



**Has The World Changed?
My Neighbor Might Know
Effects of Social Context on
Routine Deviation**

Tilman Betsch
Stefanie Lindow
Christoph Engel
Corina Ulshöfer
Janet Kleber





Has The World Changed? My Neighbor Might Know Effects of Social Context on Routine Deviation

Tilman Betsch/Stefanie Lindow/Christoph Engel/Corina Ulshöfer/Janet Kleber

August 2011

Has The World Changed? My Neighbor Might Know Effects of Social Context on Routine Deviation[†]

Tilman Betsch^{*}, Stefanie Lindow^{*}, Christoph Engel^{**},
Corina Ulshöfer^{***} and Janet Kleber^{****}

Abstract

In two experiments we studied the effects of behavioral models on routine deviation decisions in observers. Participants repeatedly chose among four card-deck lotteries together with a human model (confederate, Exp. 1) or a non-human model (computer, Exp. 2) that made correct decisions in the majority of the trials. In a learning phase, participants acquired a choice routine (preferring one deck over others). In a subsequent test phase, participants had to adapt to changes in the payoff structure that required them to deviate from their routine. We found a strong tendency to maintain the routine despite negative feedback (routine effect). In a social situation (Exp.1), models decrease routine effects more intensely than in non social situations (Exp.2). The process of adaptation follows a belief updating process. Results indicate that the model effect is not due to an increase of the sample of relevant information nor to application of a simply copy heuristic. Rather, deviation models may provide a cue for change that fosters reevaluation of the situation in the observer.

Keywords: Experienced-based decision making, routine, habit, adaptation, social influence, Bayesian updating, novelty

[†] We would like to thank Silvana Danz for serving as the confederate. We greatly appreciate the helpful comments of Marc Jekel, Emanuel Towfigh, Andreas Glöckner, Frank Renkewitz and Heather Fuchs on an earlier version of this paper. Address Correspondence to: Tilman Betsch, Department of Psychology, University of Erfurt, Nordhauser Strasse 63, D-99089 Erfurt, Phone: +49-(0) 361/737 2221, E-mail: tilmann.betsch@uni-erfurt.de

^{*} University of Erfurt, Germany

^{**} Max Planck Institute for Research on Collective Goods, Bonn, Germany

^{***} University of Bern, Switzerland

^{****} University of Vienna, Austria

Every day we all make hundreds of decisions. We possess the abilities to carefully deliberate and acquire additional information before we make up our minds. Yet, for most decisions, we do not call on these abilities. Instead, we quickly assess whether the decision problem appears sufficiently similar to problems we have encountered previously and, if so, follow our acquired decision routine. This is not only a means to reduce decision effort. It also prevents us from being led astray by marginal features of the situation instead of reacting to the important, core elements of the problem. If the situation is uncertain or even ambiguous, repeating what was reasonable in the past enables us to make any decision at all (Athay & Darley, 1981).

Adult decision makers possess a vast routine repertoire. Routines are defined here as behavioral options that come to mind as viable choice candidates whenever the individual re-encounters and recognizes a decision situation (Betsch & Haberstroh, 2005a,b). Routines are learned by experience. For example, assume you have rented a new flat. During the first weeks, you explore the possible options for going to work. Soon, however, when you are leaving the house for work, the decision situation has changed substantially. Due to your prior experience, you now have access to a viable choice candidate in memory, e.g. the route you chose the day before. Therefore, you can simply repeat your prior choice.

With increasing repetition and reinforcement frequency, routines may eventually freeze into habits (Aarts & Dijksterhuis, 2000; James, 1890/1950; Ouellette & Wood, 1998; Triandis, 1977; Verplanken & Aarts, 1999; Wood & Neal, 2007). However, not all past behaviors are over-learned habits, which can be implemented automatically. “Routine” is a broad concept capturing both automatic and non-automatic solutions to decision problems that have been learned by the decision maker (Betsch & Haberstroh, 2005a). Extensive behavior repetition is not necessary to anchor decisions to a particular alternative. Even only a few repetitions may suffice to add a behavioral option to an individual’s routine repertoire.

Sometimes, however, the world may change. A routine that promoted a person’s goals in the past may produce suboptimal or even negative outcomes in the future. An adaptive decision maker must be capable of detecting such changes and, if necessary, eventually decide on alternative courses of action.

In mundane situations, individuals often make decisions in a social context. They can observe other people’s behavior or form expectations about their own behavior based on social knowledge (Bandura, 1977). Interestingly, behavioral models are still rarely considered as a source of information in research on individual decision making in humans. Most pertinent studies consider the influence of models on risk seeking in individuals (e.g., Yechiam, Druryan & Ert, 2008). Routine decision making in social contexts, however, has largely been neglected (but see Mussweiler & Rüter, 2003; Reimer et al., 2005). This is the object of our study. We lead participants to acquire a decision routine. At a certain point, unbeknownst to our subjects, the situation changes such that the routine is no longer optimal. We then investigate how participants adapt to this change. We manipulate when participants receive information about others’ choices: during

the learning phase, during the test phase, or during both phases (compared to controls that never receive additional information).

Specifically, we study the following situation. Participants repeatedly choose among four gambles represented by four decks of cards (similar to the IOWA-task, Bechara et al. 1994). Accordingly, they face a probabilistic choice environment in which the monetary outcomes of the choices are risky. By making repeated decisions, participants eventually acquire a choice routine, i.e., to consistently prefer the deck with the highest expected value. Following this routine acquisition phase, the underlying pay-off structure of the task is changed so that routine choices no longer yield the highest profit. Instead, a different deck of cards has the highest expected value. Consequently, successful adaptation requires participants to detect this change, terminate their routines and henceforth prefer the alternative deck that now has the highest expected value. In each of the two phases, half of our participants are provided with information on choices and outcomes made by a model. During each trial, these participants can inspect the results of another draw from the same type of lottery task. These draws are made either in a social or a non-social situation. In Experiment 1, the individual works together with a human model who is a confederate of the experimenter. In Experiment 2, the sequence of draws is said to have been previously determined by a computer and is presented to the individual in the form of a so-called “card-information box”. Importantly, the sequence of draws and their outcomes are exactly the same in both experiments and strongly encourage routine deviation. We examine (a) to what extent such additional information helps participants adapt to the new situation (i.e., deviate from their routines) and (b) whether it makes a difference when the additional information is conveyed in a social or a non-social situation.

In the probabilistic context of our game, decision makers are not certain which decks yield which gains and losses with which probabilities nor whether the game remains constant over all repetitions. Consequently, a constant choice of the optimal deck in the learning phase may result from either deliberate (or intuitive) Bayesian updating that is repeated every trial or the application of a decision routine. Our within-subjects manipulation enables us to discriminate between these explications. In the learning phase, it is always optimal to choose the best deck, even if one has incurred a loss in earlier trials. In the test phase, losses with this deck are more frequent. Given that participants do not know that the world has changed, it is optimal to begin choosing the other decks. Therefore, if we find that a participant keeps choosing the former best deck despite frequent losses, this suggests that the participant is indeed applying a decision making routine.

From a normative perspective, it is appropriate to distinguish between (a) a routine effect and (b) a routine bias. A routine *effect* occurs if participants must first overcome the earlier routine by relearning before they begin reacting optimally to the changed situation. A routine *bias* occurs if participants take longer to abandon the earlier routine than one should expect based on a normative model of probability updating.

We test our participants in a decision situation in which past outcomes do not predict future outcomes. Hence, the profit maximizing reaction would be to treat the second phase of the experi-

ment as a separate decision situation and use Bayes' rule to identify the new best deck. However, our participants are not informed that there is a change at all, nor are they aware that information on past outcomes is irrelevant. An alternative benchmark, therefore, begins with the legitimate assumption that the situation has not changed. Based on this assumption, if participants choose the former best deck and incur a loss, they must integrate this new information with the information from the training phase using Bayes' rule.

If one defines Bayesian updating as the norm, the probability of an outcome for the choice at stake is a function of the choice's entire prior outcome distribution (see Betsch, Haberstroh et al., 2001, for a discussion). For example, assume that a particular option produces 8 gains and 2 losses in one sequence of choices and 2 gains and 8 losses in another sequence. Considered separately, $p_{\text{SEQUENCE 1}}(\text{gain})$ is .8, whereas $p_{\text{SEQUENCE 2}}(\text{gain})$ is .2. Now assume that sequence 2 is experienced as a continuation of sequence 1. At the start of the second sequence, the expectation of a gain is .8 due to the prior outcome distribution in sequence 1. If the individual then experiences a loss as the first outcome of the first choice in sequence 2, the base rate of losses changes from 2 out of 10 (20 %) to 3 out of 11 (27.3 %). Hence, it is justified to still expect this alternative to produce a gain in the next trial in the sequence. Due to incremental base-rate updating, the loss probability exceeds the threshold of 50% only after the 6th move in the second sequence provided each move consistently yields a loss. Accordingly, maintaining the routine over 6 trials in the second sequence (even after consistent negative feedback) would be perfectly justified if individuals assume that the situation has not changed in the second round of the game. By contrast, if decision makers assume that the situation has changed and that the past is uninformative, they must begin constructing a new base rate from scratch.

Within the Bayesian framework and under the assumption that all players are playing the same game, additional information should increase the speed of adaptation. Accordingly, if the sample of valid information is increased by allowing the individual to inspect choices and outcomes of another person, the routine effect should decrease. This effect of additional information should occur regardless of whether the information is conveyed in a social or non-social fashion, given that the information is valid and reliable.

The literature on social learning, however, gives rise to the expectation that a social model may have a stronger impact on an observer's behavior than a non-social model. Identification with a model plays a crucial role in social learning (Bandura, 1977). Specifically, the tendency to imitate a model increases with identification. Presumably, it is easier for participants to identify with another person (in the case of our study: a person of the same sex and similar age) than with a computer. Therefore, we should expect individuals to attribute more weight to socially conveyed information than to non-socially conveyed information. In our study, the additional information provided by social and non-social models suggests routine deviation. But the decrease of a routine effect should be most pronounced in the social model conditions.

We still know little, however, about the *nature of the adaptation process*. Different mechanisms may drive the information search and integration processes as well as the resulting choice. First,

consider (as a normative benchmark) a *homo oeconomicus* who employs a Bayesian updating rule for probability estimation and takes all information into account. Each piece of new evidence is aggregated with prior experience. An individual following an updating rule will maintain the routine until counterevidence compensates prior success rates. The sample of counterevidence needed for adaptation after the world has changed increases with the size of the sample of prior experience reinforcing the routine. As such, the strength of a routine effect is a function of the amount of both prior experience and new information. The adaptation process is slow because it requires step-by-step updating of prior knowledge. Observing an (ideal) model (i.e., showing the correct behavior) should speed up this process because the model *increases the sample of information* used for updating (cf. Yechiam et al., 2008). If the weight of each piece of information remains stable in this process, choices should reveal routine effects but no routine biases.

A second, more extreme possibility is that decision makers apply conservative strategies once a routine has been established. Similar effects of inertia in decision making have been described in the literature on the “Einstellungs-Effekt” (Luchins & Luchins, 1959), status quo bias (Samuelson & Zeckhauser, 1988), omission bias (Baron & Ritov, 1994) and the strategy maintenance effect (Bröder & Schiffer, 2006). The most problematic routine biases are those that occur against the individual’s own intention. Even less than 10 choice repetitions can be sufficient to cause counter-intentional routine maintenance effects (relapse errors) in subsequent decisions (Betsch et al., 2004). Routine bias can evolve from overweighting prior base rates or from confirmatory search strategies (Betsch, 2005; Betsch et al., 2001). In the most radical form, routine decision makers might entirely reject evidence counter to the routine. From the viewpoint of such a conservative decision maker, a model’s behavior would simply be irrelevant.

Browsing the literature on social modeling, however, invites a more optimistic view. Assume, as a third possibility, that individuals apply a very simple copy strategy vis-à-vis a good decision model. Adopting the model’s behavior will free individuals from the effortful processes of belief updating and information integration. Research on humans and animals indicate that organisms are prone to use copy strategies, especially if information acquisition and processing is costly (Boyd & Richerson, 1985; Laland, 2004; Webster & Laland, 2008). A copy strategy enables a decision maker to jump to new courses of actions and instantly adapt to changes in the world when a good model is present.

The possibilities discussed so far lie at the endpoints of a spectrum. The first and the second assume strong inertia effects. Adaptation may completely fail under extreme conservatism or involves effortful relearning. This process is slow and may even fail if information processing is biased. A model can promote this process only by increasing the sample size of relevant information (and this may also have no effect if the decision maker ignores this information). This adaptation strategy would manifest itself in strong routine effects or even routine biases. In contrast, it is conceivable that individuals will reveal no routine effects at all because they could simply copy the behavior of the successful model.

We suggest a fourth possibility that lies between these two extremes. The key idea is that observing a behavioral model in a changing world causes decision makers to experience a conflict. After the world has changed, the individuals suddenly realize that their own routine deviates from the model's behavior. This divergence might be considered a *cue for change* of the situation. Cues indicating novelty have a powerful impact on routine choice. They encourage the consideration of alternatives and, eventually, routine deviation (Betsch et al., 1998, 2001). If such a conflict is present - or if at least the decision makers cannot exclude this possibility - they must decide on two levels. At the behavioral level, the individuals still select between behavioral options (e.g., the decks of cards in our experiments). On a meta-level, the individuals consider different hypotheses about the situation, "the world has changed" versus the "world is still the same". Depending on which hypothesis is accepted, different types of information become relevant. If the first is accepted, only new information is relevant. If the second is accepted, prior information remains relevant. As such, the meta-decision will influence the selection of the sample of information to be considered. In this perspective, the model primarily serves as a cue for novelty, rather than a mere source of additional outcome information. If this is the approach used by all or at least the majority of participants, their speed of adaptation should lie somewhere between the extremes discussed above.

In the next sections, we report two experiments that examine routine decision makers' process of adaptation to changes in the pay-off structure. In a learning phase, all participants establish a choice routine, i.e., to prefer a certain lottery (deck of cards), during a series of routine-learning trials. In a test phase (routine-deviation phase), the task pay-off structure is changed so that it is normatively appropriate to abandon the routine and prefer a different option. In the first experiment, we study routine maintenance and routine deviation decisions under the presence (absence) of a social model (i.e. a confederate of the experimenter) who provides the participant with additional relevant and valid information about choices and outcomes. To avoid gender effects, we recruited only female participants to work with the female experimenter and the female confederate. In the second experiment, the additional information is conveyed in a non-social fashion (i.e. by a card information box that contains draws made by a computer). We assess the influence of prior information and new evidence in order to identify the mechanisms of information processing that underlie routine deviation decisions and adaptation to new environments.

EXPERIMENT 1

Method

Participants and Design

Fifty-four female undergraduates (mean age: 20.7) from different majors at the University of Erfurt participated in the experiment, which lasted approximately one hour. They were randomly assigned to one of four experimental conditions resulting from a 2 (model observable / not observable during routine-learning phase) by 2 (model observable / not observable during routine-

deviation (test) phase) full-factorial design. Based on individual performance, the six best players were rewarded with 30 Euros each (approximately 42 US dollars at that time). The reward was quite attractive to the students, as it was almost five times the typical hourly pay for a research assistant.

Materials

We employed a variant of the Iowa Gambling Task (Bechara et al., 1994), in which participants repeatedly chose between four decks of cards yielding monetary losses and gains. The decks of cards differed with regard to their pay-off distribution. In our version of the task, the outcome structure of the decks was arranged in such a fashion that one deck clearly dominated the others. Specifically, one deck had a positive expected value $EV=5$, while the EVs of the other decks were either negative or zero (see Table 1). Moreover, the range of pay-offs was larger in the two decks with negative EVs. The pay-offs were represented by poker chips. The pay-off structure encourages the decision makers to develop a routine. That is, after having identified the best alternative, decision makers should henceforth repeatedly prefer that particular deck in order to maximize their profits over time. Accordingly, we assessed this tendency towards routinization by comparing rates of draws from the four decks. The absence of routines should manifest itself in an even distribution of rates (25% of draws from each deck). With increasing routinization, the dispersion of rates would increase, revealing a preference order (high rate of draws from one deck and low rates of draws from the others).

In order to increase option distinctiveness, the decks had different colors. Combinations of deck EV's and color were counter-balanced in order to avoid systematic effects of prior preferences for color. We covertly changed the assignment of colors to the type of decks from the learning phase (routine learning) to the test phase (routine deviation, see below). With this change, the routine (e.g., the green deck) suddenly became a poor alternative. In order to continue maximizing profits, participants had to deviate from their routines in the test phase (e.g., stop choosing the green deck) and choose a deck with a different color (e.g., the red deck, see Table 1).

Table 1: Payoff Structure and Appearance of the Card Decks in the Learning and Test Phase of Experiment 1 and Experiment 2

Color of Deck* Learning Phase	green	red	blue	yellow
Color of Deck Test Phase	red	green	yellow	blue
Range of Gain Values	5; 10	5; 10	5; 10	5; 10
Range of Loss Values	5; 10	5; 10	10; 15	15; 20
Gain Probability	80%	50%	40%	50%
Expected Value	5	0	-5	-5

* Table shows examples of colors. Within each experimental condition, colors were randomly assigned to the decks

Figure 1 depicts the experimental setting. Two sets of four decks of cards (backside up) were laid out on a table; one set in front of the participant (P), the other in front of the confederate (C) who served as the behavioral model. Both sets were identical with respect to color and outcome distribution. Across from the experimenter (E), P and C were seated next to each other (distance approximately 1 m apart) with a separating screen between them. In the experimental conditions in which the model's behavior was observable, the screen was shifted backwards so that the participant could see the model's choices and the subsequent outcomes (see right part of Figure 2). In the counter conditions, the screen was moved forward so that the model's moves were hidden (see left part of Figure 1). Participants were told that the position of the screen might be changed within the experiment in order to study different conditions of team performance. A female experimenter served as the croupier. Specifically, she handled the bank, guided participants through the game, logged the outcomes, collected losses and returned gains (poker chips).




Figure 1: The experimental setting used in the “model not observable” conditions (left side) and “model observable” conditions (right side). E = Experimenter, P = Participant, C = Confederate (model). P and C each have four decks of cards in front of them. They place drawn cards face-up on the rectangular field marked by the dotted line. Poker chips are piled on the fields marked by circles.

Procedure

Participants were recruited with the understanding to take part in a study on team performance in a card game. Upon arriving in a waiting room, the participants met their team player, a confederate of the experimenter, who served as the behavioral model (female, age 22 years, student of another university, blind to the hypotheses). Participants were informed that the team's task was to gain as many casino chips as possible. After the study, the six best performing teams would be identified (representing 12 % of all participating teams) and each of the team players rewarded with € 30.

The team situation establishes a decision task in a social environment in which inspection of a model's behavior is justified. One may argue that it may also demand the participant to copy the model's behavior. Note, however, that such a tendency would work against our hypothesis of delayed adaptation.

After signing an informed consent statement, the participant and the confederate were led to the experiment room. The participant was always seated to the left of the confederate. Depending on assignment to experimental conditions, the participant could either observe the model's behavior or not (see Figure 1).

The experimenter introduced herself as the croupier and provided each team player with 100 poker chips. She then described the setting (see Figure 1) and the procedure of the game. Specifically, team players were informed about the following rules of the task:

“The team players repeatedly draw cards. They alternate whose turn it is to draw. A player must wait her turn until it is announced by the croupier. Thereafter, the player must draw one card from one of the four decks in front of her. She is free to choose any of the four decks. After she has drawn a card, she must turn it face-up and place it on a rectangular field marked on the desk in front of her. Each card announces either a gain or a loss in terms of positive and negative points, respectively. The croupier inspects the payoff, logs it into a file and pays/collects the corresponding number of chips. The team goal is to jointly maximize the total number of points and, at the same time, minimize the number of draws. One deck of cards is better than all the others. To maximize profits, team players must identify this deck as quickly as possible. The croupier adds up the payoffs achieved by the participant and will terminate the game as soon as the sum of points exceeds a specific amount¹. Team players are not allowed to verbally communicate with each other during the game.”

After having introduced the rules, the experimenter initiated the learning phase of the game in which the participant acquired a routine (i.e., a preference for one particular deck of cards). In each trial, first the learning model (confederate) drew a card and placed it face-up on the table. Subsequently, the participant was asked to make her draw. The model and the participant stacked the drawn cards in front of them so that only the last card drawn was visible. The experimenter terminated the learning phase as soon as the *participant* had gained 250 points. To reach this amount, the participant must have chosen the best deck at least 30 times. The minimal number of routine repetitions is equal to those used in the studies by Betsch and colleagues (2001, strong routine conditions). Moreover, the game was terminated if the participant failed to reach the critical number of points within 100 trials. In the case that the participant initially preferred a suboptimal deck, the number of trials needed was increased to compensate accumulated losses. The number of critical points and trials were determined based on careful pretesting and simulations of drawing strategies. The chosen combination ensures that participants passing these requirements have repeated the routine at least three times more often than any of the other alterna-

1 An exact amount could not be announced in advance, because it was dependent on the participant's performance.

tives (cf. Betsch et al., 1998). Six participants failed to gain 250 points within 100 trials (see result section).

After having successfully completed the learning phase of the experiment, the team was led to two separate cubicles where the participant was presented with a short questionnaire on a computer screen. The questionnaire assessed personal information (gender, age, academic majors, etc.); none of the items were related to the decision task. Participants were informed in advance that the decision tasks will be interrupted by a short questionnaire session. Meanwhile, the experimenter exchanged the card decks. Color, position and height of the decks remained the same but the pay-off distributions were changed so that the routine deck yielded an $EV = 0$ (see Table 1). Thus, in subsequent trials players must deviate from their routine and identify the new best performing deck in order to continue maximizing their profits.

After completing the questionnaire, the team returned to their table to begin the test phase of the experiment. The experimenter reminded the team of the ostensible purpose of the experiment to study different conditions of team performance. Depending on experimental conditions, she either changed the position of the screen between the players or left it in the original position. Following a brief rehearsal of the rules of the game, she asked the confederate to begin with the first draw. The participant received no information about the changes in the payoff structure. After 50 trials, the game was terminated.

In each condition, the model's moves followed a fixed schedule. At the beginning of the learning and test phases, the model began by drawing a sample of three cards from each deck starting with the left deck. Thereafter, she preferred the best deck in about 90% of her choices.

After completing the entire experiment, all participants received a written debriefing via email that described in detail the purpose of the study as well as the reasons why it was necessary to use a confederate as a behavioral model. In each of the two routine-learning conditions (model observable vs. not observable), the three individuals who performed best during the learning phase were identified. Results of test-phase performance were ignored. Altogether, six participants were invited to the lab and rewarded with € 30 each.

Results and Discussion

Six participants failed to successfully complete the learning phase. Specifically, they did not reach 250 points within 100 trials and were thus excluded from further analysis. The remaining 48 participants are evenly distributed across the four experimental conditions ($n = 12$ per group).

We analyzed participants' choice behavior in the test (routine-deviation) phase of the study. Figure 2 descriptively shows how the presence or absence of a behavioral model influences participants' adaptation to changed circumstances. When there was no model in both the learning and test phases ("no no" condition), more than 80 % of participants chose the former best deck in the first trials. However, this percentage steadily declines while the percentage of those choosing the

new best deck steadily increases over trials. By the 20th trial, the new best deck is chosen more often than the old one. This pattern is very similar to that observed when a behavioral model is always present (“yes yes” condition). In this case, however, adaptation occurs much more quickly. The new deck is already chosen more frequently than the old one by the 10th trial. Choices also polarize over time. By the 40th trial, almost all participants chose the new deck. When a model was present during the learning phase and absent during the test phase (“yes no” condition), the data are very noisy until trial 35. Even then, the new deck is preferred only slightly over the former one. Finally, when the model was not present during the learning phase and present during the test phase (“no yes” condition), the data fluctuate widely until trial 35, after which the choices polarize.




Figure 2: The figure shows adaptation during the test (routine deviation) phase in the four experimental conditions of Experiment 1: No learning model / no test model (“no no”), no learning model / test model (“no yes”), learning model / test model (“yes yes”), learning model / no test model (yes no). For each trial of the test phase, dots indicate the percentage of adaptive choices (new normative deck), while squares indicate the percentage of maladaptive routine choices (former normative deck).

In a non-parametric test of the mean percentage of normative choices per participant in the test phase, we find significant differences between the “no no” and “yes yes” conditions (Wilcoxon rank sum test, $N = 24, p = .0492$) as well as between the “yes no” and “yes yes” conditions ($N = 24, p = .0301$). The remaining comparisons are not significant at conventional levels. For parametric estimation, we face two challenges. First, participants could choose between four decks; the expected values of the two least attractive decks are equal ($EV = -5$). Second, participants had to infer expected values from the information that they happened to receive. For both reasons, we must treat the four options as categorical. The appropriate model for analysis, therefore, is a mul-

tinomial logit model². Each participant had 50 choices; and observations are not independent within participants. Unfortunately, there is no generally acknowledged random-effects multinomial logit model. Therefore, we clustered standard errors per individual. Note that this procedure assumes less structure than is actually contained in our data and is therefore conservative.

Both models in Table 2 use the deck that was normative in the learning phase (the routine) as the reference category. Thus, all regressors measure the likelihood that a participant deviates from the former norm in the test phase. In terms of experimental conditions, the “no no” condition is the reference category. Hence, all treatment effects measure how strongly the presence of a behavioral model in either phase affects the likelihood of deviating from the former normative choice. In model 1, all treatment effects (except for “yes no” with respect to the new normative choice) are significant. Whenever there is a behavioral model, participants are less likely to choose the two options with a negative expected value, compared with the “no no” condition. By contrast, if participants have the opportunity to observe a behavioral model in the test phase, there is a greater likelihood that they will choose the new normative deck.

Model 2 is even more revealing. It adds the time trend as well as interaction effects with time. Results show that participants are less likely to choose the two decks with negative expected values and more likely to choose the new normative deck. All treatment effects diminish over time. This follows from the fact that the condition and interaction effects have opposite signs. Interestingly, in this model for the normative choice, condition main effects are negative. In early trials, participants in all conditions are less likely to make the normative choice, as compared with the “no no” baseline. However, for the “no yes” treatment, the effect reverses due to the interaction effect from the 24th trial on ($23.846 * .052 = 1.240$). For the “yes yes” treatment, 20 trials suffice for this reversal to occur. Since the model is non-linear, coefficients cannot be directly interpreted as marginal effects. Model 2 predicts that the likelihood of making the normative choice increases from the first to the last trial of the test phase from 26 to 69 %, respectively, in the “no no” baseline. This likelihood increases from 19 to 88 % for the “no yes” condition, from 15 to 95 % for the “yes yes” condition and from 23 to 62% for the “yes no” condition.

2 Technically, we estimate
$$pr(choice_j) = \frac{\exp(x_{ij}\beta_j)}{1 + \exp(x_{i2}\beta_2) + \dots + \exp(x_{iM}\beta_M)}, j = 1..M$$
. The model simultaneously estimates all three contrasts with the reference category (which we chose to be the former normative deck) with maximum likelihood and thereby reflects that subjects only once choose between all four options. For background, see Cameron & Trivedi (2005, chapter 15).

Table 2: Results of a Multinomial Logistic Regression on Choices in the Test Phase of Experiment 1.

		model 1	model 2
Deck EV=-5 (gain probability = .40)			
	no yes	-2.222***	-2.451***
	yes yes	-1.731***	-2.135***
	yes no	-1.671***	-2.114***
	trial		-.057***
	no yes*trial		.069**
	yes yes*trial		.081***
	yes no*trial		.076***
Deck EV=-5 (gain probability = .50)			
	no yes	-1.552***	-1.579***
	yes yes	-1.764***	-1.883***
	yes no	-1.570***	-1.399***
	trial		-.052***
	no yes*trial		.054**
	yes yes*trial		.060**
	yes no*trial		.044**
Deck EV= 5 (gain probability = .80)			
	no yes	.489**	-1.240***
	yes yes	.732**	-1.619***
	yes no	.004	-.943*
	trial		.019 ⁺
	no yes*trial		.052***
	yes yes*trial		.083***
	yes no*trial		.018
N		2394	2394
chi ²		329.23	1358.10

Note: Standard errors clustered for 48 groups (i.e. participants). *** p < .001, ** p < .01, * p < .05, ⁺ p < .1

Routine Effect versus Routine Bias

To disentangle routine bias from routine effect, we calculate the posterior probability of the previously normative choice yielding a gain using Bayes' rule based on the assumption that the game has not changed. We calculate this in the following manner: in the learning phase, participants make decisions and know the color of the deck. For the analysis, we focus on the color of the deck with the highest expected value in the learning phase; c denotes the number of times that the participant chose a certain deck (e.g., the red deck) in the entire experiment, i.e. in the learning and in the test phase, up to the current trial. Through feedback participants learn whether they have made a gain in the trial in question. g denotes the number of times that a certain

participant has made a gain up to the trial in question. Participants are also able to keep track of the number of decks with the previously normative color that have led to gains. We denote this number by cg . Finally, participants know the number of the current trial in terms of the entire experiment, i.e. treating the time after the restart as a continuation of the previous experiment. This number is denoted as n . Given this information, we can calculate the posterior probability of making a gain for each trial, given that the deck has the previously normative color:

$$p(\text{gain} | \text{color}) = \frac{\frac{cg}{g}}{\frac{c}{n}} = \frac{cg}{c}$$

The upper lines in all panels are generated in the following manner: for each participant, we calculate the posterior probability that the previously normative deck is still best using the above equation. The lines in the graphs represent the means of these individual posterior probabilities. Figure 3 shows that, for the previously normative choice (the routine), the posterior probability remains high throughout the game. Even in the final trial of the test phase, i.e. 50 trials after the change in circumstances, the mean updated probability of success is still 70.32 %. Even for the participant who has experienced the most losses with the former best deck, this probability is still 50.94 %. Given that each participant has four options, the posterior probability is thus still far greater than chance (which would be 25 %). Hence, if participants were blind to the possibility of change, all should have continued choosing the deck that they had learned to be best in the routine acquisition phase. Only one participant actually showed this extreme routine behavior. Accordingly, from a normative view, we find no evidence for a general routine bias among participants.




Figure 3: The figure shows Bayesian updating during the test (routine deviation) phase in the four experimental conditions of Experiment 1: No learning model / no test model (“no no”), no learning model / test model (“no yes”), learning model / test model (“yes yes”), learning model / no test model (yes no”). Dots indicate the mean posterior probabilities for the former normative deck (routine) using both learning and test phase information for Bayesian updating. Triangles indicate the mean posterior probabilities for the new normative deck using only test phase information.

In contrast, it is also possible that participants wonder whether the situation has changed in the second decision phase. Recall, that there was a short break between the two phases. This interruption might cue concepts of change or restart in the participants. In the extreme case, participants might believe that the past is completely irrelevant (which happened to be our manipulation, but this was unknown to the participants). In this case, they must begin exploring the situation from scratch. In other words, they should not even show a routine effect. If we allow for some uncertainty in their reactions and count those participants who chose each of the four decks at least once during the first 8 trials of the test phase, one quarter of participants, i.e. 16 out of 48, did in fact behave in this manner. Note that, in the context of this experiment, such variance cannot be explained by participants merely guessing or choosing decks at random, as we only admitted participants to the second part of the experiment who had acquired the routine to choose the previously normative deck. If participants deviate from this routine in the test phase, it must mean that they question whether the situation is still the same. However, Bayesian analysis is more precise. Both lines in Figure 3 were calculated in the same manner; the lower line, however, refers to the newly normative choice and is based only on information sampled in the test phase, while the upper line refers to the routine and uses information from both phases. In the aggregate, this posterior probability is greater than 50% from the 10th trial on. Given that the new

normative deck yields a gain in 80% of all cases, the posterior probability climbs to 100% in the first trial for each individual in four of five cases. This probability then stabilizes at a slightly lower level once the participant experiences one of the rare losses. Participants are told that only one deck is optimal. Therefore, if the posterior probability of a gain is greater than 50%, one should expect that, from then on, individuals should choose this deck with certainty.

The large majority of participants, however, are substantially below this benchmark, as Figure 4 shows.




Figure 4: Sensitivity to gain from new normative deck, provided posterior probability is above p=.5

This indicates that these participants chose non-normative decks, even if the posterior probability of gaining from choosing the normative deck is greater than 50% (and hence more than twice as likely as chance). Our findings thus far show that our participants neither blindly follow the routine from the training phase nor completely neglect the information from the past and explore the situation from scratch. The pattern of findings suggests that most participants adopt a richer decision strategy. Arguably, most participants are uncertain whether the situation has changed or not. If the situation is unchanged, the fact that they now face more losses when choosing the previously normative deck is new information regarding an old phenomenon. They should therefore integrate it with the information from earlier rounds. If, however, the situation has changed since the first phase of the experiment, participants should begin searching and subsequently update their predictions in light of their new experiences. For such participants, one option would be to first try out both interpretations of the new situation in competition with each other and then switch to the new normative deck (with high gain probability) only once the predicted probabil-

ity of gaining based on the assumption that the situation has changed is greater than the predicted probability of gaining based on the alternative assumption that the situation is the same.

Adaptation and the Role of the Behavioral Model

If participants are uncertain about how to interpret the situation, they must decide between two hypotheses: (1) the situation is unchanged (routine applies to test phase) or (2) the situation has changed (routine does not apply to test phase). Each hypothesis invites different strategies of information processing. Residing with the first hypothesis, the individual should use information from both phases and assess the posterior probability of the previously normative deck still being the best choice (*upall*). Residing with the second hypothesis, the individual should ignore information from the first phase of the experiment. Given this strategy, the individual would have to estimate the winning probability for each option and update each estimate in light of the experiences she is making. For the analysis, we focus on the evolution of the posterior probability that the newly normative deck is indeed the best given only feedback from the testphase (*uptest*).

We first investigate the process of adaptation in the aggregate. We use Bayesian predictions based on both conceptions of the situation and their corresponding strategies (*upall*, *uptest*) as explanatory variables. Model 1 in Table 3 shows that both strategies indeed explain choices independently. In terms of the decision to avoid the decks with a negative expected value, only the updated probability of making a gain with the previously normative deck matters. Participants had learned during the training phase that these decks are bad; this information remains valid. By contrast, choosing the new normative deck is significantly and considerably more likely the more important the updated probability of gaining from this option is and less likely the higher the updated probability of gaining from choosing the previously normative deck is. Coefficients indicate that participants attach more weight to the possibility that the game has changed. However, the weight attached to the possibility that the game has remained stable still makes up approximately 2/3 of the weight attached to the opposite interpretation of the situation.

Model 2 in Table 3 shows two things. First, controlling for predictions does not affect condition effects. More interestingly, we learn how the presence of a behavioral model affects choices. If there is such a model in the test phase, participants are much less likely to interpret the situation as unchanged due to the condition main effects. However, as indicated by the negative interaction effects, the pull of the model becomes weaker the stronger the updated prediction of gaining from choosing the previously normative deck is. By contrast, there are no significant interactions between conditions and the updated probability of gaining from the newly normative deck. We conclude that *a behavioral model is not important for increasing the amount of information about the newly normative deck*. The key issue, therefore, is not that of providing the participant with a larger sample. Rather, she must be induced to see the situation in a new light. The fact that the model no longer follows the old norm provides a powerful cue and induces participants to explore the situation more closely.

Table 3: Results of a Multinomial Logistic Regression on Choices in the Test Phase of Experiment 1 Using Different Updating Strategies as Predictors.

	model 1	model 2
Deck EV=-5 (gain probability = .40)		
upall	-3.150***	-3.341***
uptest	0.890*	1.066
no yes		2.213
yes yes		6.339
yes no		-6.973**
no yes*upall		-3.256
yes yes*upall		-7.792
yes no*upall		9.695**
no yes*uptest		-0.233
yes yes*uptest		-0.518
yes no*uptest		0.097
Deck EV=-5 (gain probability = .50)		
upall	-2.326***	-2.126***
uptest	0.154	-0.575
no yes		-2.205
yes yes		0.755
yes no		-4.728*
no yes*upall		2.604
yes yes*upall		-1.429
yes no*upall		6.091*
no yes*uptest		1.099
yes yes*uptest		0.949
yes no*uptest		1.107
Deck EV= 5 (gain probability = .80)		
upall	-2.026***	-2.481***
uptest	2.808***	2.966**
no yes		8.494**
yes yes		13.569*
yes no		-1.393
no yes*upall		-10.668***
yes yes*upall		-17.583*
yes no*upall		2.815
no yes*uptest		-0.424
yes yes*uptest		0.122
yes no*uptest		-1.069
N	2394	2394
chi ²	349.88	3805.35

Note: Standard errors clustered for 48 groups (i.e. participants). *** p < .001, ** p < .01, * p < .05, + p < .1. The strategy *upall* uses information from the learning and test phases for updating the probability that the old normative choice remains best ; *uptest* only uses information from the test phase and updates the probability that the new objective norm is indeed the best.

By testing each participant individually on model 1 in Table 3³, we gain a number of additional insights. We count a participant as relying on the old routine if the regressor *upall* explains her choice of the new normative deck after the break with $p < .05$. We count a participant as exploring the situation from scratch if the regressor *uptest* explains her choice of the new normative deck with $p < .05$. We count the participant as using both strategies in conflict if both regressors explain her choice of the new normative deck with $p < .05$.

When a behavioral model is present in both phases, only 2 of 12 participants cannot be classified as either following their earlier routine (1 participant), exclusively exploring the new situation (5 participants) or using both approaches simultaneously (4 participants). Nine of 12 participants are influenced in their choice of the newly normative deck by an interpretation of the situation as new. Thus, the presence of a behavioral model helps participants make meaningful choices and induces participants to consider the possibility that the situation has changed.

When there is no behavioral model in either phase, the adaptation strategy cannot be identified for half of the participants. The other half (i.e. 6 participants) simultaneously use both strategies. If they do not have help from the social context, participants have a more difficult time making meaningful choices. Those who do make choices rightly seem to consider the two competing interpretations: the game continues or the game has changed. In the “no yes” condition, participants are simultaneously faced with two sources of uncertainty after the restart: the game might have changed; and they must adjust to the fact that they can now observe their counterpart.

While participants in the “yes yes” condition knew the other group member had performed very well in the training phase, no such information was available to participants in the “no yes” condition. Indeed, participants perform less well in this condition. Seven of 12 participants cannot be classified as using either strategy. Three focus exclusively on the novelty, while two are significantly influenced by both the *uptest* and *upall* regressors.

Finally, in the “yes no” condition, participants not only have to determine whether the game is still the same after the break, they are also deprived of the help by their group member, which they had observed to be instrumental during the training phase. This heightened degree of uncertainty translates itself into the fact that only 4 of 12 participants can be classified as using the updated information. Only two interpret the situation as changed, while an additional two seem to simultaneously consider both interpretations.

EXPERIMENT 2

In Experiment 1, the participant’s choice problem was embedded in a social context. In three of the four conditions, the participant also had the benefit of observing the confederate’s choices in

3 We count a participant as relying on the old routine if regressor *upall* explains her choice of the new normative deck after the break with $p < .05$. We count a participant as exploring the situation from scratch if regressor *uptest* explains her choice of the new normative deck with $p < .05$. We count the participant as using both strategies in conflict if both regressors explain her choice of the new normative deck with $p < .05$.

either one or both phases. In the “no no” condition, participants only knew that another subject was solving the same task and that performance would be assessed for both of them together. We found that participants were most likely to adaptively deviate from their routine if they could observe the social model in both phases. Our analysis gave rise to the interpretation that this effect was *not* caused by an increased sample of information provided by the model. Rather, the behavior of the social model may have alerted the participant to the possibility that the world might have changed. As such, the behavior of a model may provide a powerful cue about the nature of the situation. If individuals systematically deviate from a prior norm (here a shared behavioral routine), they can create what Jones and Davis (1965) termed a situation of non-common effects. According to their theory of correspondent inference, norm deviation is a powerful cue for intentionality and evokes increased deliberate activity in the observer. The observer might begin to wonder why the model is behaving differently and, in our case, might consider the hypothesis that the world has changed. Now assume that participants observe a non-social model, i.e. the behavior of a computer. We regularly do not attribute insight, purpose and intentional behavior to computers. Therefore, a non-social model should be less likely than a social model to alert observers to changes in the decision situation. If our conclusions from the first study are valid, we should find that the impact of the model’s behavior decreases when the model is a non-social model. In contrast, if our previous findings were caused by a mere sample size effect, we should be able to replicate the previous pattern of results in a non-social context.

To test these hypotheses against each other, we replicate the previous study in a setting that *removes all social aspects* of the decision task in order to create a condition in which merely the additional information is present. If our findings were due to a mere sample size effect (model provides additional information) then results should replicate. If, however, the social situation (working together with a human) has caused the pattern of effects, we expect our findings not to replicate.

In Experiment 2, we maintain all essential structural aspects of the task, the procedure and the design. The participants in Experiment 2 have access to information drawn by a computer. This information is provided to them by a “card information box”. The box is placed next to the participant, as the model was in the prior study. Draws from this box convey exactly the same information that the social model provided in Experiment 1.

Method

Participants and Design

Fifty-eight female undergraduates (mean age: 20.4) from different majors at the University of Erfurt participated in the experiment, which lasted approximately one hour. They were randomly assigned to one of four experimental conditions resulting from a 2 (model observable / not observable during routine-learning phase) by 2 (model observable / not observable during routine-

deviation (test) phase) full-factorial design. Based on individual performance, the six best players were rewarded with 30 € each (approximately 42 US dollars at that time).

Materials

We employed the same card drawing task and the same pay-off distribution as in Experiment 1 (see Table 1).

Setting

The experimental setting was identical to that in Experiment 1, except that there was no human model. We used a so-called card information box (“Karteninformationsbox” in the German instructions) to provide participants with additional information in a non-social context. For this purpose, the instructions referring to a team-situation were removed in order to avoid any associations related to social issues. The box was placed where the model was seated in the prior study. It contained a stack of cards of different colors representing a sequence of draws. The sequence of draws and pay-offs were exactly the same as those of the model in the former experiment. In the conditions “model observable”, the croupier drew a card from the stack contained in the box, turned it upside down and told the participant to inspect its outcome. The experimenter then added (removed) the corresponding number of poker chips to (from) the pile in front the box. As such, the presentation of additional information in this study was identical to that in Experiment 1. In the conditions without access to additional information, no cards were drawn from the box.

Procedure

Participants were recruited with the understanding to take part in a study on individual performance in a card game. All procedural steps relating to instructions about the nature of the game, its rules and its parameters (initial endowment, termination criteria, etc.) were identical to Experiment 1. Instructions relating to social aspects, however, were removed and replaced, when necessary, with instructions relating to the card information box and its usage.

In the conditions “additional information observable”, participants were read the following instruction: “During this game you have access to additional information. Based on computer simulations, we determined different variants of successful sequences of draws in this card game. Three sequences consisting of yellow, red, green and blue cards were determined by the computer.”

Participants were presented with three boxes (labeled with the numbers 1, 2, 3) that contained these three sequences. They were asked to select one by throwing a die. The selected box was then placed on the desk beside the participant. The participants were also informed that pay-offs achieved by draws from the box were only for the purpose of providing additional information and would not be added to (or subtracted from) the participant’s yield. The experimenter then drew the first card from the box. After allotting chips based on the outcome, the experimenter

asked the participant to make her own first draw. This alternate drawing procedure was continued until the experimenter announced the end of the first (second) phase of the experiment.

As in Experiment 1, the experimenter terminated the learning phase as soon as the participant had gained 250 points or if she failed to reach the critical number of points within 100 trials (these participants were excluded). Participants who successfully completed the learning phase were led to a cubicle where they were presented with the same type of questionnaire used in experiment 1 (questions relating to gender, age, academic majors, etc.).

Thereafter, the experimenter began the test phase in which the pay-off distribution was changed unbeknownst to the participants. Participants were reminded of the rules of the game and asked to play another round. In the condition “additional information not observable during routine-deviation phase” no cards were drawn from the box. In the condition in which additional information was only observable in the deviation phase, the box was introduced with the same instruction as reported above. The moves and pay-offs produced by the box were the same as those of the model in the deviation phase of Experiment 1. After 50 trials, the game was terminated. The reward procedure was the same as in Experiment 1.

Results and Discussion

Eight participants failed to successfully complete the learning phase. Specifically, they did not reach 250 points within 100 trials and were thus excluded from all analyses. The remaining 50 participants are almost evenly distributed across the four experimental conditions ($n = 12-13$ per group). Analyses are basically the same as in Experiment 1. Comparing Figure 5 (Exp. 2) with Figure 2 (Exp. 1), one sees that non-social models are considerably less effective in inviting routine deviation as compared to human models: even in late trials, the percentage of participants in all conditions choosing the deck that was normative during the learning phase remains similar to the percentage choosing the new best deck. In Experiment 1, the newly normative deck clearly outperformed the former best deck in later trials, especially in the “no no” and “yes yes” conditions. Moreover, it takes longer for the new norm to gain more support than the old norm when the card information box is used – again this is most visible in “no no” and “yes yes” conditions. In contrast to Experiment 1, participants in Experiment 2 are tempted to maintain their routines over test trials.




Figure 5: The figure shows adaptation during the test (routine deviation) phase in the four experimental conditions of Experiment 2: No learning model / no test model (“no no”), no learning model / test model (“no yes”), learning model / test model (“yes yes”), learning model / no test model (yes no). For each trial of the test phase, dots indicate the percentage of adaptive choices (new normative deck), while squares indicate percentage of maladaptive routine choices (former normative deck).

Routine Effect versus Routine Bias

To disentangle routine biases from routine effects, we calculate the posterior probability of the previously normative choice yielding a gain using Bayes’ rule based on the assumption that the game has not changed (as in Experiment 1). As in the prior study, the posterior probability for the previously normative choice remains high throughout the game (see upper lines in all sections of Figure 6). Therefore, if participants were blind to the possibility of change, they should have continued to choose the deck that they had learned to be best in the routine acquisition phase. Three participants showed this extreme routine behavior. Another four participants chose the previously normative deck at least 40 times, and an additional five chose it at least 30 times. By contrast, if there was a human model, only a single participant never deviated from the previous norm. No (additional) participant followed the old norm 40 times or more. Only six participants chose the old norm at least 30 times. Compared to experiment 1, a greater number of participants (but still not the majority) were thus sticking to their routine from the training phase. Participants were slightly less inclined to fully explore the (potentially) new situation. While a quarter chose each deck at least once during the first eight trials of the test phase when they had a human partner, only 10 out of 50 (i.e. 20%) participants interacting with a computer did so.

Again, Bayesian updating provides a more precise measure for the strength of routine effects. The lower lines in Figure 6 are calculated in the same manner as the upper lines, except that they

refer only to the newly normative choice and use only information sampled during the test phase. In the aggregate, this posterior probability is consistently greater than 50% from the 28th trial on (trial 10 in Exp. 1)⁴. A further difference to Experiment one is that the two curves cross in only one condition (“no no”). The slow increase in the posterior probabilities for the new normative deck shows that participants were less inclined to consider the new normative deck. Overall, the routine effect in this study is stronger than in Experiment 1. In other words, the non-social model was less effective than the social model in fostering observer adaptation. This finding strongly speaks against the interpretation that the findings of Experiment 1 were due to a mere sample size effect, because participants in both studies had access to exactly the same amount of additional information. Figure 6 also demonstrates that, again, we only have a routine effect, not a routine bias. The updated probability of the old deck still being best remains very high throughout the test phase as well as in all treatments. We do not find inertia, but a much lesser sensitivity to the possibility that the situation might have changed.




Figure 6: The figure shows Bayesian updating during the test (routine deviation) phase in the four experimental conditions of Experiment 2. Dots indicate the mean posterior probabilities for the former normative deck (routine), using both learning and test phase information for Bayesian updating. Triangles indicate mean posterior probabilities for the new normative deck using only test phase information.

4 It is 50.38 % in trial 26, 49.03 % in trial 27, and 52.79 % in trial 28.

Adaptation and the Role of the Behavioral Model

Using the same estimation strategy as in Experiment 1, we investigate the adaptation process in Table 4. Again, given the information from the entire experiment (*upall*), participants are less likely to deviate from the former normative deck the greater the updated probability of gaining from choosing this deck is. In contrast, given only feedback from the trial phase (*uptest*), participants are more likely to choose the new normative deck the greater the updated probability of this being optimal is. We conclude that, in the second experiment, the majority of participants are also torn between the two competing interpretations of the situation: as a mere continuation of the previous experience or as a (somewhat) altered opportunity. Yet in sharp contrast to Experiment 1, we no longer find significant treatment effects once we control for the pull of these two competing interpretations of the situation. Descriptively, the (insignificant) effect of the additional information from the card information box even makes it *less* likely that a participant chooses the new best deck. Seeing the choices from the card box does not help participants understand that the situation has changed. By contrast, while we had no significant interaction effects between *uptest* and conditions in Experiment 1, we now find that having access to the information from the box in both phases of the experiment (*yes yes*) leads to a more rapid adjustment to the new normative choice. Hence, *with a non-human model there is a sample size effect* but no improvement in reinterpreting the situation. In fact, in the “no yes” condition, we even find a significant interaction with updated information from the past that has the “wrong” sign. The more the information indicates that the choice from the past is preferable, the more participants actually switch to the new best deck.

DISCUSSION

Routines are a ubiquitous phenomenon. They help individuals reserve cognitive resources in recurrent decision situations. The ability to form routines provides an advantage in adaptation (Athay & Darley, 1981). However, the world sometimes changes, causing routines to become obsolete. In order to achieve their desires and needs in such situations, decision makers must abandon their routines and pursue alternative courses of action. Still, little is known about the mechanisms underlying routine deviation decisions. In this paper we explored the role of behavioral models for deviation. In many mundane decision contexts, social behavioral models are available and thus provide an important source of information (Bandura, 1977).

Models can impact decision making in several ways. First, they provide decision makers with the opportunity to delegate decisions. Observers can employ a very simple heuristic for choice. That is, they simply copy the behavior of the model (Boyd & Richerson, 2005). Second, models provide additional information about the consequences of behaviors. Observers may witness a model's behavior being reinforced or punished (vicarious reinforcement). In this perspective, models increase the sample size of pay-off information that observers can utilize in their decisions. Third, if there is a discrepancy between the observer's routine and the behavior of the model, this may encourage epistemic thinking in the observer (Jones & Davis, 1965). In the case of routine

decision making, the discrepant behavior of a model may invite observers to assess the situation more carefully. Specifically, observers might consider the hypothesis that the situation is novel (“world has changed”). Perceived novelty is one important factor fostering routine deviation (Betsch et al., 1998, 2001). Discrepant behaviors in models may serve as a cue for novelty. All these aspects should foster adaptive behavior in a decision maker who can observe a model making accurate decisions.

Before evaluating the three potential paths by which models may impact decision making in an observer, we begin the discussion by highlighting two major findings regarding routine effects and the adaptation process. In both studies we obtained strong routine effects but no routine biases. Accordingly, participants considered prior success rates of the routine and subsequent decisions, but they were not immune to new counter-evidence. New evidence had to compensate prior knowledge before individuals decided to abandon their routine. This process of relearning and gradual adaptation is best described by a Bayesian updating process. This result converges with prior findings indicating that in repeated decision making individuals do not neglect base-rates but rather consider both prior knowledge and new evidence in choice (e.g., Betsch et al., 2005a, 2001).

With regard to potential pathways of model effects, we found no systematic evidence for the copy-heuristic or the sample-size account. The mere presence of a good model was not a sufficient condition for improving routine deviation in our participants. Only few participants copied the model from the beginning of the test phase. As already noted, the majority of individuals followed a belief updating strategy that manifested itself in systematic routine effects (delayed adaptation compared to the learning phase). This finding speaks against a general prevalence in individuals to apply a simple copy heuristic when observing a successful model.

Models provide the possibility to learn reinforcement and punishment rates without acting. As such, models provide access to additional information about pay-offs without behavioral costs such as the risk of losses. In the case that the model provides reliable and valid information, it should foster deviation to the extent that the individual can observe this additional information. In our second experiment the draws by the computer clearly showed that the new alternative was better than the routine. Nevertheless, individuals made little use of this information. This finding speaks against a mere sample size explanation of a model effect.

Rather, our findings appear to be compatible with the third explanation - that a discrepancy between the observer’s routine and the behavior of the model may encourage epistemic thinking in the observer (Jones & Davis, 1965). Multinomial logit model analyses are in line with the interpretation that the model may have served as a cue for novelty. Participants who observed a social model appeared to perform two tasks simultaneously during the routine-deviation phase. On the one hand, they continued to update prior beliefs regarding the validity of their routine by using primarily information that they sampled themselves. On the other hand, they appeared to solve a conflict regarding the relevance of information. In a social situation (Exp.1), the model’s behavior was much more likely to instigate a reevaluation of the situation than in the non-social situa-

tion (Exp.2). More participants became suspicious about the situation and began to consider the hypothesis that the routine no longer lead to the best outcome. If this is the case, the prior information from the learning phase (the reinforcement history of the routine) would be irrelevant. Over trials, evidence accumulates in favor of this hypothesis.

According to this interpretation, our participants appeared to be engaged in two levels of analysis. On the level of the lottery decision at stake, they followed an updating strategy by integrating prior information with new evidence. This resulted in a gradual updating of posterior probabilities as well as a gradual adaptation at the choice level. On a meta-level participants may have considered two different interpretations of the situation: “the world has changed” versus “the world is still the same”. After repeatedly observing the social model, however, they appeared to arrive at the conclusion that the world has indeed changed. At the data level, this assumption is manifested in the curvature of individual Bayesian updates. Steepness, especially in the social model conditions, shows that new pay-offs from the test phase receive more weight than is justified by the normative benchmark when taking the entire pay-off history into account. At the behavioral level, this assumption is manifested in the absence of a routine bias in choice. Especially in the social model conditions, almost all participants deviated from their routine earlier than would have been justified based on comprehensive Bayesian updating. Overall, our findings indicate that routine deviation is a complex process that may involve different levels of analysis (Betsch & Held, 2011).

Another important finding was that the model’s impact was most pronounced in a social situation (Exp.1) compared to a non-social situation (Exp.2). Note, however, that at the present we cannot determine what aspect of the social situation was responsible for this effect. In the social situation, the model was a human confederate and an in-group member (team situation). Each of these two aspects or both together may have driven the model effect. In Experiment 2 we had to remove both social aspects to create a situation without social cues. This was necessary to assess the mere informational effects of a model and provide a strong test for the competing sample size hypothesis. Therefore, our experiments do *not* warrant the conclusion that non-human models are generally ineffective. Information provided by non-social models in real world settings (e.g. search engines in the Internet) presumably have a stronger impact on an observer’s choices than in our study.

On the other hand, social models in real world settings may not generally have such a strong impact as in our first experiment. We confronted participants with an “ideal” model. The model’s behavior was adaptive and outcome information provided by the model valid. Half of the participants had the opportunity to learn this in the first phase. The model was similar to the participant with respect to gender, student status and age. Moreover, the human model was said to be a team member. These features certainly do not apply to all real-world settings. For instance, we may sometimes doubt whether our neighbor’s behavior is an ideal choice because we lack confidence in his or her decision competence. We might also neglect potential models because they belong to an out-group or are perceived as being highly dissimilar to ourselves. As such, the average strength of model effects in real world decisions might be lower than in our first study.

On the other hand, one might argue that the team-situation in the first study enhanced the accuracy motivation as compared to the second study without a team condition. We do not believe that this was the case. First, the incentivization for individuals in both studies was exactly the same. Moreover, a team situation does not enhance the probability of success per se, because the behavior of the team member cannot be controlled and, hence, adds an uncertainty factor. Second, and most importantly, an alternative interpretation in terms of differences in motivation cannot account for the fact that we found no significant differences between the studies in the learning phases. In both studies, participants required fewer than 60 trials, on average, to complete the learning phase⁵.

It is important to note that one should not readily generalize the finding that participants did not show routine biases in their choices. Recall the distinction we made between a routine effect and a routine bias. Roughly speaking, a routine effect occurs when adaptive preference formation (consistently choosing the best alternative) in the test phase takes longer than in the learning phase. Routine effects are normatively justified because they reflect the consideration of all prior information during a belief updating process (e.g. by applying Bayes' rule). A routine bias occurs when the routine is still chosen although posterior probabilities would justify routine deviation. We found routine effects but no routine biases. This contradicts our observation that routine biases are pervasive in everyday behavior. For instance, consider the case of maladaptive habits such as drug consumption or stable self-destructive behavioral patterns in close relationships.

Again, we must be aware of the special features of our research paradigm. First, we induced only weak routines and *not* habits. Habits are routines that are so frequently performed that they are automatically instantiated (James, 1890/1950; Wood & Neal, 2007). In our studies the routine was repeated an average of 30 to 40 times during the learning phase. This repetition frequency is very low compared to many habits that have been consolidated over a person's lifetime. Therefore, we wish to limit our conclusions to non-automatic routines (cf. Betsch & Haberstroh, 2005a, for a discussion). Note also that our choice environment provides feedback in an ideal fashion. All feedback was precise, valid and always consistently presented directly after the choice was made. As such, we exposed participants to a "friendly" feedback environment, to use Robin Hogarth's term (Hogarth, 2001). In real world settings, individuals also face "wicked" environments in which feedback is delayed, inconsistent, opaque or invalid. For example, many of our consumer routines yield positive outcomes in the short run while the negative outcomes are postponed. The hamburger we buy in the convenience restaurant is cheap and tasty. The negative outcomes, such as health problems, are delayed. Some cannot be directly observed, such as the increase in carbon dioxide concentration by the intensive stock-rearing of the animals that yield the meat for the burger. Not surprisingly, routine deviation is less likely in such environments than in the tasks in our experiments.

5 Trials needed for completion of learning phase in Exp. 1: M Human Model Observable = 58.5 (Std = 12.5), M HM Not Observable = 60.3 (Std = 14.1), $t(46) = 0.47$, $p = .64$, $d = .139$; in Exp.2: MNon-Human Model Observable = 58.1 (Std = 22.4), MNHM Not Observable = 55.8 (Std = 13.5), $t(48) = 0.44$, $p = .66$, $d = .12$.

Another boundary condition for our results stems from the pay-off structure of the task. In our studies, the pay-offs are risky; and their values vary within the same option. In decisions under certainty and without outcome variation, routine deviation may occur instantly (Humphreys, 1939). For instance, if the routine consistently produced the same outcome in the past (say “5”), a change in outcomes (say “-5”) can be easily noticed and will likely lead to immediate routine deviation. Note, however, that many decisions in the real world must be made in probabilistic environments such as in our study.

Social situations have emergent properties. Social frames can promote adaptation. For instance, logical thinking in the Wason selection task (Wason, 1960) improves when the task is embedded in a relevant social context (e.g. cheater detection, Cosmides, 1989). Our research adds to this evidence by showing that social context matters. Specifically, it provides insights into the mechanism by which our neighbor’s behavior can impact our own decisions.

References

- Aarts, H., & Dijksterhuis, A. (2000). Habits as knowledge structures: Automaticity in goal-directed behavior. *Journal of Personality and Social Psychology*, *78*, 53-63.
- Athay, M., & Darley, J.M. (1981). Toward an interaction-centered theory of personality. In N. Cantor & J.F. Kihlstrom (Eds.), *Personality, cognition and social interaction* (pp. 281-308). Hillsdale, N.J.: Erlbaum.
- Bandura, A. (1977). *Social learning theory*. Englewood Cliffs, N.J.: Prentice-Hall.
- Baron, J., & Ritov, I. (1994). Reference points and omission bias. *Organizational Behavior and Human Decision Processes*, *59*, 475-498.
- Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, *50*, 7-15.
- Betsch, T. (2005). Preference theory – an affect-based approach to recurrent decision making. In T. Betsch & S. Haberstroh (Eds.). *The routines of decision making* (pp. 39-65). Mahwah, N.J.: Lawrence Erlbaum.
- Betsch, T., & Haberstroh, S. (2005a). The routines of decision making: Preface. In T. Betsch & S. Haberstroh (Eds.). *The routines of decision making* (pp. ix-xxv). Mahwah, N.J.: Lawrence Erlbaum.
- Betsch, T., & Haberstroh, S. (2005b). Research on the routines of decision making: Advances and prospects. In T. Betsch & S. Haberstroh (Eds.). *The routines of decision making* (pp. 359-376). Mahwah, N.J.: Lawrence Erlbaum.
- Betsch, T., Haberstroh, S., Molter, B., & Glöckner, A. (2004). Oops, I did it again -- relapse errors in routinized decision making. *Organizational Behavior and Human Decision Processes*, *93*, 62-74.
- Betsch, T., Haberstroh, S., Glöckner, A., Haar, T., & Fiedler, K. (2001). The effects of routine strength on information acquisition and adaptation in recurrent decision making. *Organizational Behavior and Human Decision Processes*, *84*, 23-53.
- Betsch, T., & Held, C. (2011). *Rational decision making: Balancing run and jump modes of analysis*. Manuscript submitted for publication.
- Betsch, T., Fiedler, K., & Brinkmann, J. (1998). Behavioral routines in decision making: The effects of novelty in task presentation and time pressure on routine maintenance and deviation. *European Journal of Social Psychology*, *28*, 861-878.
- Boyd, R., & Richerson, P.J. (2005). *The origin and evolution of cultures*. New York: Oxford University Press.

- Bröder, A. & Schiffer, S. (2006). Adaptive flexibility and maladaptive routines in selecting fast and frugal decision strategies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 904-918.
- Cameron, A. C., & Trivedi, P.K. (2005) *Microeconometrics: Methods and Applications*. New York: Cambridge University Press.
- Cosmides, L. (1989). The logic of social exchange: has natural selection shaped how humans reason? *Cognition*, 31, 187-276.
- Hogarth, R.M. (2001). *Educating intuition*. Chicago: University of Chicago Press.
- Humphreys, L. G. (1939). Acquisition and extinction of verbal expectations in a situation analogous to conditioning. *Journal of Experimental Psychology*, 25, 294-301.
- James, W. (1890/1950). *The principles of psychology*, Vol. 1. New York: Dover.
- Jones, E. E.; & Davis, K. E. (1965). From acts to dispositions: the attribution process in person perception. In L. Berkowitz (ed.), *Advances in Experimental Social Psychology* (Vol. II, pp. 219-266). New York: Academic Press.
- Laland, K.N. (2004). Social learning strategies. *Learning and Behavior*, 32, 4-14.
- Luchins, A.S., & Luchins, E.H. (1959). *Rigidity of behavior. A variational approach to the effect of Einstellung*. Eugene: University of Oregon Press.
- Mussweiler, T. & Rüter, K. (2003). What friends are for! The use of routine standards in social comparison. *Journal of Personality and Social Psychology*, 85, 467-481.
- Ouellette, J. A., & Wood, W. (1998). Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Bulletin*, 124, 54-74.
- Reimer, T., Bornstein, A.-L., & Opwis, K. (2005). Positive and negative transfer effects in groups. In T. Betsch & S. Haberstroh (Eds.), *The routines of decision making* (pp. 175-192). Mahwah, NJ: Lawrence Erlbaum Associates.
- Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1, 7-59.
- Triandis, H.C. (1977). *Interpersonal behavior*. Monterey, CA: Brooks/Cole Publishing Company.
- Verplanken, B. & Aarts, H. (1999). Habit, attitude, and planned behavior: Is habit an empty construct or an interesting case of goal-directed automaticity? In: W. Stroebe & M. Hewstone (Eds.), *European Review of Social Psychology* (Vol. 10, pp. 101-134). Chichester: Wiley.
- Wason, P.C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 12, 129-140.

- Webster, M.M., & Laland, K.N. (2008). Social learning strategies and predation risk: minnows copy only when using private information would be costly. *Proceedings of the Royal Society, Series B* 275, 2869-2876.
- Wood, W., & Neal, D. T. (2007). A new look at habits and the habit-goal interface. *Psychological Review*, 14, 843-863.
- Yechiam, E., Druyvan, M., & Ert, E. (2008). The effect of observing others on risk taking in decisions from experience. *Judgment and Decision Making*, 3, 493-500.