The Role of Feedback and Post-Error Adaptations in Reasoning

Hannah Dames (damesh@cs.uni-freiburg.de)
Clemens Schiebel (schiebec@cs.uni-freiburg.de)
Marco Ragni (ragni@cs.uni-freiburg.de)
Cognitive Computation Lab, Department of Computer Science
Georges-Kohler-Allee 52, 79110 Freiburg, Germany

Abstract

Monitoring our errors enables humans to adapt behavior when actions fail to result in desired outcomes. Post-error adaptations have been studied extensively using simple laboratory tasks where people typically slow down after errors. Few studies, however, examined such behavioral adaptations in more complex tasks such as reasoning. In two experiments we investigated how participants adapt their behavior based on evaluative feedback in syllogistic reasoning tasks. Experiment 1 demonstrates that participants' likelihood to give a logically correct response increased throughout the experiment when given feedback. This feedback effect was limited to syllogisms that have no logical conclusion and thus mostly driven by an increase in participants' "No valid conclusion" responses. Experiment 2 investigates post-error adaptations on a trial-level and shows that participants with a high accuracy slowed down after errors while participants with a low accuracy slowed down after correct responses. Implications on error-monitoring and reasoning research are discussed.

Keywords: post-error adaptations; reasoning; feedback

Introduction

Humans try to make sense of the world by inferring conclusions from the information they perceive and act accordingly. Despite our best effort, our reasoning and actions can occasionally fail to produce the desired outcome. For instance, in politics, education, or even casual conversations different parties quickly dispute any illogical argument and inform the other about errors. Even for simple actions performed in everyday life, like pouring a cup of tea, feedback provides us with the important information about whether our actions produce an intended effect or not. By monitoring our errors we are able to adapt to frequently changing environments and we can adjust our actions when they suddenly fail to produce an intended effect (Ullsperger & Danielmeier, 2016). Unsurprisingly, post-error processing is thus often assumed to be an adaptive process aimed at improving our behavior (Wessel, 2018).

Using simple, speeded response tasks in laboratories (e.g., reacting to arbitrary stimuli by left and right key presses) a great amount of research has investigated how we adapt our behavior after committing an error (e.g., Danielmeier & Ullsperger, 2011; Rabbitt, 1966; Wessel, 2018). Here, post-error adaptations generally refer to the neurophysiological and behavioral changes related with error processing (for an overview see Wessel, 2018). Recently, studies started to emphasize that it is crucial to investigate error processing and corresponding adaptations in more complex tasks which resemble daily situations of flexible behavior (Desmet et al., 2012; van der Borght, Desmet, & Notebaert, 2016). Yet, there is a lack of research in the field of error-processing concerning tasks that require higher-level cognitive processes and multiple strategies such as reasoning. In turn, the role of feedback and how people adapt following errors has not been in the focus of reasoning research. For instance, syllogistic reasoning is a greatly researched domain and to this day over twelve theories aim to explain how humans reason about syllogisms using the syllogistic reasoning task (see Khemlani & Johnson-Laird, 2012). Still, the power of these theories to consider effects of feedback or flexible adaptations when perceiving errors are not reported or not existent (but see Ackerman & Thompson, 2017, for a meta-cognitive perspective on monitoring when reasoning). Additionally, there is only one study that has investigated the role of evaluative feedback on reasoning using fourteen out of the 64 syllogisms (Khemlani & Moore, 2012). The present study aims to fill this gap investigating (1) how feedback changes participants' performance over the time-course of a reasoning experiment and (2) how people adapt their reasoning behavior in trials following an error.

The Syllogistic Reasoning Task The psychology of reasoning investigates the cognitive processes driving human inferential mechanisms. In the syllogistic reasoning task, as a core domain of human reasoning research, participants are typically instructed to derive a conclusion from two quantified statements called the premises. The two premises each have one of four quantifiers, called moods, in their syllogistic combination: *All* (abbreviated by A), *Some* (I), *None* (E), and *Some...not* (O). The terms in the premises (the sets of entities) can be arranged in four different ways called figures (see Khemlani & Johnson-Laird, 2012). In combination, there are a total 64 different kinds of reasoning problems, one illustrated in the following example:

All architects are beekeepers.
<u>*Some beekeepers are chemists.*</u>
What, if anything, follows?

The task is to generate a response about the two unrelated terms, architects and chemists, using one of the quantifiers or to conclude that "no logically valid conclusion" – NVC for short – follows. Most individuals erroneously infer from the information in the example above that "some architects are chemists" (Khemlani & Johnson-Laird, 2012). Yet, the only logical valid response is that nothing follows. Syllogisms where a conclusion with a quantifier logically follows are called *valid* and those were NVC is the logically correct response are called *invalid* syllogisms (which is the case in the given example). Thirty-seven out of the 64 syllogistic problems, hence a total of 58%, are invalid, thus, only NVC can be derived. As we will later see, some participants appear to be heavily biased against this response option (e.g., Ragni, Dames, et al., 2019; Revlis, 1975). Could such biases or aversions against NVC be overcome by providing feedback?

Measuring Post-Error Adaptations in Reasoning Tasks: Three Challenges The investigation of adaptations following feedback in the syllogistic reasoning task is novel and challenging: First, the analysis of post-error adaptations within participants requires a comparison of response characteristics in trials following either correct or incorrect responses. Thus, many trials per person are needed to gain sufficient observations for statistical analyses. In simple, speeded response-mapping tasks, this poses no challenge as participants usually respond to arbitrary stimuli (e.g., numbers, letters) with a response deadline of less than 1s (Wessel, 2018). Naturally, within an experimental session of 30 minutes, a participant can then respond to hundreds of such stimuli. However, reasoning tasks, and syllogisms in particular, require the participant to think about the presented information and to apply cognitive processes such as heuristics, strategies, or logical reasoning. As a result, participants typically take their time to infer a conclusion (around or above 20s) with no experimental response deadline (Hardman & Payne, 1995). Hence, fewer observations can be collected within on experimental session as trials naturally last longer. In the current study, we aim to collect data for all 64 syllogisms to meet this challenge.

Second, individuals differ in their ability to reason logically (e.g., Galotti, Baron, & Sabini, 1986; Khemlani & Johnson-Laird, 2016; Stanovich & West, 2000). In the given syllogistic example, for instance, most but not all individuals erroneously respond, "S*ome architects are chemists*" (Khemlani & Johnson-Laird, 2012). Consequently, there are some participants that commit errors frequently while others do not. Not only can highly imbalanced data be a consequence (we will meet this problem using mixed models) but post-error adaptations have also been observed to differ depending on participants' accuracy (Notebaert et al., 2009). In Experiment 2, we will specifically account for participants' accuracy within our hypothesis and model.

Third and most importantly, participants typically do not receive feedback in the syllogistic reasoning task and we must thus first analyze how feedback impacts this task in general. In particular, in order to investigate changes on a reaction-time (RT) level following feedback, we aim to present participants with feedback only for a short period of time and provide no additional explanation on why their response was correct or incorrect (similar to studies on post-error processing). A first study provided initial evidence that, at least for a subset of syllogisms (most difficult and valid problems), participants yield more correct responses when receiving only such evaluative feedback (44%) than when no feedback is provided (e.g., 33%, Khemlani & Moore, 2012). However, the authors used only fourteen of the most difficult valid syllogisms. It is therefore still unclear how feedback may affect participants responses to all 64 syllogisms, regardless of difficulty and validity. In addition, how people adapt their reasoning behavior when given feedback for invalid syllogisms is of most importance for the syllogistic reasoning task: A recent study demonstrated that people greatly differ in their tendency to respond NVC – to the extent that some participants avoided this response option in general (Ragni, Dames, et al., 2019). This indicates that for some participants invalid syllogisms can pose an exceptionally high challenge possibly because they avoid or misinterpret the NVC response option. We predict that participants are able to overcome the mishandling of the NVC response when provided with short and instant feedback (see Experiment 1).

In sum, these considerations highlight the need to first investigate the general role of feedback on solving reasoning tasks before exploring how people adapt following such feedback. Hence, we conducted a first experiment to investigate how feedback affects participants' performance throughout the time-course of an experiment in general (research question 1, RQ1) and then analyzed potential post-error adaptations in a separate experiment (RQ2).

Experiment 1 – The Effect of Feedback

Experiment 1 examined the influence of feedback on participants' performance in a reasoning task. To this end, we had participants solve all 64 syllogisms and provided them either with evaluative feedback (1s; feedback condition) or not (non-feedback condition). For the reasons stated in the introduction, we expected the likelihood to give a logically correct response to be generally higher in the feedback than in the non-feedback condition (**Hypothesis 1, H1**).

Furthermore, we expected reasoning performance to improve with increasing trial number. Note, that this prediction is in contrast to Khemlani and Moore (2012) who found in their second experiment that accuracy on the first five trials was not reliably lower than on the last five trials (41% vs. 45%, non-significant). Thus, in their experiment performance in the feedback condition did not increase steadily over the time-course of solving fourteen syllogistic reasoning problems with feedback (no information is provided in the non-feedback condition). Unfortunately, there is a lack of research investigating how participants' reasoning performance evolves over time when testing all 64 syllogisms (but see Ragni, Dames, et al., 2019). That the stability of the human reasoning behavior in the syllogistic reasoning task has been neglected so far is surprising as it can be assumed that individuals use different strategies to solve syllogisms. Some theories suggest that people also use simple heuristics to solve syllogisms (e.g., probabilistic heuristics Oaksford & Chater, 2007). In an experimental setup considering all 64 syllogisms, we thus hypothesized participants to improve in their accuracy with increasing trial number for the following reasons: We assumed that the development and application of such strategies and heuristics should generally increase the likelihood to give a logically correct response with increasing trial number regardless of the feedback condition (**H2a**). This prediction is also in line with a recent study that investigated the question of whether participants improve in their reasoning performance over the time-course of an experiment (Dames, Klauer, & Ragni, in preparation). Here, we demonstrated that that both within an experimental session as well as between two test sessions, the likelihood to give a logically correct response to a syllogism increased over time – however, mostly for valid and not invalid syllogisms (a summary of the main results can be found online at osf.io/x3wvf; Dames, Klauer, & Ragni, 2020).

According to Oaksford and Chater (2007) participants also respond logically incorrect to syllogisms when they apply cognitive inexpensive heuristics that are fallible. While feedback may trigger improved performance by inducing reasoners to apply all heuristics instead of just a subset (see Khemlani & Moore, 2012), reasoners that were told that the deduction they drew from a set of premises was incorrect, may be less inclined to use the same heuristic in the future and vice versa. Furthermore, feedback directly incentivises applied heuristics and strategies. Hence, we predicted that the positive influence of more problems solved (as realized by the trial-number) will be greater in the feedback than in the non-feedback group (**H2b**).

In previous work we showed that without feedback people improve mostly for valid but not for invalid syllogisms with increasing trial number (Dames, Klauer, & Ragni, in preparation). This is surprising as 58% of the syllogisms are invalid with NVC as the logically correct response. This observation can be explained when considering that some participants may be less inclined to respond NVC assigning this response a meaning of "giving up" (see Ragni, Dames, et al., 2019). Such biases against the NVC response have been proposed in earlier works already (e.g., Revlis, 1975; Roberts et al., 2001) Feedback may thus not only help participants to learn that for some types of problems a valid conclusion cannot be found (i.e., they may start to assume that the proportion of invalid problems is high), they may also become more confident in responding NVC. This could be the case, for instance, when participants in the feedback condition become aware of the great proportion of NVC responses in the task. As this process is not necessarily a result of participants becoming more logical, we should find an improvement mostly for invalid syllogisms. In addition, participants receive no further information on *why* their answers were incorrect in the current experiment. Such feedback may not help to apply the knowledge of one's response's correctness to different valid syllogisms. Reasoners however may find similarities in the structure of invalid syllogisms over time (e.g., whenever there are two *some's* in the premises, nothing can follow; Galotti, Baron, & Sabini, 1986). Consequently, we assume that the effect of trial number on the likelihood to give a logically correct response in the feedback group is stronger for invalid than for valid syllogisms (**H3**).

Method

Participants Sixty-nine participants (M_{age} = 41.1 years, SD = 10.7, 58% female, 42% male) were recruited via Amazon Mechanical Turk for the online experiment and received monetary compensation for their participation. Only participants that finished the whole experiment and agreed to the usage of their data were taken into consideration. Participants with highly unrealistic completion time of below 20 minutes (n = 11; as an indicator for careless responses) were excluded and replaced with new ones. Participants were randomly assigned to either the feedback (n = 30) or the non-feedback (n = 39) group (between-subject).

Task and Materials Participants' task was to generate a conclusion for all 64 possible syllogisms consisting of two premises each using a selection task design. That is participants were instructed to select a response out of all nine possible conclusions presented on the screen. Participants were told to assume the premises to be true and were instructed to draw a conclusion only if it necessarily followed from the two premises (see introduction). Furthermore, participants were asked to respond as quickly and accurate as possible. The content of the syllogisms depicted professions of groups of people where we assumed no influence of content on the believability of premises. More information on the task and instructions as well as all materials and data are made publicly available.[1]

Trial Structure and Procedure Each trial started with the presentation of a central fixation cross for 300 ms (the inter-trial-interval, ITI) followed by the two premises. The premises were presented until the participants used the spacebar thereby indicating they have found a conclusion to the premises. Subsequently, the nine response options were displayed. Participants selected one of them using the mouse and then pressed a "continue"-button to confirm their selection. Starting from the onset of the premises, participants were provided with a response deadline of 90 seconds. When they exceeded the response deadline at any time within a trial, the trial was aborted. Only in the feedback condition, a feedback-screen informed the participants about the accuracy of their response (1s, "correct" for correct responses; "incorrect" for incorrect ones; "too slow!" for response outside the response frame of 90 s). Overall, the experiment consisted of a practice task and four blocks, consisting of 16 trials each. In between blocks, participants could take a self-paced break. The order of the 64 syllogisms, content of the premises, and order of the presented conclusions was random per participant. At the end of the experiments, participants answered additional questions about themselves and were debriefed.

Results and Discussion

We excluded single trials with response omissions (i.e., trials where participants did not respond within the given timeframe; n = 16 trials out of all trial from all participants, 0.3% of all the data) from the analysis. Mean error rates and RTs can be taken from Table 1.

[1] osf.io/dy3gr (Dames, Schiebel, & Ragni, 2020)

Table 1: Response times (RTs) and error rates

	No Feedback	Feedback
RTs (s): Mean (SD)	22.22 (7.82)	19.70 (7.77)
Errors (%): Mean (SD)	67.4 (18.3)	56.6 (21.0)

To test our hypotheses and to account for the multi-level structure of the design (e.g., multiple measures per participant), we employed a generalized linear mixed model (GLMM; see Baayen, Davidson, & Bates, 2008) analysis and reported the odds ratio statistic (OR). Estimates were computed with maximum likelihood and binomial link functions. Effect coding was used for all dichotomous fixed effects. All continuous predictors were centered and scaled. The correctness of participants' responses on a given trial *n* (0 = incorrect vs. 1 = correct) was analyzed as a function of the fixed factors *group* (1 = feedback vs. -1 = no feedback), *trial number* (1- 64), *validity* of a syllogism (1 = invalid vs. -1 = valid), and the corresponding interactions. We implemented the maximal random-effects structure justified by the design (as suggested by Barr, Levy, Scheepers, & Tily, 2013): by-participant random intercept and random slopes for all fixed factors, their interaction including all possible correlations between the random effects, and a by-syllogisms random intercept. As this model did not adequately converge, the random slope for the interaction term was removed. The results of the GLMM can be taken from Table 2. In line with H1, the main effect of *group* (see the results for the predictor "Group" in Table 2) was significant, demonstrating that the likelihood to commit a correct response was on average higher in the feedback than in the non-feedback condition. As predicted in H2a, participants improved throughout the course of the experiment (effect "Trial_n" in Table 2). This effect was, in line with H2b, stronger for the feedback group as apparent in the significant interaction between the two factors ("Group x Trial_n", Table 1). Our results also show that this interaction was strongest for invalid syllogisms ("Group x Val. x Trial_n", Table 1): For valid syllogisms only participants in the non-feedback but not in the feedback group improved over time. In line with this finding, the effect of feedback was generally stronger for invalid than for valid syllogisms (see the significant interaction between "Group x Validity" in Table 1). All results and the hypothesized three-way interaction are shown in Figure 1.

Table 2: GLMM for the correctness of participants' responses on a given trial (correct vs. false)

Predictors	*Odds Ratios*	*Std. Error*	*z-value*	*p-value*
Intercept	0.43	0.24	-3.43	**.001**
Trial_n	1.15	0.05	2.91	**.004**
Group (feedback)	1.42	0.18	2.00	**.045**
Validity (invalid)	0.49	0.20	-3.46	**.001**
Group x Trial_n	1.10	0.05	2.06	**.039**
Validity x Trial_n	1.12	0.04	2.67	**.008**
Group x Validity	1.62	0.11	4.30	**<.001**
Group x Val. x Trial_n	1.17	0.04	3.71	**<.001**

The present findings indicate that feedback increases an individual's likelihood to respond logically correct to syllogisms. Expanding results provided by Khemlani and Moore (2012), we found that when using all 64 problems, this enhancement in accuracy is mostly driven by participants responding "NVC" more often as the effect of feedback was much stronger for invalid than for valid syllogisms (one could argue that it was not apparent for valid syllogisms). In addition, we found a clear evolution over the course of the experiment for this feedback effect. Yet again, this was only true for invalid syllogisms. In summary, these results suggest that reasoners improved not by becoming more logical over time in general but by learning the underlying distribution of NVC responses. We were thus able to demonstrate that feedback impacts participants' reasoning behavior, however, leaving the question open how participants may adapt their behavior through feedback on a trial-by-trial level.

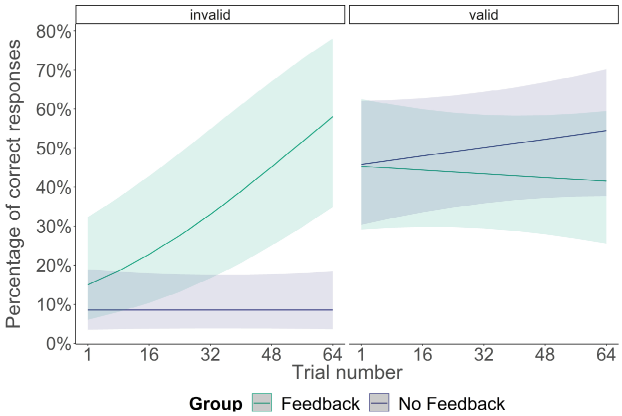

Figure 1: Marginal means separated for valid and invalid syllogisms. Error bands depict 95 % confidence intervals.

Experiment 2 – Post-error adaptations

In Experiment 2, we investigate whether humans show post-error adaptations commonly observed in simple stimuli response-mapping tasks when solving syllogisms with feedback. To the best of our knowledge, this is the first study investigating post-error adaptations in syllogistic reasoning.

Behavioral Changes Upon Error Commission How human reactions change on a behavioral basis upon the occurrence of an error received a considerable amount of attention (for an overview, see Danielmeier & Ullsperger, 2011). To prevent error repetitions, motor adaptations and attentional focusing toward task-relevant information are induced (e.g., on a neural level; Marco-Pallarés, et al., 2008). Those adaptations in cognitive control typically manifest through post-error changes in reaction times (RTs) and error rates (see Danielmeier & Ullsperger, 2011). Research suggests that errors first trigger a stereotypic, two-step cascade of automatic processing consisting of initial rapid inhibitory control processes (1) which then facilitate a subsequent shift in attentional orientation away from the current task-set representation (2), towards the source of the error (for an overview see Wessel, 2018). As a result, people tend to slow down in trials following the error (Rabbitt, 1966). Traditionally, this well-replicated post-error slowing (PES) effect has been linked to cognitive control and inhibitory processes (Botvinick et al., 2001; Ridderinkhof et al., 2004). This observation is in line with theories connecting PES to a by-product of a re-orienting process initiated by the error (Notebaert et al., 2009) or theories of capacity limitations during error monitoring (Jentzsch & Dudschig, 2009), but it still remains a discussion point. Only later in time, once the two-step cascade of control processes is completed, error-specific adaptations are assumed to occur aimed at improving the ongoing task (Wessel, 2018). In particular, with longer inter-trial-interval (ITI), attentional focusing toward task-relevant information has been observed in a *post-error improvement in accuracy* (PIA; e.g., Maier, Yeung, & Steinhauser, 2011).

Post-Error Adaptations in Complex Tasks How may reasoners react to errors in reasoning tasks? It is still unclear whether typical post-error adaptations, namely the PES and the PIA effect extend to the syllogistic reasoning task. Based on the above mentioned, well-replicated studies for the PES effect and given that we can also observe PES in more complex mental arithmetic tasks (Desmet et al., 2012; van der Borght, Desmet, & Notebaert, 2016), we assume to also find a typical increase in RTs in post-error trials as compared to post-correct trials. Note, however, the following differences between the present and previous tasks that have investigated PES: Although ITIs in the present and previous studies on error-processing are comparable (300ms), participants in the present study receive a rather long response deadline (1.5 minutes). However, we still assume that the error processing and monitoring processes elicited by an error occupy a central bottleneck (e.g., Jentzsch and Dudschig, 2009). Participants are assumed to take some of the time required to complete this automatic processing cascade on the next trial leading to an increase in RTs for post-error trials on average.

Last, the inter-individual differences in participants' reasoning ability (i.e., accuracy) need to be considered. There are some theories that provide substantial evidence that it is the infrequent event that causes PES and not the erroneous nature of the incorrect response: Notebaert et al. (2009) demonstrated that PES occurs after infrequent errors but they could also observe post-correct slowing (PCS) following infrequent correct responses. Consequently, PES can be assumed to be part of a general orienting response prior to an actual error-specific adaptation effect (Wessel, 2018). Hence, PES should occur only for participants with a high accuracy and instead a PCS effect for low-accuracy participants. In those cases, correct feedback should be the unexpected, motivationally salient event that captures participants' attention and distracts them during the processing of the subsequent syllogism. We thus predict to find a difference in RTs between post-error and post-correct trials moderated by an individual's accuracy (**H1**). Note, that the investigation of a potential PIA effect is not part of this study.

Method

Recruitment and data exclusion criteria were the same as in Experiment 1 resulting in a sample size of $n = 72$ ($M_{\text{age}} = 45.8$, $SD_{\text{age}} = 10.9$, 47.2% female, 52.8% male). Note, that prior to data collection we intended to reach our desired sample size of $n = 100$ by including the $n = 30$ participants from the feedback condition in Experiment 1 in the current analysis resulting in a final sample size of $n = 102$. The task, material, and procedure were identical to the feedback condition of Experiment 1 (all participants were provided with feedback) and are available online[1].

Results and Discussion

Results and model specifications can be found online[1]. Mean RT was 21.39s ($SD = 8.59$) and mean error rate was 55.0% ($SD = 20.7\%$). Single trials with and following response omissions ($n = 10$), error trials as well as trials both directly preceding and following an error (i.e., in-between error-trials. $n = 1980$) were excluded from the post-error analyses. Also, trials with RTs deviating more than 3 *SD*s from the mean RT within each condition and participant were removed from analyses ($n = 42$). Last, the first trial of each block was not considered as these trials are not preceded by a correct or incorrect response. After deleting those trials, mean RT was 19.42s ($SD = 8.77$). Data-analysis procedure and selection of random-effect structure were the same as in Experiment 1 except for using a linear mixed model for the analysis of RTs. We used restricted maximum likelihood estimation for slope estimates. *p*-values for effects were obtained using the Kenward-Roger (1997) approximation for denominator degrees of freedom. The RTs were analyzed with the predictors *response-correctness*$_{n-1}$ of the preceding trial (depicting participants' response accuracy in trial n-1; 1 = error, -1 = correct) and participants' *accuracy* (the relative frequency of correct responses over all trials per participant) including their interaction. The trial sequence number (1-64) was added as covariate capturing effects due to fatigue or learning. We added the *validity* of a syllogism and the corresponding interaction with *response-correctness*$_{n-1}$ to the model to control for a potential influence of validity.

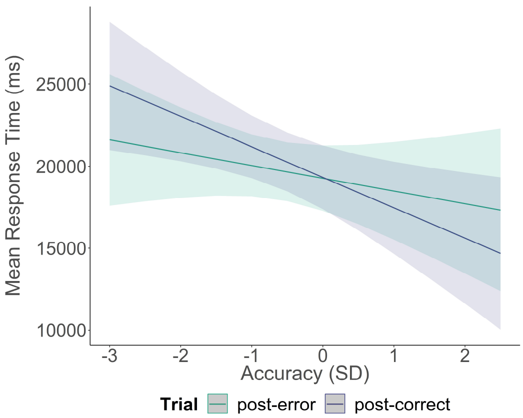

Figure 2: Marginal means for RTs by participants' standardized (SD) accuracy. Error bands: 95% confidence intervals.

The model that converged consisted of by-participant and by-syllogism random intercepts and random slopes for response-correctness$_{n-1}$, including their correlation. Participants response latency was 19.s ($p < .001$, intercept). Participants generally respond faster the higher the trial-number ($\beta = -1.97$, $p < .001$). The main effect of *response-correctness$_{n-1}$* did not reach significance ($\beta = -0.03$, $p = .921$), but as expected and in line with our hypothesis (H1), there was a significant interaction of *response-correctness$_{n-1}$* and participants' error-rates ($\beta = 0.54$, $p = .029$). As also illustrated in Figure 2, we can thus confirm that we found a PES only for participants with a low error rate and a PCS for participants that committed errors frequently. In summary, Experiment 2 demonstrates that also in the syllogistic reasoning task, participants adapt to correct and incorrect feedback on a trial-to-trial basis.

General Discussion

The aim of the present paper is to explore the behavioral changes induced by evaluative feedback in a syllogistic reasoning task. We investigated this research question twofold: In a first experiment, we investigated to what extent feedback influences participants' reasoning performance over time. In a second experiment we analyzed how people adapt to incorrect feedback on a trial-by-trial level (i.e., changes in RTs).

These are core questions for *theories of syllogistic reasoning*, as they do not yet consider such flexible adaptations based on feedback, and for *research on post-error adaptations*, as our study extends their methods to tasks from complex cognition. To this end, in Experiment 1, we had two groups of participants perform a syllogistic reasoning task receiving either no feedback or feedback. Our results show that with feedback participants' reasoning performance is enhanced. Yet, in contrast to Khemlani and Moore (2012) who used only valid syllogisms, the observed improvement in the feedback condition was driven mainly by participants responding "NVC" more often as participants with feedback only improved for invalid but not for valid syllogisms. Furthermore, our study demonstrates that the improvement on invalid syllogisms clearly evolves over the course of the experiment. We attribute this effect to participants learning the underlying distribution of NVC responses over time with feedback (i.e., they start to assume that the probability of NVC problems is high; it is 58% for all syllogisms). Consequently, this process is not necessarily associated with participants becoming more logical in general. Based on rapid supervised learning principles, reasoners may have learned that for certain structures of the premises (e.g., whenever there are two *some's* in the premises, Galotti et al., 1986) nothing can follow. These observations may also explain why participants did not substantially improve for valid syllogisms in the feedback condition. Moreover, note that participants received no information on *why* their conclusion was incorrect and were given only 300ms (ITI) to process feedback. Possibly, feedback may not help to apply the knowledge of one's response correctness to different sets of valid premises. This raises the question whether people may learn effective response strategies for valid syllogisms when they are given more time to process feedback – an aspect that should be addressed in future studies.

In Experiment 2, we investigated how feedback impacts participants' RTs on a trial-by-trial level. We predicted that we could extend findings from standard speeded response tasks to more complex reasoning problems. In line with this prediction, the results provide first evidence that for the syllogistic reasoning task, indeed, typical post-error adaptations can be observed: Upon errors, participants slowed down when their accuracy was high. When participants' accuracy was low, they slowed down after correct responses. Thereby, the present study revealed that individual differences associated with performance accuracy can substantially modulate post-error reactivity even in complex tasks. These results are in line with the assumption that an orienting response as a source of post-error adjustments is negatively related to the frequency of errors (e.g., Notebaert et al., 2009). That is, the more infrequent errors are, the better an error can be detected (e.g., Coles, Scheffers, & Holroyd, 2001) and the more likely PES can be observed. Most importantly, we could demonstrate that not only typical post-error adaptations but also their modulation by inter-individual differences can thus generalize to reasoning tasks and to a situation in which the response deadline is over 1 minute. Interestingly, a recent study (Aczel & Palfi, 2017) also suggests that cognitive control adaptations effects observed for standard speeded response tasks (e.g., congruency sequence effect) can be found in the ratio-bias reasoning task. Together with the present findings, we therefore advocate the consideration of cognitive control and adaptation processes on a trial-by-trial level in reasoning research: We believe that they provide important insights not only for cognitive but also for metacognitive models on reasoning (e.g., Ackerman & Thompson, 2017) by shedding light on an important but so far neglected mechanism.

Conclusion In conclusion, we were able to demonstrate that feedback impacts participants' reasoning behavior – an observation that needs to be addressed by theories on syllogistic reasoning. Importantly, feedback improves reasoning performance over time, but this improvement seems to occur almost exclusively for invalid syllogisms. We conclude that this effect is driven by participants learning the relative importance of NVC responses over time with feedback. On a trial-by-trial level, the present study additionally provides evidence that typical post-error adaptations observed in speeded response (low-level cognition) tasks can be found when solving syllogisms. The reported post-error and post-correct slowing effects suggest that error processing might be an important mechanism underlying performance on reasoning tasks. Investigating such trial-by-trial adaptations could potentially facilitate the refinement of predictions made by cognitive theories and facilitate the exploration of cognitive control in human reasoning.

Acknowledgements

This research was made possible by grants from the DFG in proposal RA 1934/3-1, RA 1934/4-1, and RA 1934/9-1.

References

Ackerman, R., & Thompson, V. A. (2017). Meta-reasoning: Monitoring and control of thinking and reasoning. *Trends in Cognitive Sciences*, *21*(8), 607-617.

Aczel, B., & Palfi, B. (2017). Studying the role of cognitive control in reasoning: Evidence for the congruency sequence effect in the ratio-bias task. *Thinking & Reasoning*, *23*(1), 81–97.

Baayen, R.H., Davidson, D.J., & Bates, D.M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*(3), 624–652.

Coles, M.G., Scheffers, M. K., & Holroyd, C. B. (2001). Why is there an ERN/Ne on correct trials? Response representations, stimulus-related components, and the theory of error-processing. *Biological Psychology*, *56*(3), 173-189.

Dames, H., Klauer, K. C., & Ragni, M. (in preparation). The Stability of Syllogistic Reasoning Performance Over Time

Dames, H., Klauer, K. C., & Ragni, M. (2020). Individual Differences in Syllogistic Reasoning Performance: Why and When do Individuals Differ? Retrieved from osf.io/x3wvf

Dames, H., Schiebel, C., & Ragni, M. (2020). The Role of Feedback and Post-Error Adaptations in Reasoning. Retrieved from osf.io/dy3gr

Danielmeier, C., & Ullsperger, M. (2011). Post-Error Adjustments. *Frontiers in Psychology*, *2*(233), 1-10.

Desmet, C., Imbo, I., Brauwer, J. de, Brass, M., Fias, W., & Notebaert, W. (2012). Error adaptation in mental arithmetic. *The Quarterly Journal of Experimental Psychology*, *65*(6), 1059–1067.

Galotti, K. M., Baron, J., & Sabini, J. P. (1986). Individual differences in syllogistic reasoning: Deduction rules or mental models? *Journal of Experimental Psychology: General*, *115*(1), 16–25.

Hardman, D. K., & Payne, S. J. (1995). Problem difficulty and response format in syllogistic reasoning. *The Quarterly Journal of Experimental Psychology*, *48*(4), 945-975.

Jentzsch, I., & Dudschig, C. (2009). Why do we slow down after an error? Mechanisms underlying the effects of posterror slowing. *Quarterly Journal of Experimental Psychology (2006)*, *62*(2), 209–218.

Kenward, M. G., & Roger, J. H. (1997). Small Sample Inference for Fixed Effects from Restricted Maximum Likelihood. *Biometrics*, *53*(3), 983–997.

Khemlani, S., & Johnson-Laird, P. N. (2012). Theories of the syllogism: A meta-analysis. *Psychological Bulletin*, *138*(3), 427-457.

Khemlani, S., & Johnson-Laird, P. N. (2016). How people differ in syllogistic reasoning. *In Proceedings of the 36th Annual Conference of the Cognitive Science Society.* Cognitive Science Society.

Khemlani, S., & Moore, A. (2012). *Evaluative feedback can improve deductive reasoning.*

Maier, M. E., Yeung, N., & Steinhauser, M. (2011). Error-related brain activity and adjustments of selective attention following errors. *NeuroImage*, *56*(4), 2339–2347.

Marco-Pallarés, J., Camara, E., Münte, T. F., & Rodríguez-Fornells, A. (2008). Neural Mechanisms Underlying Adaptive Actions after Slips. *Journal of Cognitive Neuroscience*, *20*(9), 1595–1610.

Notebaert, W., Houtman, F., van Opstal, F., Gevers, W., Fias, W., & Verguts, T. (2009). Post-error slowing: An orienting account. *Cognition*, *111*(2), 275–279.

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*: Oxford University Press.

Rabbitt, P. M. (1966). Errors and error correction in choice-response tasks. *Journal of Experimental Psychology*, *71*(2), 264–272.

Ragni, M., Dames, H., Brand, D., & Riesterer, N. (2019). When Does a Reasoner Respond: Nothing Follows?. In A. K. Goel,C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society.* Cognitive Science Society.

Revlis, R. (1975). Syllogistic reasoning: Logical decisions from a complex data base. In R. Falmagne (Ed.), *Reasoning: Representation and Process*. Erlbaum

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, *306*(5695),443–447.

Roberts, M. J., Newstead, S. E., & Griggs, R. A. (2001). Quantifier interpretation and syllogistic reasoning. *Thinking & Reasoning*, 7(2), 173–204.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, *23*(5), 645–665.

Ullsperger, M., & Danielmeier, C. (2016). Reducing Speed and Sight: How Adaptive Is Post-Error Slowing? *Neuron*, *89*(3),430–432.

Van der Borght, L., Desmet, C., & Notebaert, W. (2016). Strategy changes after errors improve performance. *Frontiers in Psychology*, *6*, 2051.

Wessel, J. R. (2018). An adaptive orienting theory of error processing. *Psychophysiology*, *55*(3), e13041.