decision makers were told that a new trigger substance had been developed, which reliably activated gene *A*. Thus, in the test phase participants had to choose among the known options *Do L* and *Do W* as well as the new option *Do A*. In this phase, however, decision makers received no feedback regarding the state of the intermediate variables (*A* and *B*) or the resulting payoff. Thus, participants could not simply learn about the consequences of the new option. As decision makers never experienced the outcomes of this option, they would have to rely on their causal knowledge to assess whether the new option would be superior to the existing ones. In the chain condition, due to the causal relation $A \rightarrow B$ the fact that $P(A | Do A) = 1.0$ implies that the new option has a higher expected value than *Do L* ($EV(Do A)_{Chain} = 175 > EV(Do L)_{Chain} = 131$). By contrast, in the common cause conditions intervening on *A* would not affect *B*. Therefore, $EV(Do A)_{CC1} = EV(Do A)_{CC2} = 100$, which is inferior to *Do L*. Thus, to maximize payoff decision makers should opt for *Do A* when assuming that there is a causal relation between *A* and *B* (i.e., in the chain condition) but stick with *Do L* when there is no causal relation (i.e., in the common cause conditions).

Table 1. Mean number of choices (\pm SEM) and received mean payoff (\pm SEM) in Experiment 1.

	RDM Phase						Test Phase		
	<i>Do L</i>		<i>Do W</i>		<i>no int</i>		<i>Do L</i>	<i>Do W</i>	<i>Do A</i>
	<i>N</i>	Payoff	<i>N</i>	Payoff	<i>N</i>	Payoff	<i>N</i>	<i>N</i>	<i>N</i>
Chain	74.8 (3.6)	132.4 (2.3)	21.8 (3.1)	72.8 (2.7)	3.5 (0.8)	0.0 (0.0)	2.2 (0.4)	0.1 (0.6)	7.8 (3.3)
CC 1	77.7 (2.6)	127.3 (1.8)	20.3 (2.4)	74.3 (2.0)	2.0 (0.4)	0.0 (0.0)	5.4 (0.8)	0.5 (0.8)	4.2 (0.8)
CC 2	73.5 (4.2)	151.4 (1.6)	23.7 (3.7)	70.3 (3.2)	2.9 (0.9)	0.0 (0.0)	6.2 (0.9)	0.5 (0.7)	3.3 (0.6)

Results and Discussion Table 1 (left hand side) shows the results of the initial decision making phase. The obtained choice pattern shows that, regardless of the underlying causal model, decision makers had a clear preference for option *L*. Accordingly, no difference was obtained between the three conditions. Also, the experienced payoffs closely resemble the options' actual values (cf. Figure 1). By contrast, participants' choices during the test phase revealed differential preferences. Decision makers in the chain condition exhibited a strong preference for the new option ‘*Do A*’, which indicates that they inferred that they would gain a higher

payoff when switching from *Do L* to *Do A*. By contrast, a very different pattern was obtained in the two common cause conditions. In these conditions participants preferred to stick with option *L*. Consequently, more *Do A* choices were obtained in the chain than in the two common cause conditions, $t(38) = 3.50, p < .001$ and $t(38) = 5.23, p < .001$, respectively. The two common cause conditions did not differ.

Subsequent to the test phase participants were asked to provide estimates of the expected payoff for all three options (*Do L*, *Do W*, *Do A*). Estimates for options *L* and *W* closely resembled the actual values (cf. Fig. 1), though participants in the CC 2 condition slightly underestimated the expected value of *Do W*. The crucial analyses concern decision makers estimates for option *Do A*, whose actual consequences they never observed. The obtained judgments reveal a strong sensitivity to the underlying causal structure: Estimates in the chain condition were significantly higher than in condition CC 1 [$t(40) = 2.4, p = .02$] and condition CC 2 [$t(40) = 2.3, p = .03$]. Again, the two common cause conditions did not differ from each other.

Table 2. Mean estimates (\pm SEM) for payoffs and causal model choices in Experiment 1.

	Expected Value Estimates			Model Choices	
	<i>Do L</i>	<i>Do W</i>	<i>Do A</i>	<i>Chain</i>	<i>CC</i>
Chain	134.8 (8.0)	75.0 (4.0)	140.0 (8.8)	17	3
CC 1	141.3 (6.4)	68.8 (4.8)	112.0 (7.8)	6	14
CC 2	157.0 (10.1)	59.6 (7.3)	113.5 (7.5)	5	15

The final dependent variable aimed to directly tap onto decision makers' causal beliefs about the decision problem. Participants were presented with graphs of a causal chain and a common cause model (similar to the ones shown in Fig. 1, but without the numbers). Then they were asked to indicate which of the two models would correctly describe the causal relations between options, intermediate variables, and the payoff. In all conditions a majority of participants chose the correct model: 85% in the chain condition and 70% and 75%, respectively, in the two common cause conditions (Table 2, right hand side).

Overall there was a strong concordance between the preferences decision makers revealed through their choices, their expected value estimates, and their assumptions about the causal structure of the decision problem. The kappa correlation between these three measures was $\kappa = .71$.

Taken together, the results provide strong evidence that participants learned about the causal texture of the repeated decision making task. Rather than merely encoding expected values, a majority of participants spontaneously induced a causal representation of the decision problem, though they were never requested to do so. The acquired causal knowledge, in turn, allowed decision makers to predict the expected value of the novel option whose payoffs

they had not encountered previously. In consequence, 70% of participants who were confronted with a causal chain model preferred the new option, while only 33% who acted on a common cause model did so.

Experiment 2

In this study we used a different test strategy to examine whether decision makers spontaneously induce a causal model, which would allow them to adapt to a change in the decision problem. Again the experiment consisted of two repeated decision making phases, with the second phase being the test phase. In the test phase we modified the decision problem by removing one of the variables observed during learning. Hence, we here used another procedure to examine whether participants acquired causal models and capitalized on this knowledge to adapt to a new situation.

Participants and Design Forty-eight University of Göttingen undergraduates participated for course credit or were paid 7€. The factor 'causal model' (causal chain vs. common cause) was manipulated between conditions.

Materials and Procedure We used the same materials and procedure as in Experiment 1. The only difference was that this time there were three genes (*A*, *B*, *C*) instead of two. Participants first had to make 100 decisions with the goal of maximizing their payoff, followed by a test phase comprising 10 additional decisions. Again feedback was only provided in the initial decision making phase, but not in the test phase.

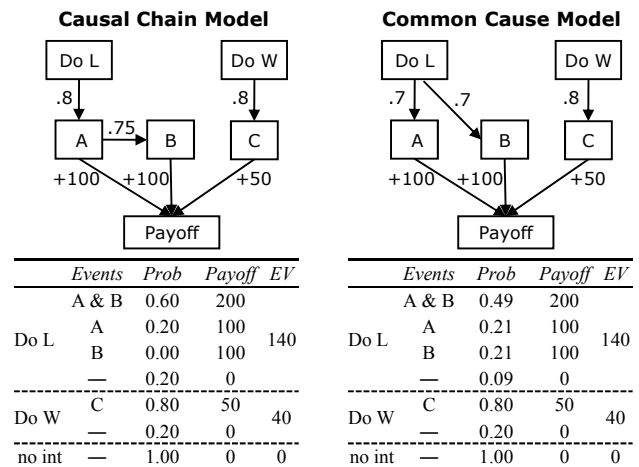


Figure 2. Causal structure and feedback of Experiment 2.

Figure 2 shows the two experimental conditions, causal chain (CH) and common cause (CC) and the associated feedback structures. Like in Experiment 1, in the causal chain condition option *L* influences *B* only by way of *A* ($Do L \rightarrow A \rightarrow B$). In the common cause condition, by contrast, *L* is directly related to both *A* and *B* ($A \leftarrow Do L \rightarrow B$). Thus, whereas in the chain condition the presence of *A* is a necessary event for the occurrence of *B* (i.e., $P(B|\neg A) = 0$), this is not the case in the common cause condition (i.e., $P(B|\neg A) > 0$). However, despite this difference the expected values (EV) of the available options were identical across

conditions: $EV(Do L) = 140$, $EV(Do W) = 40$, and $EV(no int) = 0$. Thus, to maximize the payoff one should choose option L regardless of the underlying causal model.

The instruction to the test phase informed decision makers that they now would be presented with bacteria that did not possess gene A . Thus, variable A was suddenly removed from the causal system. This removal has very different implications for the two causal systems. In the causal chain condition L affects B only by way of A , therefore, the expected value of L decreases from 140 to zero. Accordingly, $Do W$ becomes the better option ($EV(Do L | no A) = 0 < EV(Do W | no A) = 40$). In the common cause condition the removal of A entails a decrease of $EV(Do L)$, too. However, due to the direct link $L \rightarrow B$, opting for L remains the better option ($EV(Do L | no A) = 70$ vs. $EV(Do W | no A) = 40$). Thus, in order to maximize participants in the chain condition should switch from $Do L$ to $Do W$, whereas participants in the common cause condition should stick with option L . As before, no feedback was provided in the test phase.

Subsequent to making their ten decisions for the modified causal system participants were also requested to estimate the expected payoffs for all options for both decision making phases. Finally, participants were presented with a figure similar to the one depicted in Figure 2, but without any arrows (i.e., only the variables were depicted). Decision makers' task was to express their causal hypotheses by drawing all causal relations they assumed to hold between options, outcome variables, and payoff. Thus, the goal was to elicit participants' representations of the decision problem.

Table 3. Mean number of choices (\pm SEM) and received mean payoff (\pm SEM) in Experiment 2.

	RDM Phase						Test Phase	
	$Do L$		$Do W$		$no int$		$Do L$	$Do W$
	<i>N</i>	<i>Payoff</i>	<i>N</i>	<i>Payoff</i>	<i>N</i>	<i>Payoff</i>	<i>N</i>	<i>N</i>
Chain Model	73.6 (3.4)	141.6 (1.9)	21.7 (2.8)	41.3 (1.3)	4.7 (1.2)	0.0 (0.0)	2.8 (0.7)	7.2 (0.7)
CC-Model	79.25 (3.0)	143.6 (1.4)	17.9 (2.7)	42.6 (.98)	2.8 (1.2)	0.0 (0.0)	5.5 (0.7)	4.5 (0.7)

Results and Discussion Table 3 depicts participants' choices for the two decision making phases and the obtained mean payoffs. As expected, participants exhibited a clear preference for option L in the initial decision making phase regardless of condition. Statistical analyses revealed no differences between conditions for any of the options. Also, the experienced values accurately mirrored the true values; here also no difference between conditions was obtained. By contrast, a clear difference between conditions was obtained for the test phase in which variable A was removed from the system. Participants chose $Do W$ significantly more often in the chain condition than in the common cause condition, $t(46) = 2.83, p < .01$. Conversely, the mean of $Do L$ choices was higher in the common cause condition than in the causal chain condition, $t(46) = 2.72, p < .01$. The fact that participants in the common cause condition exhibited only a slight

preference for L over W is probably due a trade-off between mean and variance (i.e., opting for L gives 100 points with $p = 0.7$ while option W results in 50 points with $p = 0.8$).

Participants' choice behavior was also consistent with their expected value estimates (cf. Table 4, left hand side). As with the choices, no differences resulted for the first phase but only for the test phase. In both conditions participants realized that a removal of variable A would decrease the expected value of option L . Most importantly, they were very sensitive to the fact that the exact amount of decrease depends on the underlying causal structure. In accordance with the respective underlying causal model they gave lower estimates for $Do L$ in the causal chain than in the common cause condition, $t(46) = 2.69, p = .01$.

Table 4. Expected payoffs (\pm SEM) for the two decision making phases and indicated causal models in Exp. 2.

	Expected Value Estimates				Model Hypotheses			
	RDM Phase		Test Phase		<i>CH</i>	<i>CC</i>	<i>CH+</i>	
	$Do L$	$Do W$	$Do L$	$Do W$			<i>CC</i>	<i>other</i>
Chain Model	149.4 (5.4)	45.6 (3.8)	37.6 (7.4)	40.6 (4.5)	10	6	4	4
CC-Model	137.1 (5.9)	44.1 (2.5)	67.3 (8.2)	43.3 (2.8)	4	6	9	5

Finally, we analyzed the causal models drawn by the participants (Table 4, right hand side). Unexpectedly, a number of participants (27%) indicated that L is a common cause (i.e., $A \leftarrow Do L \rightarrow B$) and that there is a direct relation $A \rightarrow B$. Since this is not completely inconsistent with the observed statistical relations, participants probably went for the "safe option" of causal overdetermination. The exactly correct model was chosen by 33% of participants. Thus, about 60% of the decision makers induced a causal model that was consistent with the obtained feedback. The rest, however, drew a model which was inconsistent with the observations made. This finding indicates that not all participants inferred the correct causal model underlying the decision problem. However, we also suspect that the free elicitation procedure was more difficult for participants than the forced choice task used in Experiment 1. Follow up analyses revealed that there again was a substantial convergence between participants' preferences revealed through their choices, expected values and causal models. The kappa correlation between choices and expected values was $\kappa = .73$ and between expected values and model implications $\kappa = .70$.

Overall, the results demonstrate that participants were remarkably sensitive to the causal texture of the decision task. Again, many decision makers seemed to have spontaneously induced a causal model representation, which allowed them to flexibly adapt their choice behavior to the removal of a causal variable from the system. Replicating the results from the previous experiment, 68% of participants encountering a causal chain switched their preferences away from the previously favored option, while only 38% dealing with a common cause system did.

General Discussion

In the present experiments we studied the interplay between causal induction and decision making. The results provide strong evidence that human decision making is sensitive to the causal texture of a given decision problem. First, in both experiments a majority of the participants spontaneously induced a causal model representation of the decision problem, although they were never asked to do so. Second, decision makers exploited their causal knowledge to adapt their choice behavior to changes in the decision context. Participants switched to a novel option whose consequences had not been observed when the causal model entailed that this action would lead to a better outcome. They also gave up the previously preferred option when the removal of a variable from the causal system implied that this option would no longer be the superior action.

Neither of these findings can be explained by a pure reinforcement learning model. Simple models which only encode the expected values of options must fail because the intermediate variables that make the crucial difference between conditions are not represented at all. However, also reinforcement learning models which represent the intermediate variables and their associations to options on the one hand and payoffs on the other hand fail. Consider the first experiment. Even if the agent encodes the payoff generated by each intermediate variable separately, she would still fail to differentiate between the two causal models since it is the causal interrelatedness of the outcome variables that makes the crucial difference (cf. Hagmayer & Meder, 2008).

However, spontaneous causal induction during repeated decision making may also be limited by a number of factors. First, the experienced feedback must enable the decision maker to discover the underlying causal structure. Thus, feedback on the state of the variables within the system and cues to causal structure must be available (cf. Lagnado et al., 2007). Impoverished outcome feedback pertaining only to statistical relations among actions and outcomes is clearly not sufficient to build up causal model representations that go beyond action-outcome contingencies. Second, with an increasing complexity of the decision problem the induction of causal models becomes more difficult and data alone is rarely sufficient. In these cases previous causal knowledge about the domain becomes crucial. However, given a certain amount of prior knowledge even sparse and noisy data may be sufficient to determine the underlying causal structure (e.g., Griffiths, Baraff, & Tenenbaum, 2004).

One may suspect that the use of causal knowledge in decision making is limited to the simple problems examined here. However, there is a growing body of evidence indicating that this is not the case. For example, causal considerations play an important role in psychodiagnostic decision making, too (e.g., Kim & Ahn, 2002). Also, when experts cannot identify the best solution immediately they tend to construct simplified models of the domain to evaluate potential courses of action (Klein, 1998). Thus, the flexibility and adaptivity of causal model representations pays off in naturalistic decision making contexts as well.

Acknowledgments

This research was supported by a grant of the Deutsche Forschungsgemeinschaft (DFG HA 3406/3-1). We thank Iris Risse and Katharina Müller for collecting the data.

References

- Kim, N., & Ahn, W.-K. (2002). Clinical psychologists theory based representations predict their diagnostic reasoning and memory. *Journal of Experimental Psychology: General*, *131*, 451-476.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215-233.
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*, 912-931.
- Griffiths, T. L., Baraff, E.R., & Tenenbaum, J. B. (2004). Using physical theories to infer hidden causal structure. *Proceedings of the 26th Annual Conference of the Cognitive Science Society*.
- Hagmayer, Y., & Meder, B. (2008). Causal learning through repeated decision making. *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp.179-184).
- Hagmayer, Y., & Sloman, S. A. (2009). People conceive of their choices as intervention. *Journal of Experimental Psychology: General*, *138*, 22-38.
- Klein, G. (1998). *Sources of Power*. Cambridge, MA: MIT Press.
- Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation: Cues to causal structure. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 154-172). Oxford: Oxford University Press.
- Meder, B., Hagmayer, Y., & Waldmann, M. R. (2008). Inferring interventional predictions from observational learning data. *Psychonomic Bulletin & Review*, *15*, 75-80.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge, MA: Cambridge University Press.
- Sloman, S. A. (2005). *Causal models*. Cambridge, MA: Oxford University Press.
- Sloman, S. A., & Hagmayer, Y. (2006). The causal psychology of choice. *Trends in Cognitive Science*, *10*, 407-412.
- Tolman, E. C., & Brunswik, E. (1935). The organism and the causal texture of the environment. *Psychological Review*, *42*, 43-77.
- Waldmann, M. R., Hagmayer, Y., & Blaisdell, A. P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science*, *15*, 307-311.