

# Rule abstraction, model-based choice, and cognitive reflection

Hilary J. Don<sup>1</sup> · Micah B. Goldwater<sup>1</sup> · A. Ross Otto<sup>2</sup> · Evan J. Livesey<sup>1</sup>

© Psychonomic Society, Inc. 2016

**Abstract** Numerous tasks in learning and cognition have demonstrated differences in response patterns that may reflect the operation of two distinct systems. For example, causal and reinforcement learning tasks each show responding that considers abstract structure as well as responding based on simple associations. Nevertheless, there has been little attempt to verify whether these tasks are measuring related processes. The current study therefore investigated the relationship between rule- and feature-based generalization in a causal learning task, and model-based and model-free responding in a reinforcement learning task, including cognitive reflection as a predictor of individual tendencies to use controlled, deliberative processes in these tasks. We found that the use of rule-based generalization in a patterning task was a significant predictor of model-based, but not model-free, choice. Individual differences in cognitive reflection were significantly correlated with performance in both tasks, although this did not predict variation in model-based choice independently of rule-based generalization. Thus, although there is evidence of stable individual differences in the use of higher order processes across tasks, there may also be differences in mechanisms that these tasks reveal.

**Keywords** Rule vs. feature generalization · Model-based vs. model-free reinforcement learning · Cognitive control · Associative learning · Individual differences

✉ Hilary J. Don  
hdon7006@uni.sydney.edu.au

<sup>1</sup> School of Psychology, University of Sydney, Sydney, NSW 2006, Australia

<sup>2</sup> Center for Neural Science, New York University, New York, NY, USA

Theories of learning and decision making often assume that behavior can arise from two distinct systems—for example, rule-based versus associative reasoning (Sloman, 1996); goal-directed versus habitual behavior (Balleine & O’Doherty, 2010); intentional versus automatic memory (Jacoby, 1991); model-based versus model-free learning (Daw, Niv, & Dayan, 2005). While these theories share a common theme, the characterizations of the two systems vary considerably. For instance, Daw et al. described the model-based system as one that uses a rich model of the environment to prospectively evaluate actions “on the fly” in terms of their immediate goal relevance, distinguishing it from model-free learning (which simply learns to repeat rewarded actions) not only in terms of diverging neurophysiology but also in terms of its flexibility and computational expense. Rule-based reasoning, by comparison, describes a system in which symbolic representations of events and their interrelations are parsed and manipulated in mental computation to arrive at a logical choice or judgment (Sloman, 1996). In contrast to an associative reasoning system, rule-based reasoning is defined in terms of its level of abstraction, use of symbolic representation, and the computational nature of its processing. Rule-based reasoning is often characterized as effortful and purposive, but adept at coping with changed circumstances.

In building a case for the dual-system approach to cognition, authors have emphasized the general commonalities between these theories, going so far as to simply label them System 1 and System 2 (Stanovich & West, 2000; but see Evans & Stanovich, 2013). These authors rightly note that although the processes invoked in any given dual-systems theory are defined in a unique way, they tend to include one process that requires cognitive control and deliberate thought and one that is simpler and relatively automatic. This approach of treating dual-systems models as subclasses of a general theory has been highly influential over recent years, but

despite its popularity, relatively little effort has been made to test whether it is really valid to group specific cognitive capacities in this way. For instance, to date there is little evidence linking abstract rule-transfer (assumed to require relational reasoning) and model-based choice (assumed to require particular aspects of cognitive control) even though they are each analogous to System 2 thinking. This question is pertinent because dual-system theories often make use of evidence from experimental paradigms that share similar properties but are nonetheless different in important ways.

In this study, we took an individual differences approach to addressing this question. McDaniel, Cahill, Robbins, and Wiener (2014; see also Little & McDaniel, 2015) recently found consistent tendencies within individuals to use exemplar- or rule-based strategies across widely varying learning tasks. Specifically, we examined interrelations between patterns of behavior on three cognitive tasks that are specifically relevant to the different theories discussed above, with the expectation that a tendency toward rule-based learning, and greater cognitive reflection more generally, would predict greater reliance upon model-based choice. In recent years, several tasks have been successful in identifying separable response strategies that suggest the involvement of distinct psychological processes (e.g., Shanks & Darby, 1998; Daw, Gershman, Seymour, Dayan, & Dolan, 2011). These tasks (described below) originate from different but conceptually similar lines of research and share in common an important property; choices that take into consideration abstract structure, and choices that are consistent with the formation of simple associations, manifest in distinctly different patterns of behavior. Stable individual differences in these patterns across tasks may reflect a propensity to rely on a particular process. We tested whether the ability to abstract and apply rules in a causal learning context predicts reliance upon a choice strategy that takes advantage of the task structure in reward-based choice. We also examined whether cognitive reflection is predictive of performance in each of these domains.

### Rule-based and feature-based generalization

Shanks and Darby (1998) developed a “patterning” task that neatly dissociates generalization on the basis of simple cue–outcome relationships from generalization on the basis of an abstract rule. In this task, participants assumed the role of a doctor attempting to determine which foods were causing an allergic reaction outcome in a fictitious patient, learning relationships between food cues and the outcome through a process of trial and error. The task design (see Table 1) included examples of a *negative patterning* discrimination, in which two food cues individually predict an outcome (A+/B+) but the combination of those cues predicts no outcome (AB-), and

*positive patterning*, in which two cues individually predict no outcome (C-/D-), but in combination predict an outcome (CD+). Participants can perform accurately on this task by memorizing specific combinations of cues and outcomes. Alternatively, participants can learn an abstract “opposites” rule. That is, individual cues and their compounds predict opposite outcomes. The task also included several incomplete patterning discriminations, including either individual cues (I+/J+ and M-/N-) or compound cues (KL- and OP+). Following training, participants completed a transfer phase without feedback. Participants’ responses to the remaining cues of the incomplete discriminations (e.g., IJ, MN) were critical. Taking novel compound IJ, for example, if participants rely on feature-based generalization, they should predict IJ+, due to the surface similarity between IJ and I+/J+. However, participants using the opposites rule should predict IJ-, despite the fact that the elements had individually predicted an outcome. Shanks and Darby found that efficient learners were more likely to show rule-based generalization than inefficient learners, and their task has also been used to establish a relationship between rule-based transfer and rule-mediated processes in the inverse base-rate effect (Winman, Wennerholm, Juslin, & Shanks, 2005).

### Model-based and model-free choice

The two-step task is a sequentially structured choice task that dissociates model-free and model-based strategies in reinforcement learning, each of which determines how actions are evaluated and selected based on previous experiences (Daw et al., 2005). A model-free strategy repeats actions that have previously been rewarded, consistent with Thorndike’s law of effect. A model-based strategy takes into account a model of the environmental structure, reasoning about action values and current goals in order to plan behavior. To identify the separate contributions of these two strategies, Daw et al. (2011) developed a two-step sequential choice task (see Fig. 1). On each trial, participants made two sequential choices that ultimately resulted in reward, or no reward. Each binary first-stage choice (left vs. right) led probabilistically to one of two second-stage states (S1 vs. S2). In each of these states, participants made a second choice between a different pair of options, each with a different probability of reward. To ensure participants continually searched for the optimal action, probability of reward on each of the second-stage choices changed slowly during the experiment. Each of the first-stage choices led to a particular second-stage state (e.g., left–S1; right–S2) 70 % of the time (common transition), and to the other second-stage state (e.g. left–S2; right–S1) 30 % of the time (rare transition). This transition structure was consistent across the entire experiment. Model-free and model-based strategies make different predictions about the

**Table 1** Patterning task design

Training			Test		
A+	B+	AB-	A?	B?	AB?
C-	D-	CD+	C?	D?	CD?
E+	F+	EF-	E?	F?	EF?
G-	H-	GH+	G?	H?	GH?
I+	J+		I?	J?	<b><i>IJ?</i></b>
		KL-	<b>K?</b>	<b>L?</b>	KL?
M-	N-		M?	N?	<b><i>MN?</i></b>
		OP+	<b>O?</b>	<b>P?</b>	OP?

*Note.* Letters refer to individual food cues. + indicates the presence of an allergic reaction, - indicates the absence of an allergic reaction.

Transfer trials are indicated in bold. Transfer scores in the current study are based on responding to compound transfer trials, indicated in italics (see main text for further details)

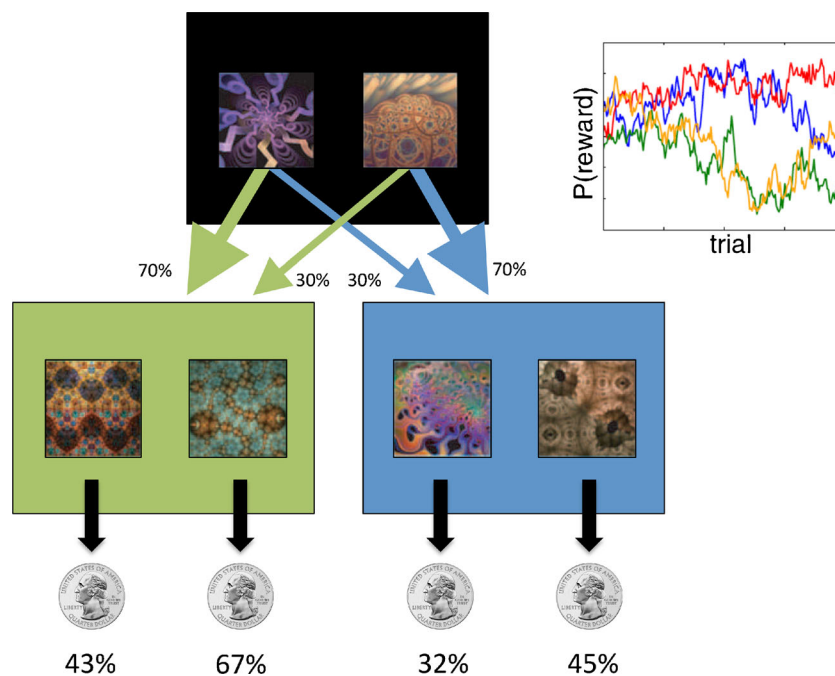
? refers to trials on which no feedback was given

way in which reinforcement and transition structure influence first-stage choice. Take, for example, a first-stage choice that results in a rare transition to a second-stage state (e.g., left-S2), in which a rewarded choice is made. A model-free strategy predicts that the rewarded first-stage choice should be repeated, regardless of whether reinforcement occurred following a common or rare transition (see Fig. 2a). Conversely, a model-based strategy predicts that the

participant should switch their first-stage choice after a rewarded rare transition, as the value of the alternative choice—which commonly leads to the rewarded second-stage state (right)—should increase. Thus, the hallmark of model-based responding is an interaction between reward and transition type (see Fig. 2b). Critically, model-based choice requires participants to learn both the second-stage reward probabilities and the transition structure of the task, and to use this information to prospectively plan subsequent first-stage choice. Typical experiments using this task find contributions of both model-based and model-free strategies, both within individuals and at the population level (Daw et al., 2011). However, a number of participants show responses consistent with purely model-free or purely model-based behavior.

### Cognitive reflection

Several factors influence strategy use in these tasks in similar ways. Under concurrent cognitive load, participants tend to display feature-based generalization (Wills, Graham, Koh, McLaren, & Rolland, 2011) and model-free choice (Otto, Gershman, Markman, & Daw, 2013). Furthermore, working memory capacity (WMC) has been shown to predict rule abstraction in a patterning task (Wills, Barrasin, & McLaren, 2011), and to protect model-based choice from stress effects



**Fig. 1** Two-step task transition structure. Participants make an initial choice (*black screen*) that leads to one of the second stage states (*green or blue screen*). Here, a second choice results in either reward (*coin image*) or nothing. Each first stage choice leads to one particular second-stage state 70 % of the time. The probability of receiving reward

on each of the second-stage states changed slowly over the course the experiment. The four coloured lines in the top-right figure represent an example of how reward probabilities for each of the four second-stage choices change according to independent Gaussian random walks over the course of the experiment

(Otto, Raio, Chiang, Phelps, & Daw, 2013). We hypothesized that an individual's reliance upon reflective, deliberative processes may also capture variation on both tasks. Therefore, we included the Cognitive Reflection Test (CRT) as a predictor of this tendency (Frederick, 2005). The CRT is a 3-item measure, including mathematical problems that trigger an intuitive incorrect answer—for example, “A bat and a ball cost \$1.10 in total. The bat costs a dollar more than the ball. How much does the ball cost?” Here, answering correctly requires the participant to engage in cognitive reflection in order to override the intuitive incorrect response of “10 cents.” The correct solution, “5 cents,” is mathematically simple to derive once the intuitive response is rejected. Performance on the CRT is therefore thought to be an indication of an individual's tendency to engage in conscious, deliberative processes, and has been shown to be a significant predictor of performance on heuristic-and-biases tasks (Toplak, West, & Stanovich, 2011).

Given previous evidence that individuals may be consistent in their tendency to engage a particular mode of processing across tasks (e.g., McDaniel et al., 2014), we expected that (a) individuals who show rule-based generalization will also show model-based choice, (b) individuals who show greater feature-based generalization will show greater model-free behavior, and (c) CRT performance will be related to strategies of both generalization and reward-driven choice.

## Method

### Participants

Sixty-one first-year psychology students from the University of Sydney participated in return for partial course credit. Seven participants were excluded for failing to show reward sensitivity in the two-step task, leaving 54 participants (36 female, mean age = 19.0,  $SD = 1.48$ ).

### Procedure

**Two-step task** Participants first completed 200 trials of the two-step task. On each trial, two fractal images representing the first-stage options appeared side by side on a black background. Participants chose between the left and right options using the “Z” or “?” key, respectively. The background then changed to either blue or green to indicate a transition to the second-stage state, and the selected first-stage image moved to the top of the screen. Each of the first-stage choices would lead to a particular second-stage state 70 % of the time (common transition) and would transition to the other second-stage state on 30 % of trials

(rare transition). Each second-stage state included a different pair of fractal images from which participants were again required to choose. Feedback was then provided while the selected image remained highlighted on-screen (see Fig. 1). Participants were presented with either an image of a coin (reward) or the number zero (no reward). The second-stage reward probabilities changed slowly over the course of the experiment according to independent Gaussian random walks. That is, Gaussian noise ( $SD = 0.025$ ) was added to each reward probability on every trial, bounded between 25 % and 75 %.

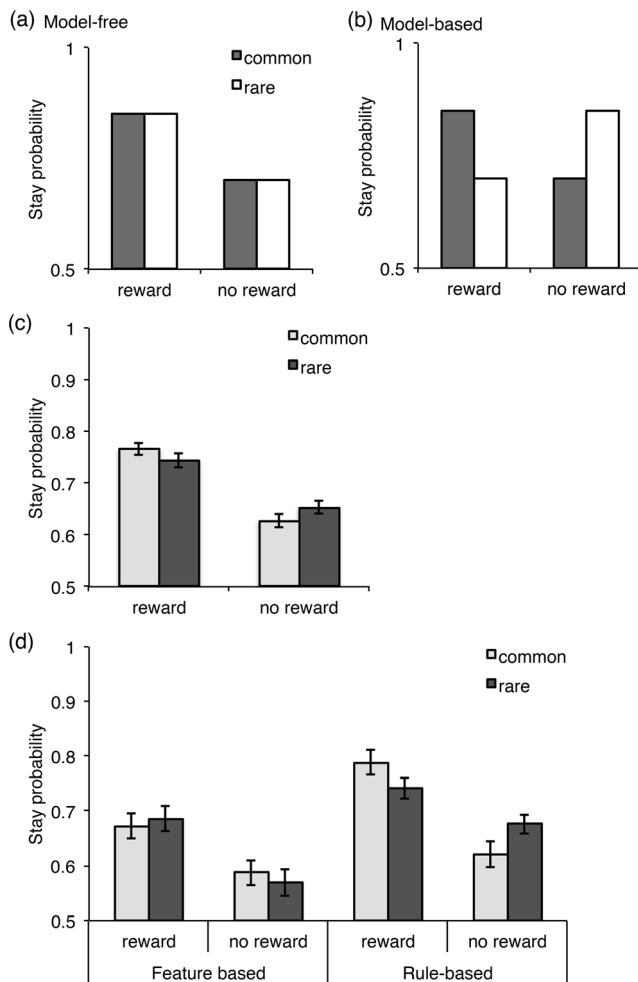
**Patterning task** Participants were then asked to assume the role of a doctor whose task was to determine which foods were causing allergic reactions in a fictitious patient, Mr. X. On each trial, one or two food cues were presented on the upper half of the screen. These cues were images of coffee, banana, fish, lemon, cheese, garlic, apple, eggs, peanuts, mushrooms, strawberry, milk, bread, avocado, broccoli, olive oil, cherries, butter, chocolate, carrots, peach, bacon, peas and prawns, with accompanying labels. Participants were required to predict whether an allergic reaction would occur by choosing from “no allergic reaction” and “ALLERGIC REACTION” boxes on the lower half of the screen. This was followed by corrective feedback provided while the foods remained on-screen. The trial design, shown in Table 1, included four complete patterning discriminations (two positive and two negative) and four incomplete discriminations. Foods were randomly allocated to cues A-P for each participant. There were six blocks of training, with each trial type presented twice per block. The position of compound cues on-screen was counterbalanced within each block. Following training, participants were instructed to use the knowledge they had gained so far in order to rate how likely it was that an allergic reaction would occur, given their patient had eaten the presented foods. On each trial, test cues were presented, and participants made their rating on a visual analogue scale, ranging from *definitely WILL NOT occur* to *definitely WILL occur*. Participants adjusted their rating before pressing the space bar to continue (without feedback). Each test trial was presented twice in random order. Participants were then given a paper-based questionnaire to assess their knowledge of the patterning rule. The first component asked participants to verbalize any rule they noticed during the task. The second required participants to give a yes/no response to two questions that assessed whether participants had noticed the negative and positive patterning rules, based on those used by Harris and Livesey (2008).

**CRT** Finally, participants were asked to complete a paper version of the CRT. The task consisted of three questions taken directly from Frederick (2005).

## Results

### Two-step task

Figure 2c shows the effects of reward and transition type of the previous trial on first-stage choice for all participants in the two-step task. We estimated a mixed-effects logistic regression (Pinheiro & Bates, 2000) to predict the probability of repeating first-stage choice on each trial (stay vs. switch), based on the outcomes of the previous trial. Predictor variables included reward (reward vs. no reward) and transition type (common vs. rare) on the previous trial. Full coefficient estimates are reported in Table 2. There was a significant main effect of reward, revealing an overall tendency to repeat rewarded first-stage choices ( $p < .001$ ), indicating a



**Fig. 2** **a** A model - free strategy predicts that a rewarded first-stage choice is more likely to be repeated regardless of whether reward occurred on a common or rare transition. **b** A model - based choice strategy predicts that reward after rare transitions will influence the following first-stage choice, leading to an interaction between reward and transition type. The center and lower panels show the probability of repeating a first-stage choice when the previous trial was rewarded/non-rewarded, following a common or rare transition, for **(c)** all participants and **(d)** participants who employed feature- and rule - based generalization in the patterning task

**Table 2** Logistic regression coefficients indicating the influence of outcome of previous trial, and transition type of previous trial, upon response repetition in the two-step task. Z-scored transfer scores from the patterning task are also included as a predictor of response repetition.

	Estimate	P value
(Intercept)	1.124	<.001*
<b>Reward</b>	<b>0.439</b>	<b>&lt;.001*</b>
Transition Type	0.032	0.300
Transfer	0.177	0.221
<b>Reward × Transition Type</b>	<b>0.115</b>	<b>0.003*</b>
<b>Reward × Transfer</b>	<b>0.068</b>	<b>0.318</b>
Transition Type × Transfer	0.001	0.978
<b>Reward × Transition Type × Transfer</b>	<b>0.07</b>	<b>0.031*</b>

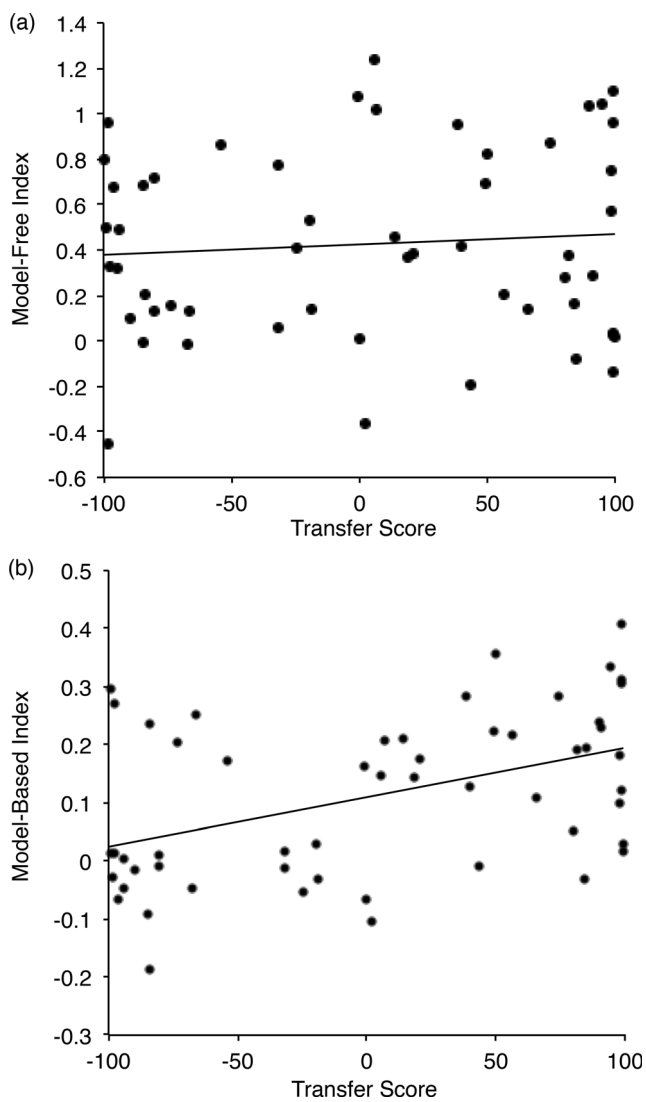
Note. A significant main effect of reward signifies model-free learning. A significant interaction between reward and transition type indicates model-based learning. Results of primary interest are shown in bold.

\* $p < .05$ .

contribution of model-free strategy to choice. There was also a significant interaction between reward and transition type, indicating a model-based contribution to choice ( $p = .003$ ). That is, participants were more likely to switch their previous first-stage choice following a rewarded rare transition than a rewarded common transition, and following a punished common transition than a punished rare transition.

### Patterning task

Analysis of transfer in the patterning task focused on responses to the compound transfer trials, as these make the clearest theoretical predictions for feature- and rule-based generalization. That is, according to associative accounts, responding to individually trained cues presented in combination will be greater than responding to individual elements of a trained compound presented alone, either due to summation of associative strength to individual cues or overshadowing of compound cues (or both; but see Verguts & Fias, 2009). A transfer score was therefore created for each participant by subtracting their average rating for IJ from their average rating for MN. This resulted in scores ranging from -100 to 100, with high scores indicating greater rule-based generalization and low scores indicating greater feature-based generalization ( $M = 3.98, SD = 74.24$ ). Although transfer scores for individual cues (K/L-O/P) are not included in any further analyses, they showed a similar pattern ( $M = 3.87, SD = 50.24$ ) and were significantly correlated with compound transfer cues ( $r = .63, p < .001$ ). Consistent with Shanks and Darby (1998), transfer scores were positively correlated with response accuracy in the final block of training (accuracy range: 58 %–100 %,  $M = 93.16, SD = 9.59; r = .43, p = .001$ ).



**Fig. 3** Scatterplots showing the relationship between the patterning rule transfer score on the x-axis and an index of (a) model-free and (b) model-based choice in the two-step task on the y-axis. These indices reflect each participant's coefficient estimate for reward and reward  $\times$  transition type interaction, respectively, from the logistic regression

In the postexperimental questionnaire, 40 participants were able to verbalize either both the positive or negative patterning rules, or a general opposites rule. Three participants were able to generate only one of the positive or negative patterning rules. In the forced-choice component, 46 participants reported that they noticed both positive and negative patterning rules during the experiment, six reported noticing only one rule, and only two participants reported that they did not notice either rule. Mean transfer scores for participants who did and did not verbalize at least one patterning rule are shown in Table 3. Of the 11 participants who did not verbalize a rule, only one showed (weak) rule-based transfer. Of the 43 participants who verbalized a rule, 29 showed rule-based transfer.

**Table 3** Number of participants showing positive and negative patterning transfer scores, and their mean transfer scores, based on ability to verbalize at least one patterning rule. The total row refers to the overall transfer scores for participants within each column

	+ Verbalize			- Verbalize		
	<i>N</i>	Mean	<i>SEM</i>	<i>N</i>	Mean	<i>SEM</i>
+ Transfer	29	64.64	6.40	1	5.50	0
- Transfer	14	-62.92	9.84	10	-78.77	6.07
Total	43	23.10	10.64	11	-71.11	9.42

### CRT

On average, participants answered 1.0 ( $SD = 1.29$ ) question correctly. Eighty-one percent of incorrect responses were the result of reporting the intuitive answer. Thirty participants (56 %) did not answer any of the questions correctly, while 14 (26 %) answered all three correctly (see Table 4).

### Relationship between tasks

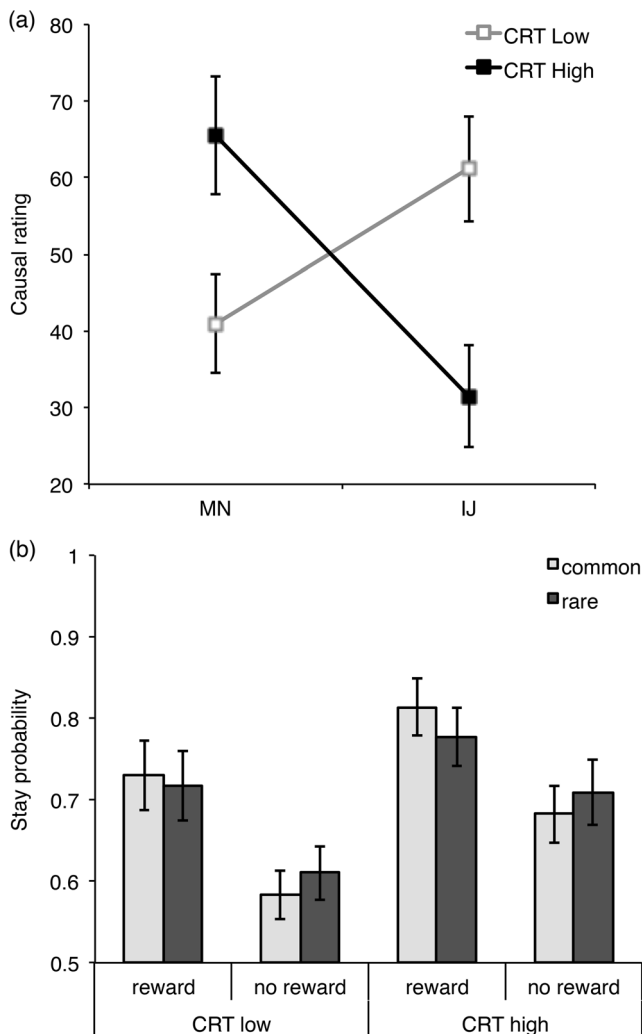
To determine the relationship between generalization and two-step performance, z-scored transfer scores from the patterning task were also included as a predictor on the logistic regression previously described. We found no significant interaction between transfer and reward, suggesting that model-free choice did not increase with increased rule transfer ( $p = .318$ ). However, a significant three-way interaction between transfer, reward, and transition type indicated that higher transfer scores were related to greater model-based choice ( $p = .031$ ). To illustrate this relationship, Fig. 3 plots the relationship between raw transfer scores and separate indices of model-free (top panel) and model-based (lower panel) responding for each participant. These indices reflect individual participants' regression coefficients for reward, and reward  $\times$  transition interaction, respectively, from the group analysis. Figure 2d depicts stay probabilities on the two-step task for participants scoring in the top ( $M = 88.43$ ,  $SD = 12.86$ ) and bottom ( $M = -85.63$ ,  $SD = 13.37$ ) tertiles of transfer scores in the patterning task.

There were significant positive correlations between CRT performance and both rule transfer ( $r = .427$ ,  $p = .001$ ) and participants' model-based index ( $r = .301$ ,  $p = .027$ ) but no relationship between CRT performance and participants' model-free index ( $r = -.018$ ,  $p = .895$ ). To quantify the possible relationships between all three measures, two separate multiple linear regressions were calculated to predict model-based choice based on raw transfer and CRT scores ( $R^2 = .21$ ,  $F(2, 51) = 6.95$ ,  $p = .002$ ), as well as transfer based on CRT and model-based choice ( $R^2 = .29$ ,  $F(2, 51) = 10.58$ ,  $p < .001$ ). The unique variance in model-

**Table 4** Patterning transfer and two-step performance based on CRT score

CRT score/3	N Participants	Patterning Transfer	SEM	MB Index	SEM	MF Index	SEM
0	30	-20.20	12.88	0.08	0.03	0.43	0.08
1	8	3.36	25.62	0.09	0.05	0.41	0.14
2	2	12.00	6.50	0.14	0.00	0.80	0.43
3	14	54.75	17.70	0.18	0.03	0.38	0.11

based performance predicted by transfer was significant, ( $sr^2 = .124, p = .007$ ), however the unique variance predicted by CRT was not ( $sr^2 = .015, p = .329$ ). The unique variance in transfer predicted by CRT was significant ( $sr^2 = .094, p = .012$ ). Performance on both the patterning and two-step tasks based on CRT performance is shown in Table 4. Figure 4 depicts patterning and two-step performance based on a split between high (>0) and low (0) CRT accuracy.



**Fig. 4** Mean (a) patterning generalization and (b) two-step choice based on CRT performance

## Discussion

We found a significant relationship between generalization in the patterning task, which dissociates rule- and feature-based strategies, and model-based choice contributions in a separate sequential choice task. Participants who showed greater rule-based transfer in the patterning task were more likely to take into account the task structure when evaluating choices after a rewarded rare transition in the two-step task. These same participants also tended to display greater cognitive reflection on the CRT.

Interestingly, the ability to verbalize the patterning rule was necessary, but not sufficient for rule transfer. Almost all participants with a positive transfer score verbalized the patterning rule in the manipulation check, and only one participant showed a positive transfer score without accurately verbalizing the rule. This indicates some participants were able to generate the rule when prompted but failed to apply it during the test phase. The significant relationship between CRT performance and rule transfer may give further insight into this finding. Application of an abstract rule may require suppression of a more impulsive feature-driven response during the test phase, even if the rule is detected during training. This explanation is, however, at odds with Wills, Graham, et al. (2011), who found that only concurrent load during training (and not at test) interfered with rule-based generalization, suggesting only extraction (not application) of a patterning rule is effortful. The results of the current study instead suggest that effortful processing may be involved in applying a rule after its acquisition. Recent studies in a similar learning task have shown that attention is relatively inflexible when attempting to overcome established biases, even when explicitly instructed that doing so is beneficial to performance (Don & Livesey, 2015; Shone, Harris, & Livesey, 2015). Therefore, the current finding suggests rule-based processes may require a level of behavioral flexibility and cognitive control in order to overcome stimulus-driven responses when planning and executing action (Evans & Stanovich, 2013; Regehr & Brooks, 1993). This is consistent with recent data connecting cognitive control abilities to model-based choice (Otto et al., 2015).

Model-based choice arguably involves overriding the tendency to repeat previously rewarded choices and thus may also share some of the key qualities of the CRT. While CRT performance was correlated with model-based choice, it did

not explain variance in responding over and above that captured by the patterning task. Specific task requirements may influence the way in which relationships between these tasks are determined. In both the CRT and patterning test phase, performance is measured by a few critical responses after relevant information has been acquired. On the other hand, model-based choice is determined by a continuous succession of discrete responses that may engage choice processes in a different way. This finding also suggests that rule-transfer and model-based choice are related beyond the involvement of reflective thinking processes. Understanding the component mechanisms underlying these tasks requires further research.

Surprisingly, this relationship between the two-step and patterning task was specific to model-based contributions, with no relationship observed between model-free choice and generalization in the patterning task. Greater feature-based generalization did not predict model-free choice behavior (akin to merely repeating previously rewarded actions), despite the view that both of these response patterns are markers of the same effortless, associative system. Selective effects of higher order processes on model-based but not model-free contributions to choice have been demonstrated previously (Gillan, Otto, Phelps, & Daw, 2015; Otto, Skatova, Madlon-Kay, & Daw, 2015). Higher order processes tend to be characterized as flexible and less context specific (Evans & Stanovich, 2013). If the capacity to use these processes varies reliably across individuals, we would expect this use to be consistent across contexts. On the other hand, as feature-based transfer and model-free choice are generally characterized as less flexible and stimulus driven, differences in task requirements may have a greater impact on the expression of these processes, such that associations between them may be less clear, irrespective of whether they are served by a common system.

In summary, rule-transfer and model-based choice appear to be related. Some, but not all, of this relationship is explained by our index of cognitive reflection. Perhaps the largest unanswered question is whether individual differences across these and similar tasks reflect (perhaps unmodifiable) cognitive capacity or (relatively flexible) cognitive strategy. As we did not include measures of WMC or intelligence, it is not possible to make this distinction from the current study. The existing evidence across the literature is mixed. Research in categorization has shown some correlation between rule learning, WMC, and fluid intelligence (McDaniel et al., 2014), but other work has shown that rule-based strategies can be employed independent of cognitive capacity (e.g., Little & McDaniel, 2015). Indeed, Sewell and Lewandowsky (2012) suggested that WMC is more important for categorization response strategy in some tasks than in others. Processing speed (a reliable correlate of fluid intelligence) is related to model-based choice, but only in participants with high WMC (Schad et al., 2014). While WMC can

protect model-based learning from deleterious influences (Otto et al., 2013), WMC itself is not correlated with model-based choice. However, model-based choice is predicted by the use of proactive strategy in cognitive control (Otto et al., 2015), and these control strategies can be modified with some training (Braver, Paxton, Locke, & Barch, 2009). Campitelli and Gerrans (2014) attempted to deconstruct performance on the CRT into components of inhibitory response ability, mathematical competency, and differences in cognitive style, but concluded they could not distinguish cognitive capacity from style.

Although the present study is correlational, manipulating cognitive strategies experimentally may provide further insight into the relationship between these processes and ways to improve effortful performance (e.g., Braver et al., 2009; Alter, Oppenheimer, Epley, & Eyre, 2007). Experimental research has primarily used theory-driven tasks, whereas individual differences research has been primarily data driven, prioritizing high-quality psychometrics. Integrating these approaches is important for the advancement of both. Complete process models require specifying how the same abilities may be differentially deployed in different tasks, for instance, by holding recent trials in mind to either update reward probabilities in one task or to compare past and present stimulus features in another. The approach taken here—characterizing what kinds of individual performance are stable across tasks—will be a critical piece in addressing these questions in the future.

**Author note** This research was supported by an Australian Research Council Discovery Grant DP150104267 to M. B. G. and E. J. L., and a Visiting International Collaborator Support grant from the University of Sydney awarded to A. R. O.

We thank Damian Birney for helpful discussion regarding the analyses.

## References

- Alter, A. L., Oppenheimer, D. M., Epley, N., & Eyre, R. N. (2007). Overcoming intuition: Metacognitive difficulty activates analytic reasoning. *Journal of Experimental Psychology: General*, *136*, 569–576.
- Balleine, B. W., & O'Doherty, J. P. (2010). Human and rodent homologues in action control: Corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology*, *35*, 48–69.
- Braver, T. S., Paxton, J. L., Locke, H. S., & Barch, D. M. (2009). Flexible neural mechanisms of cognitive control within human prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(18), 7351–7356.
- Campitelli, G., & Gerrans, P. (2014). Does the cognitive reflection test measure cognitive reflection? A mathematical modeling approach. *Memory & Cognition*, *42*(3), 434–447.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*, 1204–1215.



- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*, 1704–1711.
- Don, H. J., & Livesey, E. J. (2015). Resistance to instructed reversal of the learned predictiveness effect. *The Quarterly Journal of Experimental Psychology*, *68*, 1327–1347.
- Evans, S. J. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives in Psychological Science*, *8*, 223–241.
- Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, *19*, 25–42.
- Gillan, C. M., Otto, A. R., Phelps, E. A., & Daw, N. D. (2015). Model-based learning protects against forming habits. *Cognitive, Affective, & Behavioral Neuroscience*, *15*, 523–536.
- Harris, J. A., & Livesey, E. J. (2008). Comparing patterning and biconditional discriminations in humans. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*, 144–154.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541.
- Little, J. L., & McDaniel, M. A. (2015). Individual differences in category learning: Memorization versus rule abstraction. *Memory & Cognition*, *43*, 283–297.
- McDaniel, M. A., Cahill, M. J., Robbins, M., & Wiener, C. (2014). Individual differences in learning and transfer: Stable tendencies for learning exemplars versus abstracting rules. *Journal of Experimental Psychology: General*, *143*, 668–693.
- Otto, A. R., Gershman, S. J., Markman, A. B., & Daw, N. D. (2013). The curse of planning: Dissecting multiple reinforcement-learning systems by taxing the central executive. *Psychological Science*, *24*, 751–761.
- Otto, A. R., Raio, C. M., Chiang, A., Phelps, E. A., & Daw, N. D. (2013). Working-memory capacity protects model-based decision-making from stress. *Proceedings of the National Academy of Sciences*, *110*, 20941–20946.
- Otto, A. R., Skatova, A., Madlon-Kay, S., & Daw, N. D. (2015). Cognitive control predicts use of model-based reinforcement learning. *Journal of Cognitive Neuroscience*, *27*, 319–333.
- Pinheiro, J. C., & Bates, D. M. (2000). *Mixed-effects models in S and S-PLUS*. New York, NY: Springer.
- Regehr, G., & Brooks, L. R. (1993). Perceptual manifestations of an analytic structure: The priority of holistic individuation. *Journal of Experimental Psychology: General*, *122*, 92–114.
- Schad, D. J., Jünger, E., Sebold, M., Garbusow, M., Bernhardt, N., Javadi, A.-H., . . . Huys, Q. J. M. (2014). Processing speed enhances model-based over model-free reinforcement learning in the presence of high working memory functioning. *Frontiers in Psychology*, *5*, 1450.
- Sewell, D. K., & Lewandowsky, S. (2012). Attention and working memory capacity: Insights from blocking, highlighting, and knowledge restructuring. *Journal of Experimental Psychology: General*, *141*, 444–469.
- Shanks, D. R., & Darby, R. J. (1998). Feature- and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 405–415.
- Shone, L. T., Harris, I. M., & Livesey, E. J. (2015). Automaticity and cognitive control in the learned predictiveness effect. *Journal of Experimental Psychology: Animal Learning and Cognition*, *41*, 18–31.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*(1), 3–22.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, *23*, 645–726.
- Toplak, M. E., West, R. F., & Stanovich, K. E. (2011). The Cognitive Reflection Test as a predictor of performance on heuristics-and-biases tasks. *Memory & Cognition*, *39*, 1275–1289.
- Verguts, T., & Fias, W. (2009). Similarity and rules united: Similarity- and rule-based processing in a single neural network. *Cognitive Science*, *22*, 243–259.
- Wills, A. J., Barrasin, T. J., & McLaren, I. P. L. (2011). Working memory capacity and generalization in predictive learning. In L. Carlson, C. Hölscher, & T. Shipley (Eds.), *Proceedings of the 33rd annual conference of the Cognitive Science Society* (pp. 3205–3210). Austin, TX: Cognitive Science Society.
- Wills, A. J., Graham, S., Koh, Z., McLaren, I. P. L., & Rolland, M. D. (2011). Effects of concurrent load on feature- and rule-based generalization in human contingency learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *37*, 308–316.
- Winman, A., Wennerholm, P., Juslin, P., & Shanks, D. R. (2005). Evidence for rule-based processes in the inverse base-rate effect. *The Quarterly Journal of Experimental Psychology*, *58A*, 789–815.