Betting on Transitivity in an Economic Setting

Dennis Hebbelmann (Dennis.Hebbelmann@psychologie.uni-heidelberg.de)
Momme von Sydow (Momme.von-Sydow@uni-heidelberg.de)
Department of Psychology, University of Heidelberg, Hauptstr. 47
69117 Heidelberg, Germany

Abstract

Theories of causal reasoning and learning often implicitly assume that the structural implications of causal models and empirical evidence are consistent. However, for probabilistic causal relations this may not be the case. Mismatches between structural implications and empirical evidence may lead to distortions of empirical evidence. Previous work has shown that people may use the generative causal relations *A* → *B* and *B* → *C* to infer a positive relation between events *A* and *C*, despite data showing that these events are actually independent (von Sydow et al., 2009, 2010). Here we used an economic trial-by-trial learning scenario to investigate how transitive reasoning in intransitive situations with even negatively related distal events may relate to betting behavior. Experiment 1 shows that transitive reasoning does affect not only probability estimates but betting as well. Experiment 2 shows that the effect remains stable even after repeated betting and feedback.

Keywords: causal induction, transitivity, Markov condition, betting, economic reasoning, coherence-based induction

Transitive Reasoning in Probabilistic Causal Chains

Causal model theory and causal Bayes nets (Pearl, 2000; Sloman, 2005; Spirtes, Glymour, & Scheines, 2001; Waldmann, 1996; Waldmann, Cheng, Hagmayer, & Blaisdell, 2008) build on the assumption of the Markov condition, stating that any node in a causal model is conditionally independent of all upstream nodes, given its parents (Hausman & Woodward, 1999; Spohn, 2001). This entails transitivity in causal chains: If *A* causes *B*, and *B* causes *C*, then *A* causes *C* (via *B*). If in a probabilistic causal chain the Markov condition holds, then the strength of the global relation *A* → *C* can be infered from the strength of the local relations *A* → *B* and *B* → *C*, which means that transitivity holds: When using the causal strength estimate ΔP ($\Delta P_{AB} = P(B|A) - P(B|\neg A)$; Jenkins & Ward, 1965), the global ΔP can be calculated by multiplying all local ΔPs that make up the causal chain (e.g., $\Delta P_{AC} = \Delta P_{AB} * \Delta P_{BC}$). It is therefore not necessary to observe the global relation directly.

Related research on transitive reasoning in the induction of causal chains has shown that people assume a transitive causal relation based on integrating single causal links (Ahn & Dennis, 2000; Baetu & Baker, 2009). This research corroborated the hypothesis that people reasoned transitively even if no information on the distal event was shown.

Subsequent research started to investigate *intransitive* chains (von Sydow, Meder, & Hagmayer, 2009; von Sydow, Meder, Hagmayer, & Waldmann, 2010). This allows bringing bottom-up evidence (*correspondence*) and top-down inferences based on the structural assumptions about causal models (*coherence*) into conflict.

They suggest a *causal coherence hypothesis* that coherence-based induction may distort bottom-up evidence about causal relations considerably if the bottom-up data do violate the structural assumptions of Bayes Nets. People are taken, at least by default, to assume a modular integration of single causal relations into larger causal networks, for instance implying transitivity in causal chains. This is predicted even when evidence to the contrary is available, but people may give up this default belief if the mismatch between coherence-based induction and correspondence-based induction gets very evident.

Intransitive chains are at odds with structural implications of Bayes Nets and involve a violation of the Markov condition. In the philosophical debate it has been put into question whether all causal relations necessarily adhere to the Markov condition and, as a consequence, whether chains need to be transitive (Cartwright 2001, 2006; Sober & Steel, 2012). However, even strict advocates of the Markov condition have pointed out that on the level of our *actually used* categories causal chains may not adhere to the Markov condition (Hausman & Woodward, 1999; Spohn, 2001). For instance, this may be the case if a category is the product of mixing subclasses for which different causal relations hold.

Von Sydow et al. (2009, 2010) showed in several formats (overview format, trial-by-trial format) that participants may infer the relation *A* → *C* from *A* → *B* and *B* → *C*, even if this is not warranted by the data presented to them: In the materials used, *A* → *B* and *B* → *C* were positive, while *A* and *C* were statistically independent from each other ($\Delta P_{AB} = \Delta P_{BC} = .5$, $\Delta P_{AC} = 0$). Participants in accordance with the causal coherence hypothesis still judged *A* → *C* in line with transitivity if they were presented with *A* → *B* and *B* → *C* first. This effect remained stable even when participants were able to directly assess the data about *A* → *C*. However, in many regards the boundary conditions of the causal coherence hypothesis need further exploration. For instance it is not clear whether people continue to infer a positive distal causal relation from positive local relations if it is clearly negative or if they completely switch to bottom-up induction due to the obvious mismatch.

Causal Reasoning and Decision Making

We here transfer the idea of the causal coherence hypothesis and intransitive chains to the field of decision making.

It has recently increasingly been emphasized that valid human decision making involves causal reasoning, because it allows for accurate predictions and effective interventions in an agent's environment. It has been shown that if an intervention changes a causal system, people base their decisions on their causal beliefs and assess utilities not only based on previously observed contingencies but on causal inferences (Hagmayer & Meder, 2013, cf. Hagmayer & Sloman, 2009). The same should be the case in betting tasks investigated here. When asked to bet on the occurrence of an event people should take their knowledge about the presence or absence of possible causes into account.

However, we here investigate non-transitive causal chains to test the causal coherence hypothesis. Although betting may also reduce the coherence-based distortion of evidence, we predict that people's bets on the occurrence of a possible effect are informed by both bottom-up learning and top-down inferences based on the assumption of transitivity, two sources of information that contradict each other in this case. Their betting may either correspond to their probability judgments (probability matching; cf. Vulkan, 2000) or to an optimal exploitation of their given probability judgment: If optimizers use bottom-up induction and realize the negative distal relation they should put all stakes on the negative prediction, if they use a top-down approach to infer the distal relation, they should put all stakes on the positive prediction.

Goals and Hypotheses

In the two experiments presented here we investigated the influence of causal coherence on people's decision making in an environment where transitivity is violated. We further examined whether causal coherence still affects participants' judgments if the distal events in the chain are not only *independent* of each other, but their relation even runs *contrary* to the assumption of transitivity (a negative global relation when transitivity suggests a positive one and vice versa). In an economic trial by trial learning scenario participants first observed co-occurrences of four events in a non-transitive causal chain and afterwards judged the statistical relations between events.

In Experiment 1 we hypothesized that the causal coherence should not only influence participants' judgments of the global relation but also the amount of money bet in line with a transitive causal model, thereby performing worse than a control group in which causal coherence should not have an effect.

In Experiment 2 we examined whether this effect remains stable after repeated betting on the global relation.

Experiment 1: Betting Biases in Learning Relations

Participants

We tested 84 participants (50 female, age M = 23.6) who were recruited at the University of Heidelberg as part of a multi-experiment session. Participants received 6€ / hour or course credit for taking part in the experiment.

Material and Procedure

Participants were told to observe individual companies and their development during learning blocks with each trial representing an individual company (cf. von Sydow et al., 2009). Each company's development consisted of four events represented by four pictures (Figure 1): Each company either buys or does not buy stocks of a second company (A vs. $\neg A$), then rises or falls on a general performance index (B vs. $\neg B$), is positively or negatively evaluated by the *Economist* (C vs. $\neg C$), and in the end either increases or decreases in stock market value (D vs. $\neg D$). The instruction stressed the temporal order of the events, which is a known cue inducing causal structure (Lagnado & Sloman, 2006). However, we neither suggested that the chain is transitive nor that specific relations were positive or negative.

The local relations between all four events were positive ($\Delta P_{AB} = \Delta P_{BC} = \Delta P_{CD} = .5$), while the global relation $A \rightarrow D$ was negative ($\Delta P_{AD} = -.5$). Figure 2 illustrates the contingencies shown to the participants. Each of the four events occurred with a probability of $P = .5$ (dark shaded segments in Figure 2). Combining four events (and their negations) results in 16 possible trial types. Figure 2 shows all types of trials used in Experiment 1. Trial type 1, for instance, consisted of A, $\neg B$, $\neg C$, and $\neg D$ etc. (Segment 1 in the circles of Figure 2). Each of the eight types of trials was used twice in each of the learning blocks. There were 12 learning blocks, resulting in 196 learning trials.

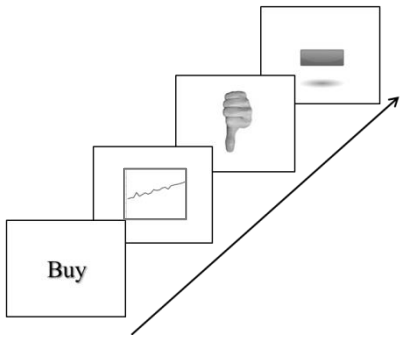

Figure 1: Exemplary trial representing one company (A, B, $\neg C$, $\neg D$).

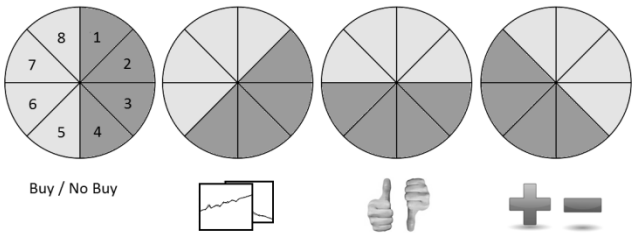

Figure 2: Structure of statistical relations between events A, B, C, and D.

The experiment consisted of learning and testing blocks. Before each learning block participants were instructed which relation to focus on and were tested on only this relation afterwards. Each learning block consisted of the same 16 trials in randomized order, regardless of condition; therefore participants did not differ in the learning material presented to them. Each trial was 4 seconds long with each picture being presented for 1 second. Participants started each trial by clicking a "Next" button on the screen.

In each test phase participants judged the relation they had focused on during the preceding learning trial on a 21-point scale ranging from -100, indicating a deterministic negative relation (e.g.: "If a company is positively evaluated then its stock market value will always decrease"), to 100, indicating a deterministic positive relation (e.g.: "If a company is positively evaluated then its stock market value will always increase"), with a middle point of 0, indicating statistical independence.

Participants were randomly assigned to one of three conditions which only differed in learning instructions and testing (Figure 3): Participants in the *local only* condition only focused (and were only tested) on the local relations $A \rightarrow B$, $B \rightarrow C$ and $C \rightarrow D$. Participants in the *local + global* condition were tested on both the local relations and the global relation $A \rightarrow D$. In the *global only* condition participants were only instructed to focus on the global relation.

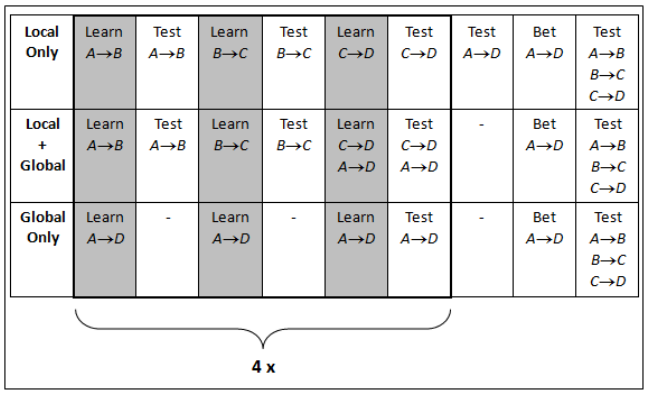

Figure 3: Temporal structure of Experiment 1.

After the learning blocks and the test blocks all participants rated the perceived relation $A \rightarrow D$ on the same 21-point scale again. They were then told that they would see one more company drawn randomly from the ones they had seen so far during the experiment. This time they only saw the company's buying decision (A or $\neg A$). They could then bet 100 cents on the development of the company's stock market value (D vs. $\neg D$). Participants could split their money between the two options and would win the amount of money they bet on the right outcome. The outcome was shown afterwards and participants were paid the amount of money they had won on top of their usual reimbursement.

After the betting trial all participants rated the local relations one last time.

Results

Estimates of the Global Relation We expected participants in the local only group to judge $A \rightarrow D$ to be positive, in line with the transitivity assumption, even though they could have seen the negative relation during 196 trials. We further predicted the global only group to judge $A \rightarrow D$ to be strongly negative, in line with the data. As the local + global group's estimates should be informed by both bottom-up learning and top-down inferences we expected their estimates to fall between the other two groups. Figure 4 shows participants' mean estimates of the global relation $A \rightarrow D$, with a positive value indicating a positive relation and vice versa. A one-way ANOVA comparing the groups mean estimates confirms this hypothesis[1]: We found a significant main effect of condition, $F(2, 81)= 23.99$, $p < .001$. Participants in the local only group judged $A \rightarrow D$ to be positive, $M = 36.2$, $SD = 29.4$, the global only group judged it to be negative, $M = -27.6$, $SD = 37.8$, with the local + global group falling between the other two, $M = 5.2$, $SD = 34.2$. A Bonferroni-corrected post-hoc comparison of group means revealed significant differences between all three groups, $ps < .01$.

Note that the local only group's estimates are even considerably higher than predicted by a perfectly transitive inference (which would correspond to an estimate of +12 on our scale). The global only group's mean estimate is closer to the estimate predicted by bottom-up processing alone (corresponding to -50 on our scale).

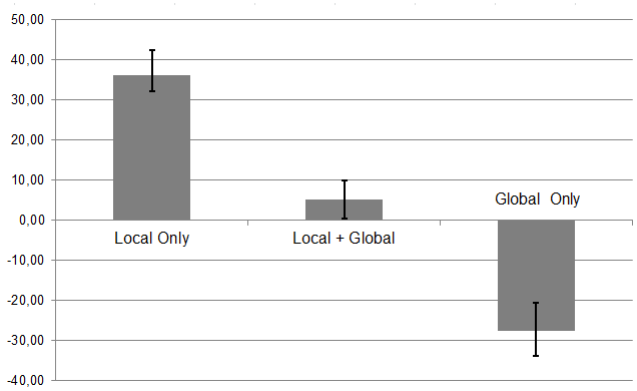

Figure 4: Mean estimates for $A \rightarrow D$.

Betting on the Global Relation To compare participants' betting performances we first calculated how much money each participant bet on the most likely outcome given the information about A vs. $\neg A$ (*ideal bet*), i.e. if participants saw an instance of A, their ideal bet would be the amount they bet on $\neg D$ (and for $\neg A$ vice versa). Figure 5 shows participants' mean ideal bets by condition. A one-way ANOVA with the ideal bet as the dependent variable again

[1] Although normal distribution was violated within conditions we still report the results of parametric tests as they have proven to be robust against this deviation. In all cases analyses using non-parametric tests led to comparable results to those reported.

showed a significant main effect of condition, $F(2, 81) = 9.73$, $p < .001$. The local only group bet significantly less money on the ideal bet, $M = 32.1$, $SD = 27.1$, than the global only group, $M = 66.6$, $SD = 29.2$, with the local + global group falling between the two, $M = 51.5$, $SD = 30.4$ (Figure 5). A Bonferroni-corrected post-hoc comparison of group means revealed a significant difference between the local only group and the other two, ps $< .05$, but not between the local + global and the global only group, $p = .15$.

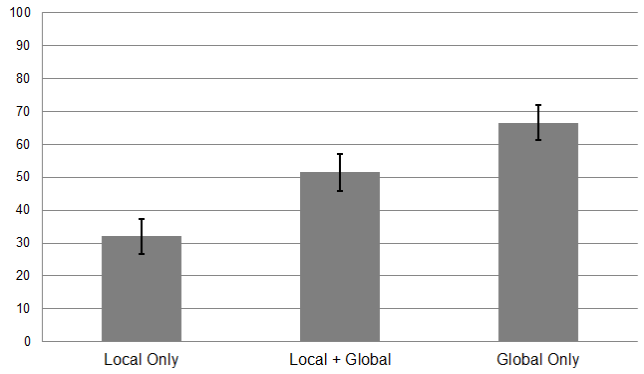

Figure 5: Mean ideal bets in ct on D vs. $\neg D$.

Discussion

In the first experiment we could replicate and expand von Sydow et al.'s (2010) findings. Participants stick with their assumptions of transitivity even if they are presented with strongly contradicting evidence during a total of 196 learning trials. However, the results in the local + global group show that this assumption is not impervious to experience: Participants' mean judgments near the point of statistical independence may either reflect an averaging of top-down assumptions and bottom-up data collection or participants' confusion about the true nature of the relation. In any case the judgment differed considerably from the value predicted by bottom-up data alone. Transitive interpretations must have had a strong impact on participants' estimates.

In the local only group participants' estimates of $A \rightarrow D$ were considerably higher than predicted by an inference purely based on transitivity. People may not use an as fine-grained scale to convey that there is a positive relation than is suggested by a fully parameterized Bayesian model. Alternatively, this deviation may be linked to previously found deviations from the Markov condition in experimental paradigms not directly assessing transitive reasoning, but showing too positive relations in functioning chains (Rehder & Burnett, 2005).

The results of the betting trial show that participants' assumptions of transitivity not only influence their judgments but also their decision making in a betting task: Participants in the local only group were willing to bet most of their money in line with the belief that $A \rightarrow D$ is positive. Participants in the global only group accurately judged $A \rightarrow D$ to be negative and bet most of their money accordingly, leading to more money bet on the most likely outcome. The observed bets suggest a sort of probability matching reflecting participants' beliefs about the probabilities of D vs. $\neg D$ given A vs. $\neg A$.

Experiment 2: Repeated Betting

Participants

We tested 94 participants (67 female, age $M = 23.3$) who were recruited at the University of Heidelberg as part of a multi-experiment session. Participants received 6€ / hour or course credit for taking part in the experiment.

Material and Procedure

Experiment 2 followed a structure similar to Experiment 1 (Figure 3), with each testing phase replaced by one betting trial as described in Experiment 1. Repeated betting on $A \rightarrow D$ should incentivize accurate learning even more, therefore putting the causal coherence hypothesis to a stronger test.

Each phase of Experiment 1 in which participants judged the relation $A \rightarrow B$ was replaced by a betting trial where participants saw A or $\neg A$ and were asked to bet on B vs. $\neg B$, etc. In each betting trial participants bet 100 points they could split between the two possible outcomes. Participants won the amount of points they bet on the right outcome and received immediate feedback about the points they won. At the end of the experiment participants were paid up to 3 € on top of their usual reimbursement depending on how many points they had collected.

Participants were again randomly assigned to either the local only, local + global, or global only group, analogously to the design of Experiment 1. At the end of the experiment participants judged the local relations and the global relation on the same scale as used in Experiment 1.

Due to the naturalistic material used in Experiment 1 participants might have had prior beliefs about $A \rightarrow D$ being positive. Their responses in the local only group may therefore not indicate transitive reasoning but rather participants' resorting to prior beliefs in the absence of further knowledge. To control for the effect of a general tendency to judge $A \rightarrow D$ positively we counterbalanced between participants whether $A \rightarrow B$ was positive or negative ($\Delta P_{AB} = .5$ vs. $\Delta P_{AB} = -.5$). With $A \rightarrow B$ being negative and the other local relations remaining positive, $A \rightarrow D$ was *positive*, $\Delta P_{AD} = .5$, but the transitive top-down prediction is *negative* for this relation. In both cases we expected participants in the local only group and the local + global group to bet in line with transitivity.

Results

Estimates of the Global Relation If $\Delta P_{AB} = -.5$ participants' answers were reverse coded. We expected participants in the local only group to judge $A \rightarrow D$ in line with the assumption of transitivity, as predicted by the causal coherence hypothesis. The global only group should judge $A \rightarrow D$ in line with the data presented during learning trials (represented by negative values in Figure 7). We expected the local + global group's estimates to fall between

the two other conditions as they should be driven by both top-down assumptions and bottom-up learning. A one-way ANOVA comparing the groups' mean estimates confirms this hypothesis: We found a significant main effect of condition, $F(2, 91)= 4.96$, $p < .01$. Participants in the local only group judged $A \rightarrow D$ to be positive, $M = 8.4$, $SD = 38.7$, the global only group judged it to be negative, $M = -25.5$, $SD = 52.4$, with the local + global group falling between the other two, $M = -8.1$, $SD = 34.0$ (Figure 7). A post-hoc comparison of group means revealed significant differences between all three groups, ps $< .01$.

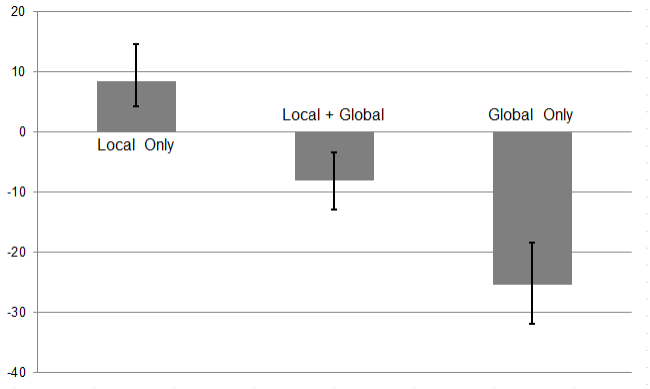

Figure 6: Mean estimates for $A \rightarrow D$.

Betting on the Global Relation Again all participants bet on $A \rightarrow D$ after the last learning block. To compare participants' betting performances we first calculated how much money each participant bet on the most likely outcome given the information about A vs. $\neg A$. A one-way ANOVA again showed a significant main effect of condition, $F(2, 91) = 18.9$, $p < .001$ (Figure 6). The local only group bet significantly less money on the ideal bet, $M = 43.5$, $SD = 19.5$, than the global only group, $M = 78.5$, $SD = 24.0$, with the local + global group falling between the two, $M = 52.7$, $SD = 25.9$. A post-hoc comparison of group means revealed a significant difference between the global only group and the other two, ps $< .001$, but not between the local only and the local + global group, $p = .12$.

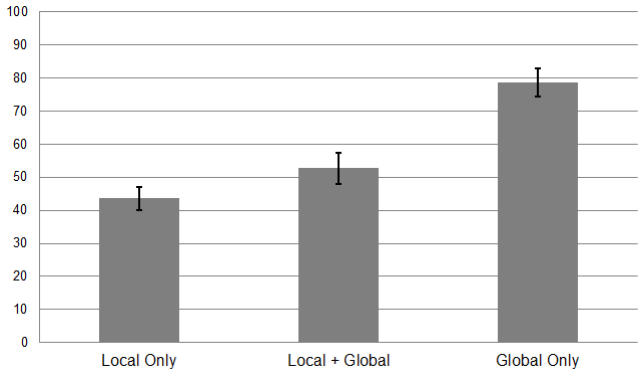

Figure 7: Mean ideal bets on D vs. $\neg D$.

Discussion

Experiment 2 replicates and extends the findings of Experiment 1. Even after betting on $A \rightarrow D$ three times (and receiving feedback about their performance) the local + global group failed to perform substantially better than the local only group during betting trials, showing that the tendency to bet in line with transitivity is hard to overcome, even if it consistently leads to non-optimal outcomes. This was also reflected in participants' judgment of $A \rightarrow D$. Again the local + global group's betting behavior might also reflect confusion about the nature of $A \rightarrow D$, because betting 50 points on both results in each bet would be the safest bet that ensures to win at least half of the possible amount of money.

General Discussion

In two studies we found compelling evidence that the causal coherence hypothesis seems to generalize to decision making in an economic context, even when using incentivized repeated betting tasks. Additionally we found that the causal coherence hypothesis holds not only for intransitive chains where the distal events are independent of each other, but also when the global relation strongly contradicts transitive inferences. In both experiments global relations were strongly negative (positive) while transitivity suggested a positive (negative) relation. All three groups differed considerably in their estimates of the global relations, showing that both sources of information, correspondence and coherence, play an important role in judging causal relations.

In Experiment 1 participants of the local only group performed significantly worse in a one-shot bet on $A \rightarrow D$, even though they had the chance to learn about $A \rightarrow D$ in a total of 196 trials.

Experiment 2 demonstrates this tendency's strength and stability: Even after repeated betting trials that led to consistently bad results for betting on transitivity, participants in the local + global group still performed significantly worse than the global only group, showing no improvement over the four betting trials.

Similar research on pseudocontingencies has likewise previously shown effects of people distorting contingencies (Fiedler, Freytag, & Meiser, 2009). However, pseudo-contingencies are usually explained based on the matching of two skewed distributions. This factor is excluded here, since we did not use skewed distributions for the single events ($P(A) = P(B) = P(C) = P(D) = .5$). Hence the postulated coherence-based inference effects cannot be explained by traditional explanations of pseudocon-tingencies and seem to add another explanation to those based on skewed distributions.

In any case, the results suggest that judgments as well as bets about distal relations in a potential causal chain can be distorted in the direction implied by transitive inferences even if transitivity is violated when sequentially inducing several local relations. Depending on the conditions, we found strong distortions in the judgments and bets largely

coherent with top-down inferences, even if direct bottom-up inferences now went into an opposed direction while people in almost 200 trials saw evidence for this relation.

Further research should address whether this is also true for situations in which participants can actively intervene, e.g. suggesting a company to buy or not buy (A vs. $\neg A$) in order to achieve rising or falling stock prices (D vs. $\neg D$). Active engagement in causal systems may both ensure participants' engagement in the task and effective encoding of predictions and outcomes of their decisions, eventually overcoming the assumption of transitivity in cases where it is invalid.

Acknowledgements

We are grateful to Shirin Betzler, Julia Folz, Alina Greis, Kamala Grothe, Vera Hampel, Antonia Lange, and Alexander Wendt for their valuable work during data collection. We would also like to thank anonymous reviewers for helpful comments and suggestions. This research was supported by the grant Sy 111/2-1 to Momme von Sydow from the Deutsche Forschungsgemeinschaft (DFG) as part of the priority program *New Frameworks of Rationality* (SPP 1516) as well as the DFG grant FI294/23-1 to Klaus Fiedler.

References

Ahn, W., & Dennis, M. (2000). Induction of causal chains. *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society* (pp. 19–24). Lawrence Erlbaum Associates, NJ: Mahwah.

Baetu, I., & Baker, A. G. (2009). Human judgments of positive and negative causal chains. *Journal Of Experimental Psychology: Animal Behavior Processes*, *35*(2), 153-168.

Cartwright, N. (2001). What is Wrong with Bayes Nets? *The Monist*, *84*, 242-264.

Cartwright, N. (2006). From metaphysics to method: comments on manipulability and the causal Markov condition. *British Journal for the Philosophy of Science, 57*, 197-218.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*, 367-405.

Fiedler, K., Freytag, P., & Meiser, T. (2009). Pseudocontingencies: An integrative account of an intriguing cognitive illusion. *Psychological Review*, *116*(1), 187-206.

Hagmayer, Y., & Meder, B. (2013). Repeated causal decision making. *Journal Of Experimental Psychology: Learning, Memory, And Cognition*, *39*(1), 33-50.

Hagmayer, Y. A., & Sloman, S. A. (2009). People conceive of their choices as intervention. *Journal Of Experimental Psychology: General, 138*22-38.Hausman, D., & Woodward, J. (1999) Independence, Invariance, and the Causal Markov Condition, *British Journal for the Philosophy of Science, 50*, 521-583.

Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs: General and Applied, 79*, 1-17.

Lagnado, D. A. & & Sloman, S. A. (2006). Time as a Guide to Cause. *Journal of Experimental Psychology, Learning, Memory, and Cognition*, *32*, 451-460.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference.* Cambridge, MA: Cambridge University Press.

Rehder, B., & Burnett, R. C. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, *50*, 264 –314.

Sloman, S. (2005). *Causal Models. How People Think about the World and Its Alternatives*. Cambridge, MA: Oxford University Press.

Sober, E. & Steel, M. (2012). Screening-Off and Causal Incompleteness: A No-Go Theorem. *The British Journal for the Philosophy of Science*, *64*, 1-38.

Spirtes, P., Glymour, C., & Scheines, R. (2001), *Causation, Prediction, and Search*. 2nd edition. New York: Springer-Verlag.

Spohn, W. (2001), Bayesian Nets Are All There Is To Causal Dependence (pp. 157-172), in: M.C. Galavotti, P. Suppes, and D. costantini (eds.), *Stochastic Dependence and Causality*, CSLI Publications, Stanford.

von Sydow, M., Meder, B., & Hagmayer, Y. (2009). *A Transitivity Heuristic of Probabilistic Causal Reasoning.* In Proceedings of the Thirty-First Annual Conference of the Cognitive Science Society (pp. 803-808). Austin, TX: Cognitive Science Society.

von Sydow, M., Meder, B., Hagmayer, Y. & Waldmann, M. R. (2010). How Causal Reasoning Can Bias Empirical Evidence. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2087-2092). Austin, TX: Cognitive Science Society.

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, *14*, 101-118.

Waldmann, M. R. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation, Vol. 34: Causal learning* (pp. 47-88). San Diego: Academic Press.

Waldmann, M. R., Cheng, P. W., Hagmayer, Y., & Blaisdell, A. P. (2008). Causal learning in rats and humans: a minimal rational model. In N. Chater, & M. Oaksford (Eds.), *The probabilistic mind. Prospects for Bayesian Cognitive Science* (pp. 453-484). Oxford: University Press.