



Toward a formal theory of proactivity

F. Lieder¹ · G. Iwama¹

Accepted: 22 February 2021 / Published online: 15 March 2021
© The Author(s) 2021

Abstract

Beyond merely reacting to their environment and impulses, people have the remarkable capacity to proactively set and pursue their own goals. The extent to which they leverage this capacity varies widely across people and situations. The goal of this article is to propose and evaluate a model of proactivity and reactivity. We proceed in three steps. First, we model proactivity in a widely used cognitive control task known as the AX Continuous Performance Task (AX-CPT). Our theory formalizes an important aspect of proactivity as meta-control over proactive and reactive control. Second, we perform a quantitative model comparison to identify the number and nature of meta-control decisions that are involved in the regulation of proactive behavior. Our findings suggest that individual differences in proactivity are governed by two independent meta-control decisions, namely deciding whether to set an intention for what to do in a future situation and deciding whether to recall one's intentions when the situation occurs. Third, we test the assumptions and qualitative predictions of the winning model against data from numerous experiments varying the incentives, cognitive load, and statistical structure of the task. Our results suggest that proactivity can be understood in terms of computational models of meta-control. Future work will extend our models from proactive control in the AX-CPT to proactive goal creation and goal pursuit in the real world.

Keywords Proactivity · Meta-control · Cognitive control · Proactive control · Computational modeling

Introduction

Neuroscience and psychology have extensively studied how the brain processes and reacts to external stimuli and how those propensities are shaped by learning. This has led to significant progress in our understanding of important processes, such as habit formation, and the underlying neural mechanisms, such as reinforcement learning (Schultz et al., 1997). However, these theories do not address an essential element of what makes us human. That is, people do not merely react to their immediate environment and their drives and impulses. They also take initiative to set and pursue their own goals even if nothing in their immediate surroundings would suggest it. This quality is known as *proactivity* (Parker et al., 2010). So far, proactivity has been primarily

studied in organizational studies, management sciences, and applied psychology, whereas part of its underlying cognitive mechanisms have been studied in the field of cognitive neuroscience (Braver, 2012). Proactivity is strongly associated with motivation and positive outcomes in real-world settings. For example, self-reported differences in proactivity in everyday life have been found to be predictors of job performance, career success, and career satisfaction (Judge & Kammeyer-Mueller, 2007). Not surprisingly, proactivity also is correlated with conscientiousness, self-efficacy, and responsibility for change (Tornau & Frese, 2013). Despite all the benefits of proactivity, high levels of proactivity are relatively rare (Seibert et al., 1999). Whether people think and act proactively depends on motivational, dispositional, and situational factors (Parker et al., 2019). In the absence of proactivity, human behavior is frequently controlled by stimulus-driven habitual or Pavlovian mechanisms instead of reflective and goal-directed decision mechanisms (Dolan & Dayan, 2013; van der van der Meer et al., 2012). We refer to this mode of behavioral control as *reactivity*.

By contrast, we define proactivity as the set of mechanisms that generate goal-directed behavior through the exertion of some form of cognitive control. This includes at least two

✉ F. Lieder
falk.lieder@tuebingen.mpg.de

¹ Rationality Enhancement Group, Max Planck Institute for Intelligent Systems, Tübingen, Max Planck Ring 4, D-72076 Tübingen, Germany

distinct classes of mechanisms. The first pathway to proactive behavior is to set intentions in anticipation of future situations, actively maintain them in working memory, and then enact them when the anticipated situation occurs. For instance, if your fridge is empty, you may anticipate the need to override your habit to go home straight from work by the goal-directed behavior of stopping at the supermarket on your way home. Based on the anticipation that you will pass by the supermarket, you may start exerting cognitive control while you are still at work to create and actively memorize the implementation intention: “When the bus announces the stop next to the supermarket, then I will request it to stop and get out.” Recent research in cognitive neuroscience has begun to study this capacity under the heading of *proactive control* (Braver, 2012). Proactive control and its neural underpinnings have been studied in working memory paradigms (Braver, 2012; Burgess et al., 2011) and cognitive control tasks (Mäki-Marttunen et al., 2019a, 2019b). In working memory paradigms, proactive control manifests as the active maintenance of the intention to press a button if the test item matches one of the items to be memorized throughout the delay period. In cognitive control tasks, proactive control manifests as people initiating their response to an anticipated stimulus before it has even appeared based on the information provided by a predictive cue. A second pathway to proactive behavior comprises the stimulus-triggered recall of a goal or intention and the exertion of cognitive control to resolve conflicts between the recalled intention and default activities (automaticity). This mechanism is known as *reactive control* (Braver, 2012). For instance, if you are daydreaming on your way home, then the announcement of the next bus stop might reactivate your goal to go shopping. This, in turn, might prompt you to exert cognitive control to stop daydreaming, set the intention to stop the bus, and enact it immediately. Our definitions of proactivity and reactivity should not be confused with the notions of proactive control and reactive control as defined in the DMC framework. Rather, our notion of proactivity subsumes both proactive and reactive control and reactivity denotes the stimulus-driven automatic behavior that occurs in the absence of either form of cognitive control.

Proactive and reactive control have been extensively studied in a paradigm called the AX Continuous Performance Task (AX-CPT). As illustrated in Fig. 1, the AX-CPT presents participants with a stream of letters that are grouped into pairs. The first letter of each pair is called the *cue*, and the second letter is called the *probe*. There are two types of cues, called A cues and B cues, and two types of probes, called X probes and Y probes. The participant’s task is to detect AX trials, that is trials in which the pair comprises an A cue and an X probe by pressing button 1 and to press button 2 for all other pairs. Critically, in the standard AX-CPT, the frequencies of the four different trial types (AX, AY, BX, and BY) are such that an A is much more likely to be followed by an X than by a Y, and there is a long

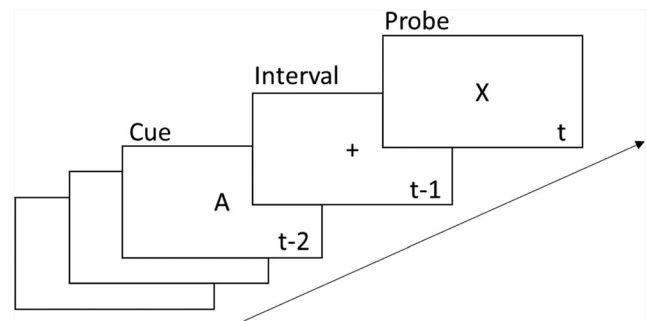


Fig. 1 Illustration of the AX Continuous Performance Task

delay between the cue and the probe. This allows participants to mentally prepare their response to probe even before it appears. For instance, when a participant sees an A cue, they might resolve to press button 1 as soon as the probe appears. This is an example of proactive control. On BX trials, by contrast, participants often have to engage in reactive control to override their habit to press button 1 when they see an X probe.

Proactivity has been found to be highly variable in laboratory paradigms (Braver, 2012) and in the real-world (Seibert et al., 1999). The Dual Mechanisms of Control framework (Braver, 2012) seeks to explain this variability in terms of differences in the extent to which people rely proactive control and reactive control. Despite initial modeling work (De Pisapia & Braver, 2006; Reynolds et al., 2006), the computational principles and algorithmic mechanisms of how people decide when to engage proactive control and when to engage reactive control remain unknown, and several theoretical conundrums remain to be resolved. For instance, it is still debated whether proactive control and reactive control are competing or complementary forms of control (Gonthier et al., 2016; Mäki-Marttunen et al., 2019a). Furthermore, it remains unclear how many and which meta-control decisions govern the variability in cognitive control within and across people. For instance, it is currently controversial whether people inhibit the intentions they have set in anticipation of one event when another event occurs that renders their intentions maladaptive (Mäki-Marttunen et al., 2018). The findings reviewed by Braver (2012) suggested that the frequency with which people engage proactive control increases with factors that make it more valuable (e.g., incentives) and decreases with factors that make it costlier (e.g., working memory load). These findings are congruent with the theory that people make rational use of their limited cognitive resources (Lieder & Griffiths, 2020) and the rational cost-benefit analysis postulated by the Expected Value of Control theory (Lieder et al., 2018; Shenhav et al., 2017). According to this theory, the identity and intensity of cognitive control signals are chosen to maximize the expected reward of performing the task minus the cost of control. The optimal control signal usually has an intermediate intensity, because stronger control signals are costlier. We hypothesize that the exertion of proactive and reactive control is governed by an equivalent rational cost-

benefit analysis. Whether this hypothesis holds up to the scrutiny of testing its quantitative predictions remains to be seen. Last but not least, there is still a large gap between the low-level concepts of proactive and reactive control studied in cognitive neuroscience and the high-level concepts of proactivity and reactivity studied in organizational psychology and the management sciences.

We address these open theoretical questions by formalizing the foundational ideas of the Dual Mechanisms of Control framework (Braver, 2012) with a formal computational model of meta-control over proactive control and reactive control in the AX-CPT and testing its predictions against human performance in five previously conducted experiments. Our model builds on previous findings suggesting that the brain selects between alternative control mechanisms (*meta-control*) according to some kind of cost-benefit analysis (Boureau et al., 2015; Daw et al., 2005; Keramati et al., 2011; Lieder et al., 2018; Lieder & Griffiths, 2017; Shenhav et al., 2017). Viewing proactive control and reactive control as two complementary mechanisms of proactivity and goal-directed behavior, we developed and tested a formal computational model of the meta-control decisions that determine whether, when, and how a person engages proactive and/or reactive control, and how those meta-control decisions depend on situational and personal factors. The resulting meta-control model allowed us to explain individual differences in proactivity and how people's propensities to engage proactive control and reactive control depend on incentives for speed and accuracy (Mäki-Marttunen et al., 2019b), cognitive load (Mäki-Marttunen et al., 2019b), the statistical structure of the task environment (Gonthier et al., 2016; Redick, 2014), and working memory capacity (Redick, 2014).

The outline of this paper is straightforward: we first introduce four alternative meta-control models of proactivity. We then test these models against each other and apply the best model to explain the findings of numerous experiments that investigated proactive control using the AX Continuous Performance Task (AX-CPT). We find that the available data is best explained by the assumption that people make two independent meta-control decisions about the engagement of proactive control and reactive control, respectively. Our findings suggest that individual differences in proactivity can be understood in terms of formal, rational models of how people tradeoff the costs and benefits of engaging in cognitive control. We close the paper by discussing directions for future work.

Modeling the meta-control mechanisms of proactivity in the AX-CPT

The coexistence of proactive and reactive control in the AX-CP task makes it a good testbed for modeling proactive

control, reactive control, and the meta-control processes that determine whether a person acts proactively or reactively. We therefore formulate our meta-control model of proactivity for the AX-CPT.

To understand the meta-control mechanisms in the AX-CPT, we developed a computational-level theory model of meta-control over proactive and reactive control. Proactive control has costs and benefits. Following previous work (Griffiths et al., 2019; Lieder et al., 2018; Shenhav et al., 2017), we model the function of meta-control over proactive/reactive control as performing a cost-benefit analysis to determine whether the benefits of proactive control outweigh its costs. In this section, we develop a model of how people make meta-control decisions about i) whether to set an intention during the cue presentation and ii) whether to engage reactive control when the probe is present. To illustrate these meta-decisions, we return to the grocery shopping example from the introduction. In this example, the meta-control decision whether to set an intention to go shopping in the evening might be made when remembering the shortage of food at noon. This decision will be informed by the subjective importance of having more groceries, the expected increase in the probability of going grocery shopping if an intention is set, and the cost of setting the intention and remembering throughout the day. The meta-decision whether to boost or inhibit this intention occurs when the bus stop near the supermarket is announced. This decision will depend on whether the encountered situation matches the anticipated situation or not (e.g., supermarket closed or medical emergency), the cost of exerting control, and how likely it is that the intention will be enacted without boosting or inhibiting it. Finally, the meta-decision whether to engage reactive control would be made when a person who did not set an intention is reminded of their shortage of groceries when they see the supermarket. This meta-decision would be informed by the expected benefit of exerting control to stop daydreaming and hit the stop button and the effort that this would take. The meta-decision whether to engage reactive control might also occur when the presence inhibited their intention to go grocery shopping due to an unforeseen event (e.g., a medical emergency) and now faces a new situation that they were not prepared for.

We based the model of proactivity in the AX Continuous Performance Task on previous studies suggesting that reactive and proactive control are independent (Mäki-Marttunen et al., 2019a; Mäki-Marttunen et al., 2019b) and that people use different strategies in different trial types (Irlbacher et al., 2014). Furthermore, Mäki-Marttunen et al. (2019b) found that people's performance in the AX-CPT decreases with the participant's cognitive load, which they manipulated by varying whether there were 1, 2, or 3 letters that could instantiate the A-cue (load = 1, load = 2, and load = 3, respectively).

According to our model the process of response selection includes two stages. The first stage begins with the

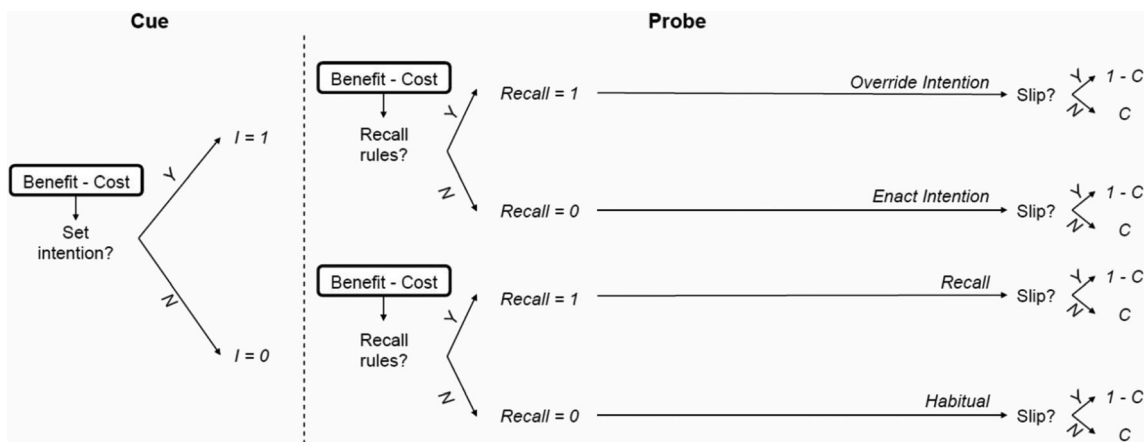


Fig. 2 Our model of meta-control over proactive and reactive control. The variable *I* denotes whether the participant set an intention during the presentation of the cue. *C* is participant’s intended choice, that is the response that they will give unless their finger slips

presentation of the cue and the second stage begins with the presentation of the probe. Figure 2 illustrates the meta-control decisions that our model makes in the first stage and in the second stage, respectively. In the first stage, the meta-controller decides whether to proactively set an intention for how to respond to the probe (X or Y) while the cue (A or B) is being presented. In the second stage, the model’s behavior depends on whether an intention was set in the first stage. If no intention was set in the first stage, then the second stage decides whether to recall the rules and the cue and apply the rules or to react automatically to the probe based on habit. Regardless of whether an intention was set, the second stage decides whether to recall the rules or to enact the intention previously set. Thus, according to our model, there are several qualitatively different levels of proactivity in the AX-CPT and a person’s level of proactivity depends on two meta-control decisions that they make during each trial. In the remainder of this section, we detail our mathematical model of these meta-control decisions and the resulting response distributions.

Figure 2 illustrates how these meta-control decisions (i.e., intention setting and recalling the cue and rules) work together to determine the response to a cue-probe pair. According to our model, the probability of making the correct choice ($C = 1$) is the weighted average of the accuracies entailed by having set an intention ($I = 1$) versus not having set an intention ($I = 0$) given the model parameters θ , that is

$$\begin{aligned}
 P(C = 1 | \text{probe, cue, load, } \theta) &= P(I = 1 | \text{cue, probe, } \theta) \\
 &\cdot P(C = 1 | I = 1, \text{load, cue, probe, } \theta) \\
 &+ P(I = 0 | \text{cue, probe, } \theta) \\
 &\cdot P(C = 1 | I = 0, \text{load, cue, probe, } \theta).
 \end{aligned}
 \tag{1}$$

The model parameters $\theta = (u_+, u_-, u_{\Delta t}, \lambda, \gamma, \delta)$ are summarized in Table 1 and will be explained one by one as we develop the model throughout the remainder of this section. Whether or not an intention is set ($P(I = 1 | \text{cue, probe, load, } \theta)$) is determined by the meta-control decision made in response to the cue (Stage 1). Upon the presentation of the probe, the model makes a meta-control decision about whether or not to recall the rules (Stage 2). These meta-control decisions jointly determine the predicted accuracies in the scenario where an intention was set ($P(C = 1 | I = 1, \text{load, cue, probe, } \theta)$) and the scenario where no intention was set ($P(C = 1 | I = 0, \text{load, cue, probe, } \theta)$).

Stage 1: Deciding whether to proactively set an intention

To model the first stage (Fig. 2), we assume that, during the presentation of the cue, the person have the decision to set or not a behavioral intention, to respond affirmatively or not when the probe will be presented (e.g., set the intention to

Table 1 Explanation of the model’s parameters $\theta = (u_+, u_-, u_{\Delta t}, \lambda, \gamma)$

Parameters	Explanation
u_+	subjective utility of correctly detecting an AX trial.
u_-	subjective utility of correctly reporting that a trial was not an AX trial.
$u_{\Delta t}$	subjective utility of responding quickly.
λ	intensity of the deleterious effect of cognitive load on controlled processing
γ	cost of setting an intention and maintaining it in working memory

press “left” if a B cue is presented). We formalize this as a binary decision whether or not to proactively set an intention ($I = 1$ vs. $I = 0$) based on the cue (A or B) made according to an approximate cost-benefit analysis in which the meta-controller evaluates whether the expected value (\mathbb{E}) of the benefits of engaging in proactive control outweighs its costs, that is whether

$$\mathbb{E}[\text{Benefit}(I = 1) \mid \text{cue}, u_+, u_-, u_{\Delta t}, \text{load}] > \text{cost}(I = 1),$$

where the parameters u_+ and u_- determine the utility or reward of making an accurate response specifically for AX or other trial types respectively and $u_{\Delta t}$ determines the utility of

making a fast response, and *load* is the contextual load determined by the number of A-cues.

The benefit and the cost of proactively setting an intention ($I = 1$) are measured relative to the reward and cost of not setting an intention ($I = 0$). According to our model, the meta-control decision should be based on the expected benefit of proactive control given the information provided by the cue. Given that the cost-benefit analysis is approximate and that there is uncertainty about the relevant values, we model the probability that the meta-control system decides to engage in setting an intention as

$$P(I = 1 \mid \text{cue}; \theta, \text{load}) = \frac{\exp(\mathbb{E}[\text{Benefit}(I = 1) \mid \text{cue}, u_+, u_-, u_{\Delta t}, \text{load}] - \text{cost}(I = 1))}{1 + \exp(\mathbb{E}[\text{Benefit}(I = 1) \mid \text{cue}, u_+, u_-, u_{\Delta t}, \text{load}] - \text{cost}(I = 1))}. \quad (2)$$

The cost-benefit analysis and the resulting probabilities of setting an intention in response to the A cue and the B cue are presented in the [Supplementary Material](#).

Figure 3 shows our model’s predictions of the probability that people will proactively set an intention when they see a cue depending on the identity of the cue, cognitive load, and the reward for accuracy, u_+ and u_- , assuming that the cost of setting an intention and maintaining it in working memory is $\gamma = \frac{1}{3}$ and the utility of responding faster is $u_{\Delta t} = 0.2$. Note that the addition of reward for accurate responses increases the probability to set an intention for B cues and decreases it for A cues because B cues are more informative of the correct response for that trial. Overall, it increases the model’s propensity towards proactive control. Thus, overall, proactivity should increase with reward and decrease with cognitive load.

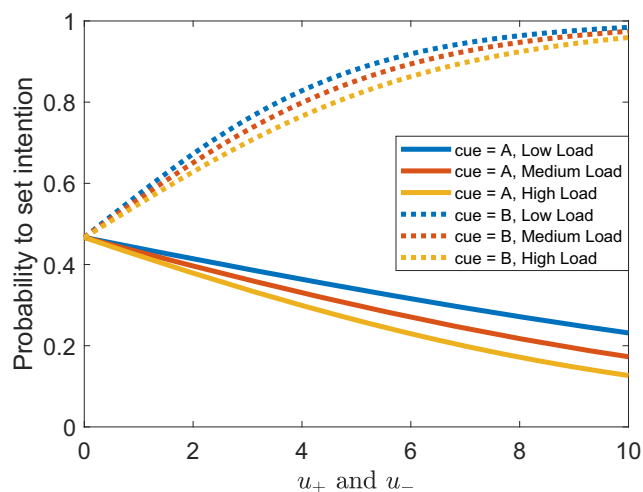


Fig. 3 Probability of proactive intention setting in response to the cue (A vs. B) across low, medium, and high cognitive loads (load = 1, 2, and 3, respectively) depending on the reward for responding accurately (u_+ and u_-), assuming $\gamma = \frac{1}{3}$, $\lambda = 0.05$, and $u_{\Delta t} = 0.2$

Our cost-benefit analysis model predicts that if people can switch between proactive and reactive control on a trial-by-trial level, then we should expect to see more proactive intention setting in response to B-cues than to A-cues. Furthermore, the cost-benefit analysis suggests that there could be interaction effects between cost and reward such that for very high and very low reward, the effect of load should be smaller than for intermediate levels of reward for B trials. Furthermore, the effect of cognitive load on intention setting should be higher on A-trials than on B-trials.

If the meta-control process has determined that an intention should be set, then intention setting proceeds by predicting the next stimulus and initiating the response to the predicted stimulus before it has even appeared. For instance, because the A is followed by an X in 87.5% of the time, proactive control will usually predict that the next stimulus will be an X and initiate an affirmative response to the upcoming probe. We therefore formally model proactive control as follows:

1. Make a prediction \hat{s} about a future state (e.g., S_{t+1}) by sampling from the predictive model ϑ ($\hat{s} \sim P(S_{t+1} \mid S_t = \hat{s}_t; \vartheta)$) that has been learned through experience¹ (Vul et al., 2014).
2. Plan one or more actions, A_{t+1} , to be taken in the predicted future state(s), \hat{s} , by maximizing the expected utility, \mathbb{E} ,

¹ This learning process can be formally modeled as Bayesian inference on the probabilities $\theta_{j,k}$ of the transitions from state s_j (e.g., stimulus A) to state s_k (e.g., stimulus X). This inference is drawn from all previously observed state transitions (E_1, \dots, E_{t-1}) according to Bayes theorem (i.e., $P(\theta \mid E) = \frac{P(E \mid \theta) \cdot P(\theta)}{P(E)}$). The likelihood function can be modeled in terms of one multinomial distribution for each previous state (i.e., $P(S_{t+1} \mid S_t = s_t) = \text{Multinomial}(\theta_{i,1}, \dots, \theta_{i,n})$) and the prior distribution can be modelled as a Dirichlet distribution (i.e., $P(\theta_i) = \text{Dirichlet}(1, \dots, 1)$). In this way, the learning process can be modelled according to the analytic update equations of the Dirichlet-multinomial model.

of the resulting outcomes, O_{t+1} , (e.g., $A_{t+1} = \operatorname{argmax}_a \mathbb{E}[u(O_{t+1}) | S_{t+1} = \hat{s}, a]$). To simulate the AX-CPT, we model the utility of correctly detecting the AX stimulus by the free parameter u_+ , and we model the utility of correctly withholding that response in its absence by the free parameter u_- , assuming that the utility of incorrect responses is 0.

3. Exert cognitive control to create the intention to execute the planned action(s) (e.g., take action a when the next stimulus appears), commit it to working memory, and actively maintain the memory of that intention.

For instance, upon seeing an A-cue in the AX-CPT with 70% AX trials and 10% AY trials, there is a 7 out of 8 chance that our model will predict that the probe will be an X. In that case, it will set the intention to give an affirmative response (e.g., “When the probe appears, I will click the button for reporting an AX-trial”). Alternatively, there is also a 1 in 8 chance that our model will predict that the A cue will be followed by a Y probe. In that case, it will set the intention to respond negatively (e.g., “When the probe appears, I will click the button for reporting that this is NOT an AX-trial”).

Stage 2: Meta-control over the response to the probe

When participants encounter the probe, they may recall the cue and apply the rule for how to respond to the observed

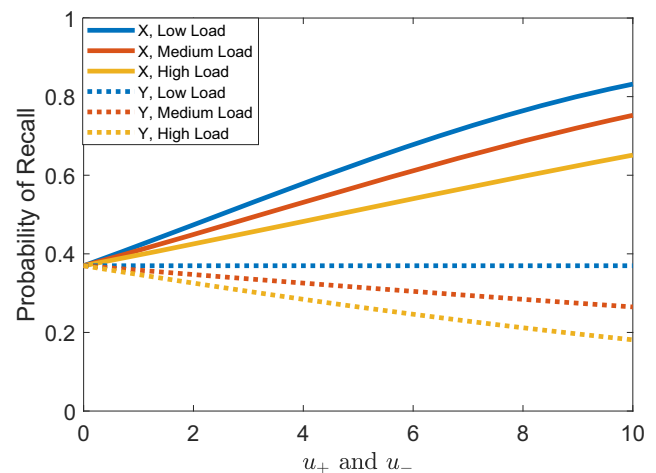


Fig. 4 Probability of recalling the cue and rules in response to the probe (X vs. Y) across low, medium and high cognitive loads (load = 1, 2 and 3, respectively), depending on the reward for responding accurately (u_+ and u_-) assuming $\gamma = \frac{1}{3}$, $\lambda = 0.05$, and $u_{\Delta t} = 0.2$

cue-probe ($c_{\text{recall}} = 1$) or not ($c_{\text{recall}} = 0$). We postulate that which of these two modes govern people’s responding is determined by a rational cost-benefit analysis (see Section 1.3 of the [Supplementary Material](#)). In brief, we model the meta-control decision whether or not to recall the cue and rules (reactive control) according to the softmax decision rule

$$P(c_{\text{recall}} = 1 | \text{probe}, u_+, u_-, \gamma, u_{\Delta t}, \text{load}) = \frac{\exp(\mathbb{E}[\text{Benefit}(c_{\text{recall}} = 1) | \text{probe}; u_+, u_-, u_{\Delta t}, \text{load}] - \gamma)}{1 + \exp(\mathbb{E}[\text{Benefit}(c_{\text{recall}} = 1) | \text{probe}; u_+, u_-, u_{\Delta t}, \text{load}] - \gamma)}, \tag{3}$$

where γ is the cost of recalling the cue and rules. The derivation of the benefit term is presented in the [Supplementary Material](#). Since recalling and relying on the cue and the rules is a binary event, the probability of responding habitually is $P(c_{\text{recall}} = 0 | \text{probe}; u_+, u_-, \gamma, u_{\Delta t}, \text{load}) = 1 - P(c_{\text{recall}} = 1 | \text{probe}; u_+, u_-, \gamma, u_{\Delta t}, \text{load})$.

Figure 4 shows our model’s predictions of the probability of recalling the cue and rules across different rewards for being accurate (u_+ and u_-) for a fixed reward for being fast ($u_{\Delta t} = 0.2$) and a fixed cost of holding information in working memory ($\gamma = \frac{1}{3}$). The probability of recalling in response to X probes is always predicted to be higher than for Y probes, since recalling the rules has no benefit when the automatic response would be correct as well. The model also predicts a higher probability of recalling for low contextual load conditions, since the probability of recalling the correct response decreases as the contextual load increases.

Predicting people’s accuracy in the AX-CPT

To complete our model, we specify how likely the response is to be correct depending on whether or not an intention was set in response to the cue. The probability that the choice will be correct, $P(C = 1)$, after an intention has been set ($I = 1$) is

$$P(C = 1 | I = 1, \text{cue}, \text{probe}, \text{load}, \theta) = P(c_{\text{recall}} = 1 | \text{probe}; u_+, u_-, \gamma, u_{\Delta t}, \text{load}) \cdot P(C_R = 1 | \text{cue}, \text{probe}, \text{load}) + P(c_{\text{recall}} = 0 | \text{probe}; u_+, u_-, \gamma, u_{\Delta t}, \text{load}) \cdot P(C_I = 1 | \text{cue}, \text{probe}, \text{load}) - \lambda \cdot (\text{load} - 1), \tag{4}$$

where $C_I = 1$ means that the intention was correct and the term $-\lambda \cdot (\text{load} - 1)$ models the deleterious effect of cognitive load on cognitive control with the parameter λ measuring the

severity of this effect. The probability of the choice being correct given that an intention was set is given in Table 2. The probability that a correct choice will be made in the absence of an intention is

$$\begin{aligned}
 &P(C = 1 | I = 0, \text{cue, probe, load}, \theta) \\
 &= P(c_{\text{recall}} = 1 | \text{probe}; u_+, u_-, \gamma, u_{\Delta t}, \text{load}) \\
 &\quad \cdot P(C_R = 1 | \text{cue, probe, load}) \\
 &\quad + P(c_{\text{recall}} = 0 | \text{probe}; u_+, u_-, \gamma, u_{\Delta t}, \text{load}) \\
 &\quad \cdot P(C_M = 1 | \text{cue, probe, load}) - \lambda \cdot (\text{load} - 1),
 \end{aligned}$$

where $C_R = 1$ means successfully recalling and applying the rule and $C_M = 1$ is the accuracy of probability matching. The probabilities of the choice being correct given successfully recalling and applying the rule are given in Table 2.

Furthermore, we assume that a random error, in which the motor slips cause the button press to be incorrect, $R = 0$, even if the participant’s choice was correct ($C = 1$) and vice versa. We therefore model the probability that the button press is correct as

$$\begin{aligned}
 P_{\text{model}}(R = 1 | \text{cue, probe}, \theta) &= P(C = 1 | \text{probe, cue}, \theta) \cdot (1 - p_{\text{slip}}) \quad (5) \\
 &\quad + P(C = 0 | \text{probe, cue}, \theta) \cdot p_{\text{slip}}.
 \end{aligned}$$

Finally, we assume that the probability that the participant’s hand slips is equal to the error rate in BY trials, that is $p_{\text{slip}} = 0.0125$ for the Mäki-Marttunen et al. (2019a) dataset. This model formalizes the key assumption of the dual mechanisms of control (DMC) framework that people make two independent meta-control decisions about whether to engage proactive control and whether to engage in reactive control according to a rational cost benefit analysis. We therefore refer to this model as our DMC model (m_{DMC}).

Table 2 Probability that the response is correct if it is driven by an intention (2nd column), by automaticity based on probability matching (3rd column), or reactive control (4th column) depending on the trial type. In this example there are 70% AX trials, 10% AY trials, 10% BX trials, and 10% BY trials. The first term of the accuracy of proactive control

Cue, probe	Accuracy of proactive control $P(C_I = 1 \text{cue, probe, load})$	Accuracy of automaticity $P(C_M = 1 \text{cue, probe})$	Accuracy of reactive control $P(C_R = 1 \text{cue, probe, load})$
AX	$\frac{7}{8} - \lambda \cdot (\text{load} - 1)$	$P(AX X) = \frac{7}{8}$	$1 - \lambda \cdot (\text{load} - 1)$
AY	$\frac{1}{8} - \lambda \cdot (\text{load} - 1)$	$P(AX Y) = 1$	$1 - \lambda \cdot (\text{load} - 1)$
BX	$1 - \lambda \cdot (\text{load} - 1)$	$1 - P(AX X) = \frac{1}{8}$	$1 - \lambda \cdot (\text{load} - 1)$
BY	$1 - \lambda \cdot (\text{load} - 1)$	$P(AX Y) = 1$	$1 - \lambda \cdot (\text{load} - 1)$

Alternative meta-control mechanisms

Inhibition of prepotent intentions

Recent physiological data and reaction time data suggest that there might be an additional control mechanism influencing people’s responses in the AX-CPT, namely the inhibition of prepotent responses (Mäki-Marttunen et al., 2018; Mäki-Marttunen et al., 2019b). This means that a participant might see an A cue and set the prepotent intention to report an AX trial in anticipation of an X probe and then inhibit that intention when they see the Y probe. In support of this view, the pupillometry findings of Mäki-Marttunen et al. (2018) and the Locus Coeruleus and dlPFC activation found by Mäki-Marttunen et al. (2019a, 2019b) support the involvement of inhibitory control in the AX-CPT. Furthermore, the especially long response times on AY trials might suggest that participants sometimes override their intention to report an AX pair when they see the Y probe (Mäki-Marttunen et al., 2018; Mäki-Marttunen et al., 2019b). For these reasons, we developed an extension of the model illustrated in Fig. 2 that includes an additional meta-decision about whether to inhibit the proactively set intention when it encounters the probe.

As illustrated in Fig. 6a, this model assumes that if people have proactively set an intention in response to the cue (A or B) then the control system may boost or inhibit reactive control in response to the probe (X or Y). For instance, if the participant set the intention “Click the button for AX trials when the probe appears” in response to the A-cue, then they might inhibit this incorrect intention when they see the Y probe. We assume that the probability that a participant will do this increases with the benefits of being correct and decreases with the cost of control. Conversely, when the participant encounters the anticipated X probe, they might boost their intention to report an AX trial to further increase their probability of being correct if they are highly motivated to be fast and accurate. Furthermore, inhibition might be especially important on no-go trials where a third type of probe signals that the response should be withheld.

(e.g., $\frac{7}{8}$) is the probability that the person predicted the more likely probe when they set their intention. The second term is the loss in accuracy due to cognitive load. *load* is the intensity of the cognitive load given by the number of different letters that could serve as the A-cue (i.e., 1, 2, or 3).

We model this decision as the specification of a control signal $c \in [-c_0, 1 - c_0]$ that shifts the probability that the decision will be determined by the proactively formed intention away from its default probability of $c_0 = 0.7$. A positive control signal boosts the effect of the proactively formed intention, whereas a negative control signal inhibits it. We further assume that cognitive load has an interference cost, λ , that affects the probability to successfully inhibit a proactively set intention. Formally, we assume that the probability that the intention will be inhibited is $P(\text{Inhibit} = 1 | c) = (1 - (c_0 + c)) \cdot (1 - \lambda \cdot \text{load})$.

According to our model, the intensity of the control signal c is chosen according to a cost-benefit analysis. This cost-benefit analysis assumes that the control signal intensity c linearly interpolates between the expected performance of responding with versus without the intention (see Section 1.2 of the [Supplementary Material](#)). Following previous work, we model the cost of inhibition as an exponential function of the absolute value of the control signal intensity, in which δ is the control cost parameter that determines how quickly the cost of control increases with the absolute value of the control signal, that is

$$\text{cost}(c; \delta) = \exp(\delta \cdot |c|) - 1.$$

Therefore, the optimal control signal c^* is

$$c^*(\text{probe}, u_+, u_-, u_{\Delta t}, \delta) = \arg \max_c (\mathbb{E}[\text{Benefit}(c) | \text{probe}, u_+, u_-, u_{\Delta t}, \text{load}] - \text{cost}(c; \delta)).$$

Assuming that the cognitive control system chooses the optimal control signal, the probability that the person will inhibit the intention is

$$P(\text{Inhibit} | \text{probe}, \text{load}, \theta) = (1 - (0.7 + c^*(\text{probe}, u_+, u_-, u_{\Delta t}, \delta))) \cdot (1 - \lambda \cdot \text{load}). \tag{6}$$

Figure 5 shows the optimal control signal intensity as a function of the subjective utility of correct and fast responses depending on the currently presented probe. The plot shows that, given a sufficiently high reward, the model boosts the effect of the proactively set intention for X probes and inhibits the intention for Y probes. The optimal control signal intensity remains zero for a wide range of values for $u_+ = u_- = u_{\Delta t}$, for which the cost of control outweighs any potential benefits.

One can think of this second step as modulating the intensity of the effect of setting a proactive intention in the first step. The control signal can diminish the increase in proactivity or enhance it further. Regardless of the chosen control signal, the level of proactivity should always be higher when a proactive intention was formed in the first step than when it was not.

The inclusion of this additional meta-control decision leads to the following equation for the probability of making a correct choice:

$$P(C = 1 | I = 1, \text{cue}, \text{probe}, \text{load}, \theta) = (1 - P(\text{Inhibit} | \text{probe}, \theta, \text{load})) \cdot P(C_1 = 1 | \text{cue}, \text{probe}, \text{load}) + P(\text{Inhibit} | \text{probe}, \theta, \text{load}) \cdot P(C = 1 | I = 0, \text{cue}, \text{probe}, \text{load}, \theta) - \lambda \cdot (\text{load} - 1).$$

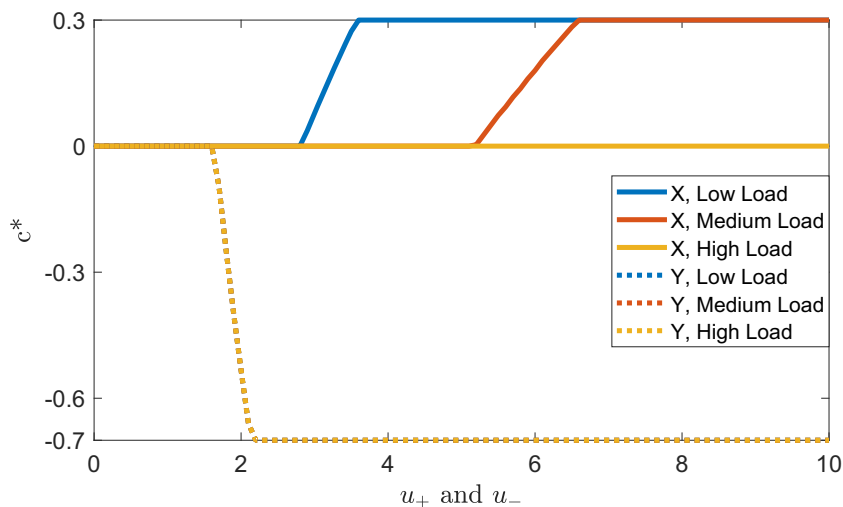


Fig. 5 Optimal control signal intensity depending on the reward for responding accurately (u_+ and u_-) across low, medium and high cognitive loads (load = 1, 2, and 3, respectively) assuming $\delta = \frac{1}{2}$, $\gamma = \frac{1}{3}$, $\lambda = 0.05$, and $u_{\Delta t} = 0.2$

Models according to which proactive and reactive control are mutually exclusive

A key assumption of the model illustrated in Fig. 2 is that people can always override the intention that they have set in response to the cue (proactive control) by invoking reactive control in response to the probe (reactive control). According

to the extended meta-control model illustrated in Fig. 6a, this is possible only when the proactively set intention is inhibited first. A third alternative is that proactive control and reactive control might be mutually exclusive in the sense that people have choose between one or the other but cannot use both. The mutual exclusivity models illustrated in Fig. 6b-c formalize this assumption in two different ways. According to the

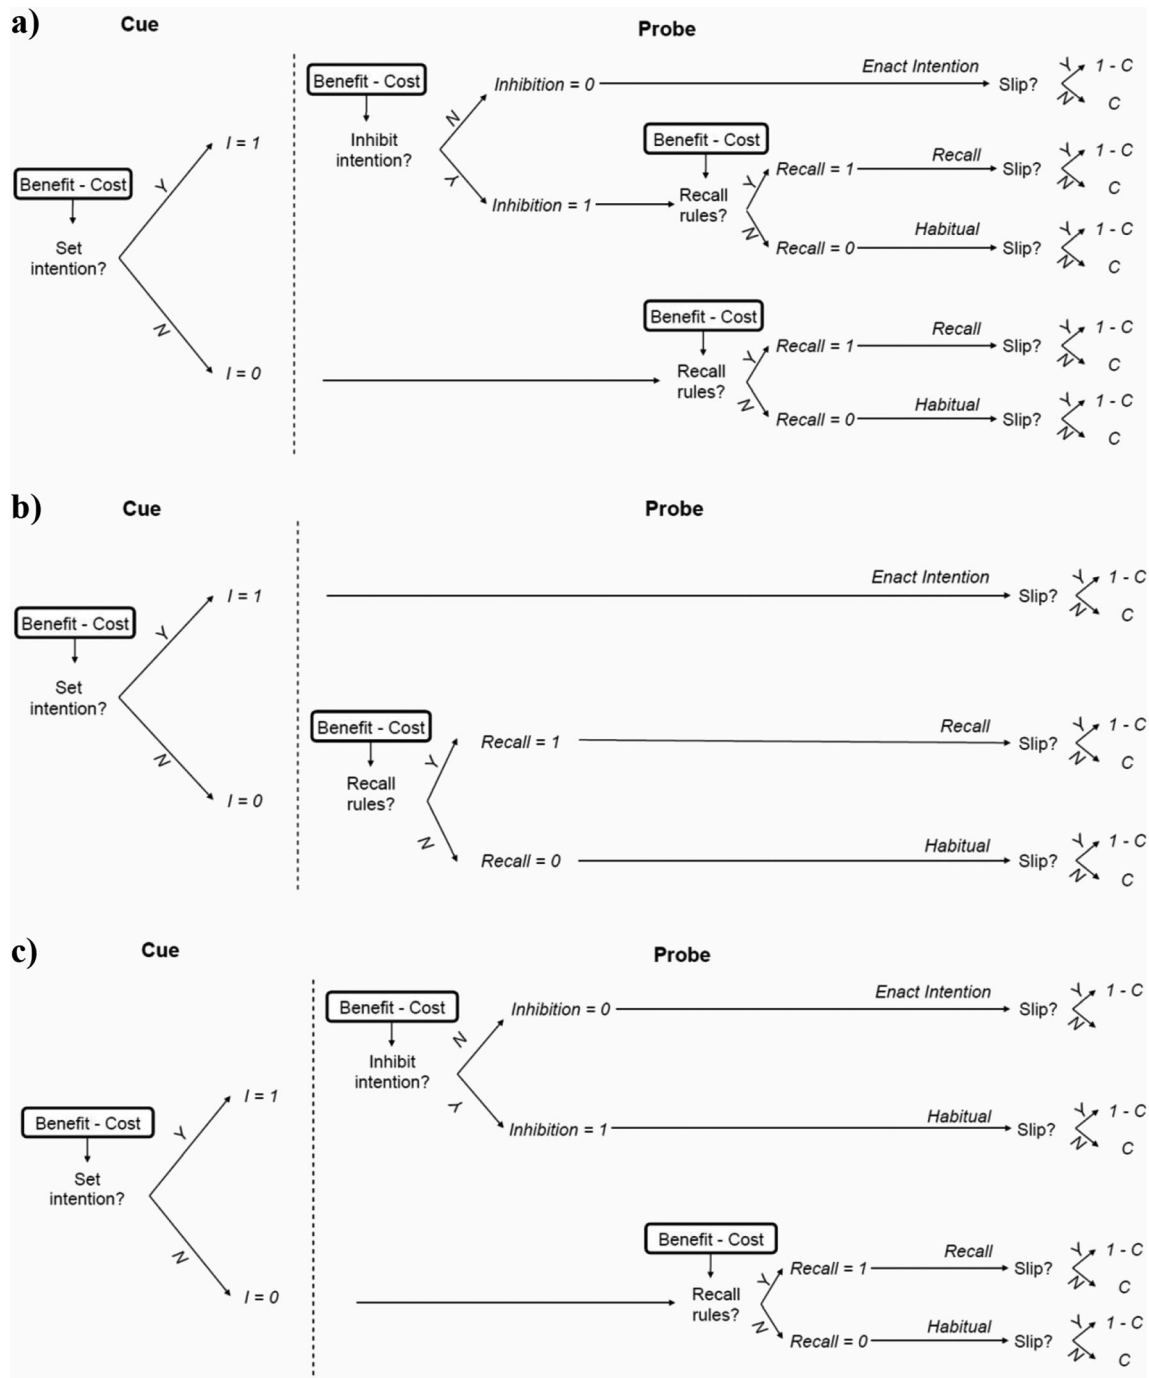


Fig. 6 Alternative models. **a** Extended meta-control model according to which intentions can be inhibited and reactive control is not considered when the person has an active intention. **b** Simpler model according to which the meta-control decision about recalling the rules (reactive

control) is only made when no intention was set (exclusivity model without inhibition). **c** An extension of the model shown in Panel b according to which the proactively set intention can be inhibited (exclusivity model with inhibition)

mutual exclusivity model without inhibition (Fig. 6b) the participant either sets an intention and then enacts it (proactive control) or does not set an intention and then chooses whether to recall the rules (reactive control). According to the mutual exclusivity model with inhibition (Fig. 6c), a participant who has set an intention can choose to inhibit it, but when they do their response is determined by their habits (no control).

Summary

Our model postulates that meta-control over proactivity is governed by two rational cost-benefit tradeoffs that governed whether intentions are set proactively, and whether the cue and rules are recalled and followed. The first meta-decision is made based on the cue, whereas the second one is made based on the probe. All meta-decisions are informed by the incentives for speed and accuracy, cognitive load, and the cost of control. Before testing the predictions of the meta-control model, we will present and test the assumptions about the proactive process.

Results: Explaining variability of control in the AX Continuous Performance Task

Our meta-control model of the variability in proactive behavior comprises three components: 1) its assumptions about which meta-control decisions govern the variation in people's proactivity, 2) its predictions of the accuracies given the outcomes of those meta-control decisions, and 3) its predictions of the meta-control decisions about intention setting and recalling the rules. In this section, we leverage previously published empirical data to examine each of these three components in turn. Each time we use the method described in Section 2 of the [Supplementary Material](#) to fit the parameters of our models to the data.

Testing alternative theories of meta-control over proactivity

To identify which of the four models introduced above best describes meta-control over proactivity in the AX-CPT, we performed a formal model comparison. The distinguishing feature of two of our models is that they inhibit a previously set intention. This feature is most relevant when the task includes no-go trials. We therefore compared the four models on the data set from Experiment 2 by Gonthier et al. (2016). Assuming that the mechanisms of meta-control might differ across participants, we performed one formal quantitative model comparison for each individual participant. To select between the models, we used the Bayesian Information Criterion (Schwarz, 1978). As summarized in Table 3, the results of our model comparison strongly support the DMC model according to which people make two independent meta-control decisions about invoking proactive and reactive control respectively. In fact, the DMC model was the best model for all 92 participants when compared to the extended model and the exclusivity model with inhibition and for 73.9% of the participants compared with the exclusivity model without inhibition.

Testing the model's predictions about the effects of meta-control

If our model of how people set intentions and how intentions affect behavior is correct, then people's response frequencies on AY trials and BX trials should be a weighted average of the response frequencies that our model predicts for the case when an intention is set (proactive control) and the case when no intention is set (no proactive control). As illustrated in Fig. 7a, the data from Mäki-Marttunen et al. (2019a) confirmed this prediction. This suggests that we can understand the average response frequencies across people and experimental

Table 3 Results of the participant-level ($N = 92$) model comparison on the data set by Gonthier et al. (2016) between the DMC model (m_{DMC}) shown in Fig. 2 and the alternative models ($m_{\text{alternative}}$) shown in Fig. 6a-c. The model comparisons were performed based on $\Delta\text{BIC} = \text{BIC}_{m_{\text{DMC}}} - \text{BIC}_{m_{\text{alternative}}}$.

	Extended model (Fig. 6a)	Exclusivity without inhibition (Fig. 6b)	Exclusivity with inhibition (Fig. 6c)
% participants whose data is best explained by this model:	0%	26.1%	0%
% participants for whom the evidence for m_{DMC} over $m_{\text{alternative}}$ is			
very strong ($\Delta\text{BIC} > 10$)	7.6%	-	8.7%
strong ($6 < \Delta\text{BIC} < 10$)	50.0%	-	46.7%
positive ($2 < \Delta\text{BIC} < 6$)	40.2%	10.9%	41.3%
weakly positive ($0 < \Delta\text{BIC} < 2$)	2.2%	63.0%	3.3%
weakly negative ($-2 < \Delta\text{BIC} < 0$)	-	26.1%	-
negative ($-6 < \Delta\text{BIC} < -2$)	-	-	-
strongly negative ($-10 < \Delta\text{BIC} < -6$)	-	-	-

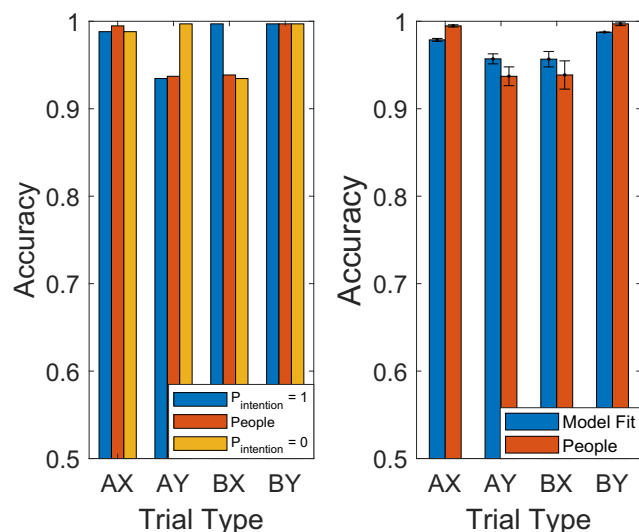


Fig. 7 The plot on the left compares the accuracies of people to the accuracies that our model predicts if people were to always set intentions ($p_{\text{intention}} = 1$) or never set intentions ($p_{\text{intention}} = 0$) for low contextual load (load = 1), and assuming and $p_{\text{recall}} = 0.93$. On the right, mean accuracy when fitting the model across trials of the no reward and low cognitive load condition for each participant with $p_{\text{intention}}$, ($M = 0.58$, $SD = 0.30$) and p_{recall} ($M = 0.93$, $SD = 0.09$) as free parameters. All error bars in this article convey the standard error of the mean

conditions as a mixture between these two modes of control. To gauge the relative contributions of proactive control and reactive control, we used maximum likelihood estimation to fit the parameters $p_{\text{intention}}$, and p_{recall} to the accuracies of individual participants. As Fig. 7b shows, this model achieves an impressively good fit with the most likely $p_{\text{intention}} = 0.58$. This suggests that, on average, people engage in proactive control in about 58% of the trials of this task. Consistent with the dual mechanisms of control framework (Braver, 2012), we found substantial interindividual differences in the propensity to engage proactive control. Concretely, individual participants' propensities to set intentions ranged from 0.0001% to 100% ($M = 0.58$, $SD = 0.30$) Participants' propensity to engage in reactive control appeared to be less variable. That is, people's propensity to recall the rules when the encountered the probe without a prepared intention ranged from 54% to 100% with a mean of 93% and a standard deviation of 9% ($M = 0.9281$, $SD = 0.09$).

Testing the model's predictions about when people set intentions

To test whether our model can capture when people set intentions, we simulated the experiments conducted by Mäki-Marttunen et al. (2019a), Redick (2014), and Gonthier et al. (2016). Each time, we compare how often people should set an intention or recall the rules according to our model to how often they actually do. To estimate how often people engage in these processes, we estimate the parameters $p_{\text{intention}}$ and

p_{recall} of a simple measurement model from people's responses. This measurement model describes participants' response probabilities, R , by

$$P_{\text{measurement}}(R|\text{cue}, \text{probe}, \text{load}, \theta; m) = \sum_{I, \text{recall} \in \{0, 1\}} P(I) \cdot P(\text{recall}) \cdot P(R|\text{cue}, \text{probe}, \theta, I, \text{recall}, \text{load}; m),$$

where $P(I = 1) = p_{\text{intention}}$, $P(\text{Recall} = 1) = p_{\text{recall}}$, and m is our model of how proactive and reactive control affect people's accuracies.

Prediction 1: Effect of incentives on proactivity

The reward condition of the experiment reported by Mäki-Marttunen et al. (2019a) incentivized fast accurate responses in a standard AX-CPT task with 70% AX trials, 10% AY trials, 10% BX trials, and 10% BY trials. Assuming that these incentives increase the subjective reward for identifying AX pairs (u_+) and responding faster (u_{Δ}) by one unit, Equation 2 of our model predicts that incentivized participants should be more likely to proactively set intentions than unincentivized participants on A-trials (Figs. 8a and 12) as well as on B-trials and become more proactive overall. Concretely, our model predict that incentives should increase the probability of intention setting on A trials from 20% to 22% and increase the probability of intention setting on B trials from 80% to 84%. To test our meta-control model's prediction about the probability of proactive intention-setting, we used the maximum likelihood estimation method described in Section 2 of the Supplementary Material to estimate the relative frequencies with which participants proactively set intentions ($p_{\text{intention}}$) separately to A-trials and B-trials, respectively. As shown in Fig. 8b, the maximum likelihood estimates showed that the incentives increased the probability of intention setting on A-trials from 20% ($SD = 0.06$) to 27% ($SD = 0.15$; $t(264) = 4.91$, $p < 0.001$). The incentives' effect on the probability of intention setting on B-trials was very close to the model's prediction (80% ($SD = 0.23$) vs. 83% ($SD = 0.22$)), but this effect was not statistically significant ($t(264) = 0.97$, $p = 0.33$).

Prediction 2: Effect of incentives on accuracy on AY trials

To predict the effect of incentivizing participants to rapidly generate correct responses, we first fitted the parameters of our meta-control model to individual participants' accuracies in the nonreward condition of the experiment by Mäki-Marttunen et al. (2019a) using the method described in Section 2 of the Supplementary Material and then used our model to simulate what their accuracies would have been if each participant's subjective utilities for responding fast (u_{Δ})

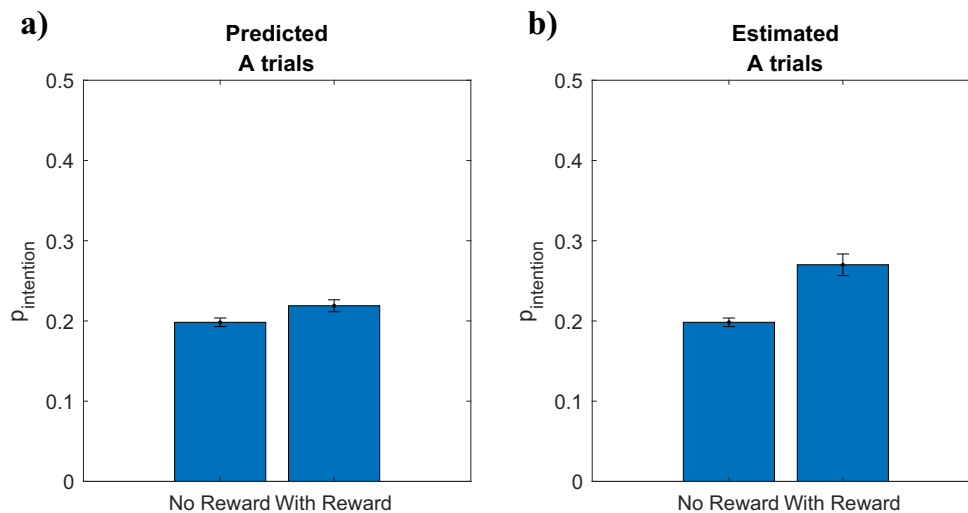


Fig. 8 Predicted and estimated effects of incentivizing fast, accurate responses on the probability of proactive intention setting

and correctly identifying AX pairs (u_+) had been one unit higher. As illustrated in Fig. 9a, our model predicts that an increased reward for accurately identifying AX pairs (u_+) and responding faster ($u_{\Delta t}$) should have a negative effect on people’s performance on AY trials (Fig. 9a). Consistently with this prediction, Mäki-Marttunen et al. (2019a) found that reward incentives decreased people’s performance on AY trials (78% vs. 92%, $t(264) = -6.43$, one-tailed $p < 0.0001$; Fig. 9b). As our model had predicted (see Fig. 9a), this inverted the relationship between people’s performance on AY trials versus AX trials. That is, people’s average performance in the experiment by Mäki-Marttunen et al. (2019a) was higher on AY trials than on BX trials in the no-reward condition (92% vs. 89%; $t(260) = 1.41$, $p = 0.0799$), but the opposite was the case in the reward condition (78% vs. 91%; $t(268) = -4.90$, $p < 0.0001$; Fig. 9b).

Prediction 3: Effect of the relative frequency of AX trials on proactive intention setting

As illustrated in Fig. 14a, our meta-control model predicts that people’s propensity to engage proactive control should decrease as the relative frequency with which an A is followed by an X drops to 50%. This is intuitive because the proactively set intention in the AX-CPT is correct only when the A is followed by an X. Because proactive control increases the frequency of errors on AY trials, decreasing the frequency of AX trials therefore should increase participants’ accuracy on AY trials, as illustrated in the top row of Fig. 10. As shown in the bottom row of Fig. 10, the experiments reported in Redick (2014) confirmed this prediction. Furthermore, people remain highly accurate on AY trials as their frequency exceeds the frequency of AX trials. According to our model, this is because people will then prepare to respond to Y when they see the A.

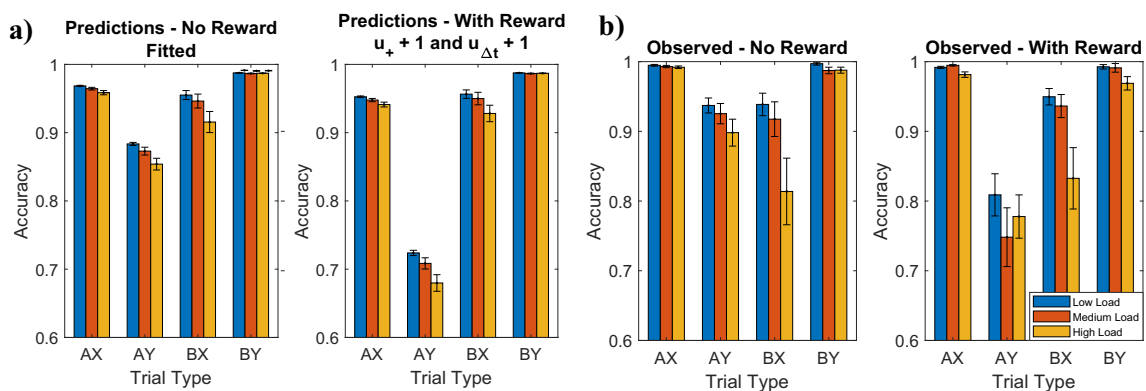


Fig. 9 Comparison of predicted versus observed effects of incentivizing fast, accurate responses in the behavioral experiment by Mäki-Marttunen et al. (2019a). **a** Prediction of our meta-control model. **b** Human performance

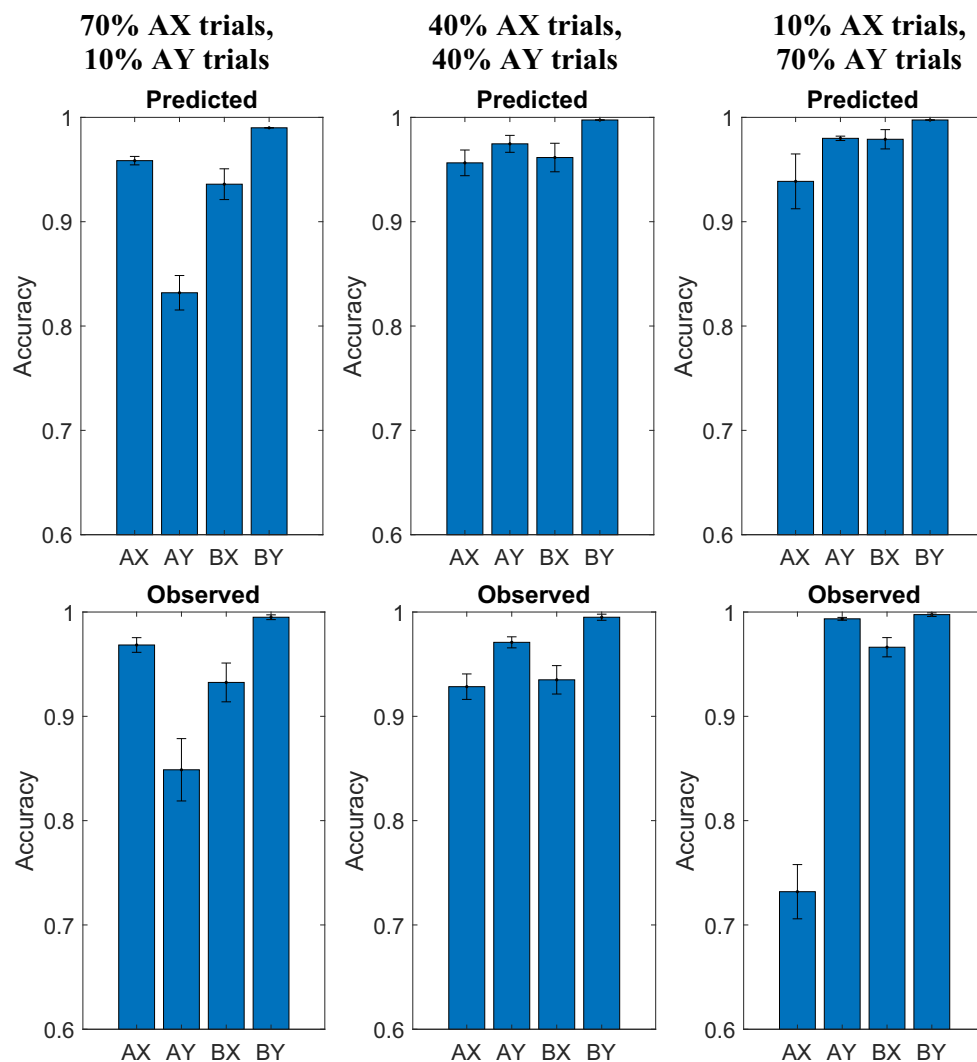


Fig. 10 Mean predicted (upper) and observed (lower) accuracies for Redick (2014) datasets

Prediction 4: Effect of no-go trials

To predict the effect of adding no-go trials on people's performance on AX, AY, BX, and BY trials, we fitted our model's parameters to data from a standard AX-CPT without no-go trials (i.e., the control condition of Experiment 1 from Gonthier et al., 2016) and then used the estimated parameters to simulate how the accuracies should change when no-go trials are added (i.e., Experiment 2 of Gonthier et al., 2016). If a participant set an intention on a no-go trial, then they either have to inhibit that intention in response to the no-go probe, which costs time and effort, or accept the penalty for giving an incorrect response. This is why the possibility of no-go trials reduces the expected benefit of intention setting and reduces its expected cost. Therefore, our rational model predicts that adding no-go trials should decrease proactive intention setting in response to both A cues and B cues (Fig. 11a). Figure 11b shows that this is indeed the case. Because of this effect our

model also predicts that no-go trials should increase people's accuracy on AY trials and decrease their accuracy on BX and BY trials (Fig. 11c). As Fig. 11d shows, Experiment 2 from Gonthier et al. (2016) confirmed this prediction.

Prediction 5: Intention setting on BX versus AY trials

According to our meta-control model, on BX trials the probability to set an intention should be higher and responses should be the faster than on AY trials according to our model. This is because B cues induce the highest probability of engaging proactive control and X probes boost proactive control, whereas A cues induce a lower probability of engaging proactive control and Y probes trigger an inhibition of the proactively prepared intention (Fig. 12).

Confirming our model's prediction, the reaction times reported by Mäki-Marttunen et al. (2019a) are significantly lower for BX trials than for AY trials across all six experimental

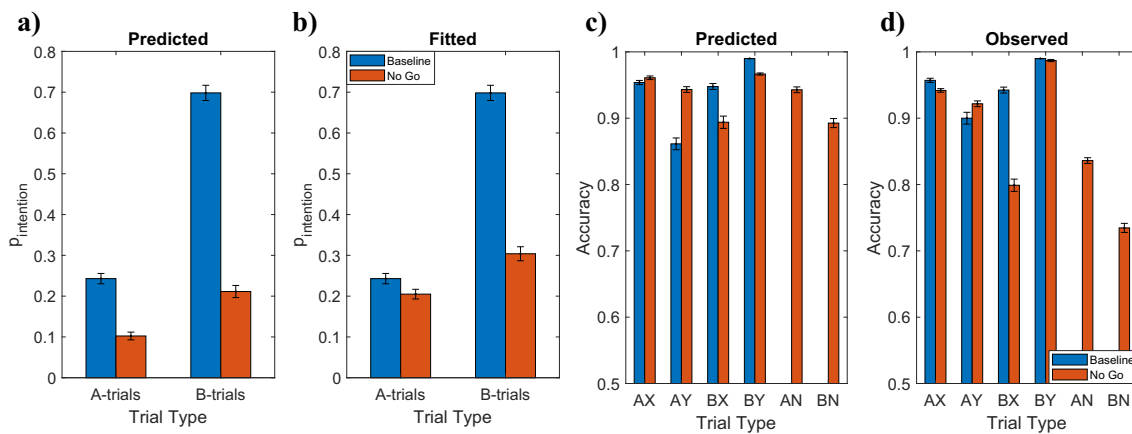


Fig. 11 **a** Predicted $P_{intention}$ for baseline and no-go conditions. **b** Fitted accuracies for baseline and no-go conditions. **c** Predicted accuracies for the conditions with versus without no-go trials using the maximum

likelihood estimates for the baseline condition and only changing the relative frequency of the trial types for no-go condition. **d** Observed accuracies in Experiment 2 by Gonthier et al. (2016).

conditions ($t(530) = -3.67, p < 0.001$), and the average probability of intention setting was significantly larger for BX trials ($M = 0.7846; SD = 0.2138$) than for AY trials ($M = 0.2630, SD = 0.1418; t(530) = 33.1610, p < 0.0001$).

Prediction 6: Effect of contextual load on accuracy

To predict the effect of contextual load on accuracy, we fitted our model’s parameters the individual participants’ response in the condition of the experiment by Mäki-Marttunen et al.

(2019a) that had a contextual load of 2 while constraining λ to be at least 0.025. We then simulated how the accuracies should change when the contextual load is increased to 3 or decreased to 1. We performed this procedure separately for the reward condition and the no reward condition. As shown in Fig. 13, our model predicted that cognitive load should reduce people accuracy on AX trials, BX trials, and BY trials, but not on AY trials.

Confirming this prediction, Mäki-Marttunen et al. (2019b) found that on AX, BX, and BY trials people’s accuracy was

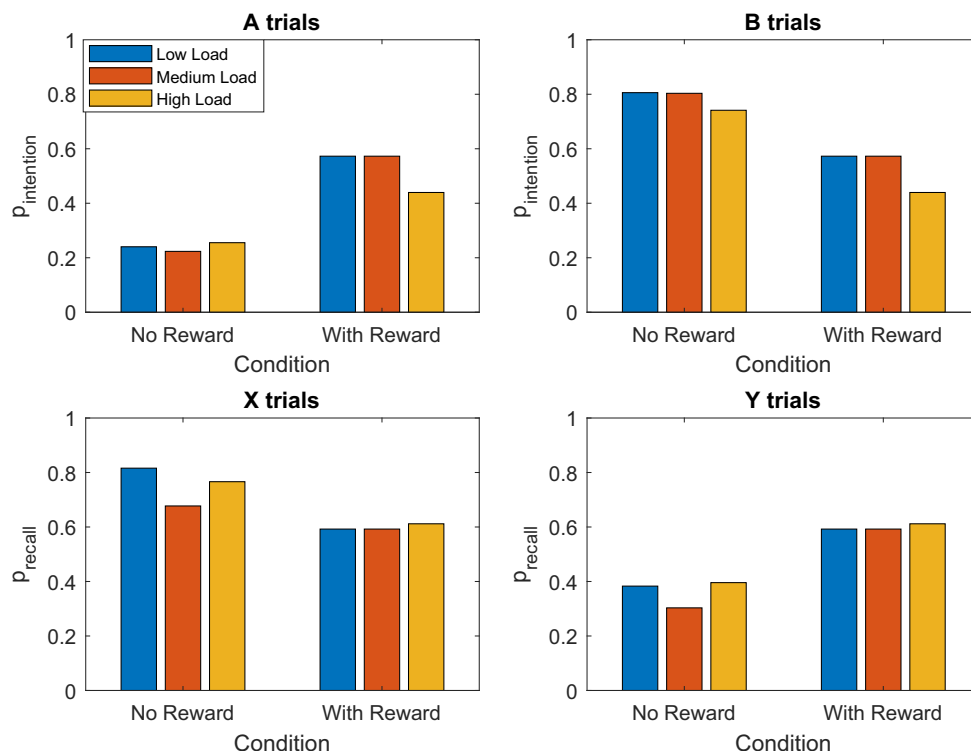


Fig. 12 Estimates of the probability to set an intention and the probability of recalling the rules by trial type and experimental condition. These estimates were obtained by fitting our meta-control process model to all

accuracies of individual participants using the method described in Section 2 of the Supplementary Material.

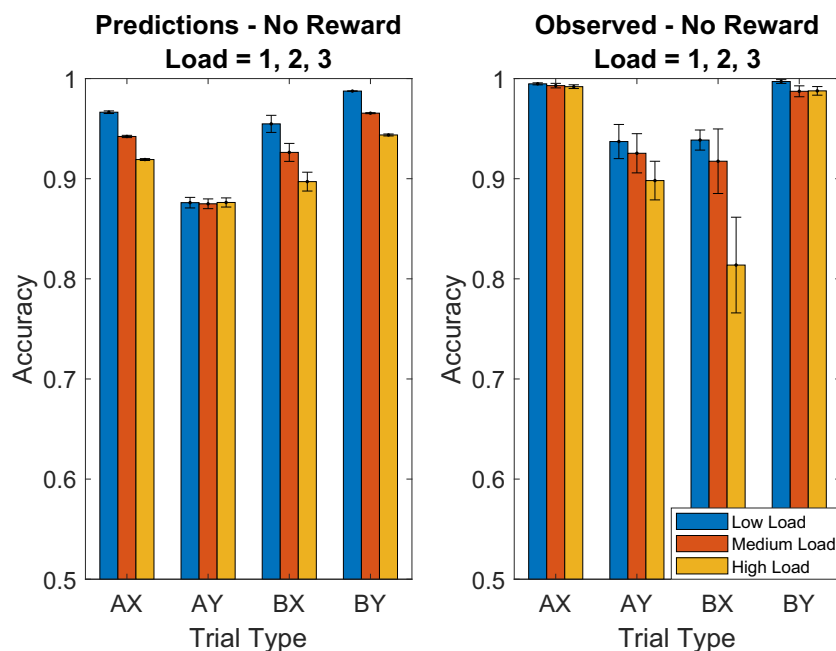


Fig. 13 Predicted effects of contextual load on accuracy

significantly lower for high cognitive load than for low cognitive load and Mäki-Marttunen et al. (2019a) found this effect for AX and BX trials when comparing the high load condition to either the low load condition or the intermediate load condition. We found that to obtain this prediction with our model, it is critical to assume that there is a direct interference of load on the accuracy of controlled processing ($\lambda > 0$). As our model predicted, no effect of load was found for AY trials. The model-based estimates shown in Fig. 12 suggest that the effect of cognitive load on accuracy might be mediated by a reduction in intention setting in the high load condition.

Discussion

In this article, we introduced and validated a computational level theory of the meta-control decisions that determine people's level of proactivity in the AX-CPT. In doing so, we have instantiated the dual mechanisms of control framework (Braver, 2012) in terms of a precise computational model of the meta-control decisions that drive variability in cognitive control. Our model predicts when people engage proactive control, reactive control, or both. Our model captures that proactivity involves setting intentions based on predictions about the future, allocating control to those intentions when opportunities arise, and inhibiting competing automatic responses. The basic idea of our model is that proactivity is governed by the allocation of control according to a rational cost-benefit analysis. Empirical data from previous experiments supported numerous predictions of our model, including its predictions about the effects of incentives, contextual

load, adding no-go trials, and changing the ratio of AX trials to AY trials. This suggests that our model is a promising step towards unraveling the computational mechanisms of proactivity. Understanding proactivity, in turn, is an important step towards understanding what it takes to live a successful life and how we can assist people in this challenging process (Lieder & Prentice, 2020).

Implications for our theoretical understanding of proactivity in the AX-CPT

Our model formalizes central ideas of the Dual Mechanisms of Control framework (Braver, 2012) in terms of rational tradeoffs between the costs of exerting cognitive control and its benefits (Lieder & Griffiths, 2020; Shenhav et al., 2017). Its success in predicting the effects of adding incentives and increasing cognitive load (Mäki-Marttunen et al., 2019b), reducing the frequency of AX-trials (Redick, 2014), adding no-go trials (Gonthier et al., 2016), therefore lends some support to viability of those theoretical assumptions. Most strikingly, our model's generalization at predicting the effects of changing the statistical structure of the task by reducing the frequency of AX trials (Redick, 2014) or adding no-go trials (Gonthier et al., 2016) supports its central assumption that people engage proactive control according to a rational cost-benefit analysis. This lends further support to the Dual Mechanisms of Control framework (Braver, 2012), the expected value of control theory (Shenhav et al., 2013), and the theory of resource-rationality (Lieder & Griffiths, 2020).

Our model comparisons strongly supported a meta control model in which proactive and reactive control are independent

and can both occur on the same trial. This version of our model makes exactly two independent meta-control decisions: one about whether or not to engage proactive control and a second one about whether or not to engage reactive control. Supporting the view that proactive control and reactive control are independent, the meta-control decision about reactive control is always made regardless of whether an intention was previously set. It decides whether to recall the rules based on a cost-benefit analysis that only considers the probe. If the rules are recalled then whatever mechanisms would have determined the person's choice otherwise, be it a proactively set intention or automaticity, will be inhibited and overridden by reactive control. Supporting the view that the inhibition of prepotent intentions is an important part of reactive control, the data favored the recall-override model over an alternative simplified model that does not allow for the inhibition of proactively set intentions. In addition, the simulation results summarized in Fig. 5 of the [Supplementary Material](#) show that proactive control and reactive control are two complementary mechanisms of proactivity that are both needed to capture people's capacity for proactive goal-directed behavior.

Future Directions

Our model allows us to derive a number of predictions that go beyond the phenomena studied by Mäki-Marttunen et al. (2019a), Mäki-Marttunen et al. (2019b), Gonthier et al. (2016), and Redick (2014). These predictions can guide the design of future experiments. In general, our rational model of meta-control predicts that proactivity increases with situational factors and personal characteristics that make proactivity more beneficial and decreases with situational factors and personal characteristics that make it costlier and less beneficial. The value of proactivity increases with the predictability of the environment. Therefore, our model predicts that for AX frequencies in-between the extremes used by Redick (2014) people should gradually become less proactive as more AX trials are converted into AY trials. Conversely, and this has not been explored yet, people should become more proactive when the frequency of AX trials is increased beyond 70% or the frequency of AY trials is decreased below 10% (Fig. 14a). Based on this prediction, proactive intention setting should be most frequent in response to a cue that is always followed by the same probe and inverts the required response to that probe. Proactivity also becomes more valuable as reflexive responding to the probe becomes less effective. Reflexive responding to the X probe is least effective when both responses are correct equally often. This is the case when the frequency of BX trials equals the frequency of AX trials. Our model therefore predicts an inverse-U shaped effect of the proportion of B-cues on proactivity when the ratio of AX trials to AY trials and the ratio of BX trials to BY trials is held constant (Fig. 14b). Furthermore, when the incentives or

instructions emphasize accuracy over speed, then the reward manipulation would have the opposite effect on people's performance on AY trials than the one observed by Mäki-Marttunen et al. (2019a); that is, as the emphasis of the reward criterion shifts from speed towards accuracy, the effect of reward on people's performance on AY trials should become increasingly positive (Fig. 14c), in contrast to the experiment by (Mäki-Marttunen et al., 2019a) where the effect was negative. The interested reader can download the code of our model and simulations are available on the Open Science Framework (<https://osf.io/ng65r>) to generate these and other predictions or to fit our model to their own data sets by following the instructions in the readme file.

The model introduced above is a computational-level theory. It defines the function of meta-control over proactivity. As proposed by Lieder et al. (2018) and confirmed by Bustamante et al. (2021), we postulate that the meta-control system *learns* to approximate the optimal solution proposed by our model. Applied to proactivity in the AX-CPT, the Learned Value of Control (LVOC) model predicts that if we were to create different versions of the AX-CPT task that require more versus less proactivity, then we should see people's proactivity gradually increase versus decrease over time.

In addition to testing these predictions, future work will extend our investigation of proactivity from proactive control in the AX-CPT task to proactivity in the real world. This will include investigating whether our formal mathematical measure of goal-directedness (see Section 4 of the [Supplementary Material](#)) based on people's performance in simple laboratory paradigms is predictive of proactivity in the real world. In a parallel line of work, we will extend our model by the aspects of proactivity that are currently missing from it. Referring back to the definitions of proactivity that we started from (Crant, 2000; Parker et al., 2010; Seibert et al., 1999; Siebert & Kunz, 2016), we can see that our model captures that proactivity includes the self-initiated and future-oriented setting of intentions and the active pursuit of those intentions over time. However, the important cognitive processes of deriving long-term goals from personal values, planning, motivation, self-improvement, progress monitoring, self-regulation, and reflection are still missing from our model. Incorporating these additional mechanisms will likely allow our model to capture even higher levels of proactivity that people are likely to exhibit in more naturalistic scenarios. As an intermediate step, we will model goal setting and extend our model of proactivity to other experimental paradigms that have been used to study proactive control, including working memory paradigms (Braver, 2012; Burgess et al., 2011).

Furthermore, future work will supplement our model from a computational-level theory of the function of proactivity with mechanistic models of the underlying meta-decision-making processes and investigate how those processes are shaped by learning. Understanding how proactivity is shaped

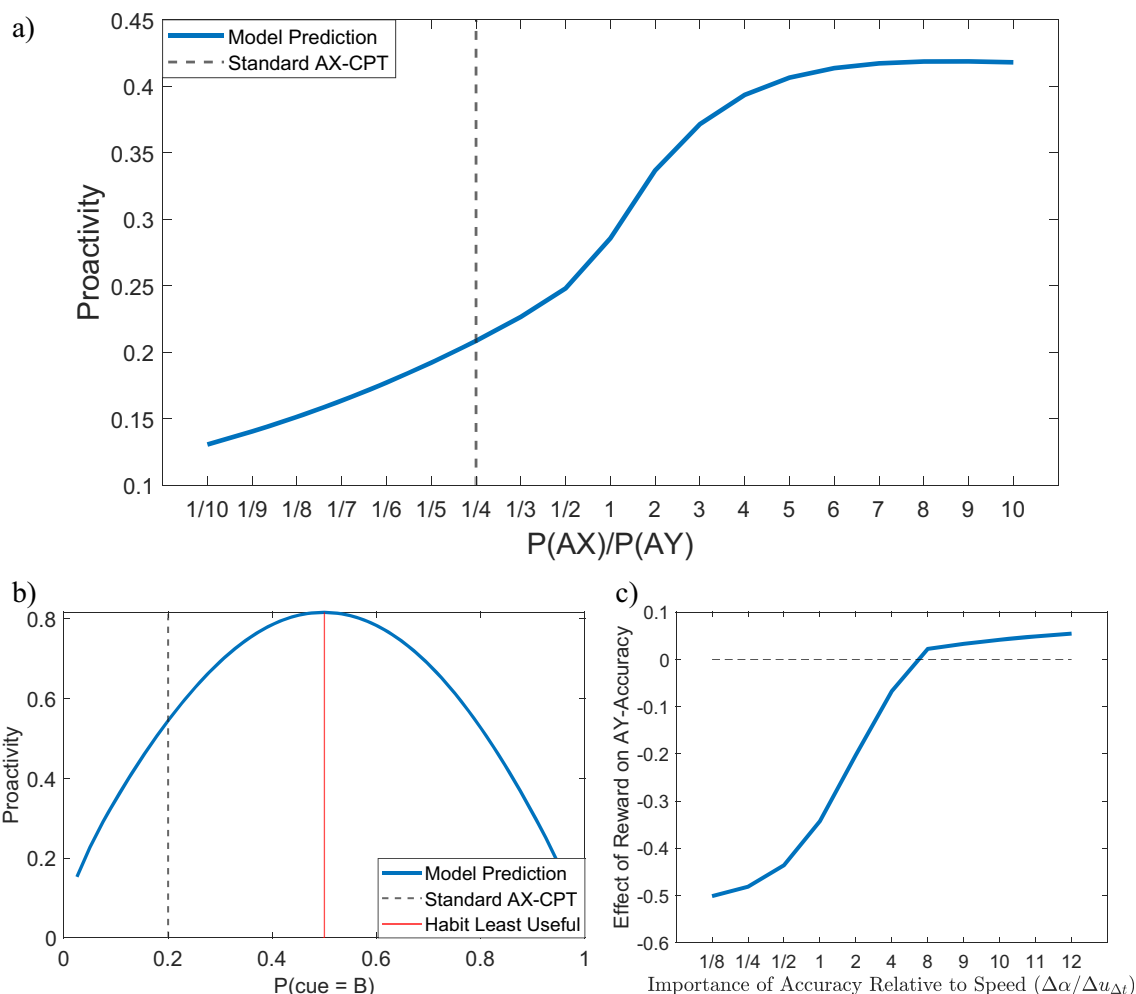


Fig. 14 Model predictions to be tested in future work. The vertical dashed lines in **a** and **b** mark the values corresponding to a standard AX-CPT with 70% AX trial, 10% AY trials, 10% BX trials, and 10% BY trials. **a** Our model predicts an increase of proactivity with the frequency of AX trials relative to the frequency of AY trials. In this simulation there are 10% BX trials, 10% BY trials, and AX and AY trials jointly make up the remaining 80% of the trials. A ratio of 1 means that there are 40% AX trials and 40% AY trials. A ratio of 20 means that there are 3.5% AY trials and 76.5% AX trials. The model parameters are those used to simulate the reward condition. Proactivity is measured according to Equation 13 in the [Supplementary](#)

Material. b Our model predicts an inverse-U-shape relationship between the proportion of B-cues (p_B) when the relative frequencies of AX, AY, BX, and BY trials are $\frac{7}{8} \cdot (1-p_B)$, $\frac{1}{8} \cdot (1-p_B)$, $\frac{7}{8} \cdot p_B$, and $\frac{1}{8} \cdot p_B$ respectively. The model predictions suggest that people should be most proactive when habits are least useful. **c** For a sufficiently high reward (4 times the subjective value of getting it correct without any reward) our model predicts that the effect of reward on people's performance on AY trials should switch from decreasing people's performance to increasing people's performance as the emphasis of the reward criterion shifts from speed to accuracy

by learning will be an important step towards developing training interventions for helping people become more proactive.

Relevance to cognitive neuroscience

Proactive control is an important topic in cognitive neuroscience and investigating its neurocomputational mechanisms will further our understanding of the function of prefrontal cortex and the neuromodulatory systems underlying cognitive control (Braver, 2012; Mäki-Marttunen et al., 2019b). We anticipate that our model will become an important asset in the search for the neural mechanisms of proactivity, akin to how the

expected value of control model (Shenhav et al., 2013) has helped us elucidate the neural underpinnings of cognitive control and mental effort (Shenhav et al., 2017).

While previous research has demonstrated that the dorsolateral prefrontal cortex plays an important role in the implementation of proactive control (Mäki-Marttunen et al., 2019b), the neural basis of meta-control over proactivity is less well understood. Based on previous work, we postulate that 1) the cost-benefit analyses that govern the meta-control over proactive control are implemented in the dorsal anterior cingulate cortex (dACC, Shenhav et al., 2013; Shenhav et al., 2017), and 2) that the meta-control decisions made by the dACC are then implemented by the dorsolateral prefrontal cortex (dlPFC, Badre,

2008; Mäki-Marttunen et al., 2019b). Future work might employ neuroimaging methods and pupillometry to test our models and our hypothesis about its neural substrates in the three following ways. First, future work might use fMRI to evaluate how well the event-related activity of the dACC is predicted by the cost and benefit terms postulated by our model. Second, fMRI or fNIRS could be used to investigate how well the control signals selected by our model predict the task-dependent activation of the dlPFC (Mäki-Marttunen et al., 2019b) and the locus coeruleus (Mäki-Marttunen et al., 2019b). Pupillometry could be used to test our model's predictions about the working memory load and mental effort entailed by the control signals and strategies chosen by our model (Kahneman & Beatty, 1966). Third, EEG or MEG could be used to test our models' assumptions about the number and nature of meta-control decisions are involved in people's responses on AY trials. This would allow for a more direct comparison of our basic meta-control model (Fig. 2) against the extended meta-control model (Fig. 6a) and simpler meta-control models (Fig. 6c). If the neural correlates of these meta-control decisions can be identified then future work might also leverage them to test our model's predictions about how those three stages are affected by the statistical structure of the task, the incentives, and cognitive load.

Conclusions

We believe that our computational level theory of variation in pro- and reactive control is an important step towards a formal theory of proactivity. We hope that by connecting the concept of proactivity from the management and personal development literatures to laboratory paradigms and computational models of meta-control, our article will help to make the computational challenges of living a good life amenable to rigorous scientific investigation in the laboratory.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.3758/s13415-021-00884-y>.

Acknowledgments The authors thank Veronica Mäki-Marttunen for sharing the data from Mäki-Marttunen et al. (2019a), Thomas S. Redick for sharing the data from Redick (2014), Corentin Gonthier for sharing the data from Gonthier et al. (2016), Sebastian Musslick for pointing us to the literature on proactive control, Todd Braver for his feedback and advice, and Emily Corwin-Renner for her feedback.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were

made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences*, 12(5), 193–200
- Boureau, Y.-L., Sokol-Hessner, P., & Daw, N. D. (2015). Deciding how to decide: Self-control and meta-decision making. *Trends in Cognitive Sciences*, 19(11), 700–710
- Braver, T. S. (2012). The variable nature of cognitive control: a dual mechanisms framework. *Trends in Cognitive Sciences*, 16(2), 106–113
- Burgess, G. C., Gray, J. R., Conway, A. R., & Braver, T. S. (2011). Neural mechanisms of interference control underlie the relationship between fluid intelligence and working memory span. *Journal of Experimental Psychology: General*, 140(4), 674
- Bustamante, L. A., Lieder, F., Musslick, S., Shenhav, A., & Cohen, J. D. (2021). Learning to Overexert Cognitive Control in a Stroop Task. *Cognitive, Affective, & Behavioral Neuroscience*.
- Crant, J. M. (2000). Proactive behavior in organizations. *Journal of Management*, 26(3), 435–462
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12), 1704
- De Pisapia, N., & Braver, T. S. (2006). A model of dual control mechanisms through anterior cingulate and prefrontal cortex interactions. *Neurocomputing*, 69(10–12), 1322–1326
- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2), 312–325
- Gonthier, C., Macnamara, B. N., Chow, M., Conway, A. R. A., & Braver, T. S. (2016). Inducing Proactive Control Shifts in the AX-CPT. *Frontiers in Psychology*, 7(1822). <https://doi.org/10.3389/fpsyg.2016.01822>.
- Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M., & Lieder, F. (2019). Doing more with less: meta-reasoning and meta-learning in humans and machines. *Current Opinion in Behavioral Sciences*, 29, 24–30
- Irlbacher, K., Kraft, A., Kehrer, S., & Brandt, S. A. (2014). Mechanisms and neuronal networks involved in reactive and proactive cognitive control of interference in working memory. *Neuroscience & Biobehavioral Reviews*, 46, 58–70
- Judge, T. A., & Kammeyer-Mueller, J. D. (2007). Personality and career success. *Handbook of Career Studies*, 59–78
- Kahneman, D., & Beatty, J. (1966). Pupil diameter and load on memory. *Science*, 154(3756), 1583–1585
- Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS computational biology*, 7(5), e1002055
- Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, 124(6), 762
- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43
- Lieder, F., & Prentice, M. (2020). Life improvement science: a manifesto. *Technical Report*.

- Lieder, F., Shenhav, A., Musslick, S., & Griffiths, T. L. (2018). Rational metareasoning and the plasticity of cognitive control. *PLoS computational biology*, 14(4), e1006043
- Mäki-Marttunen, V., Hagen, T., Aminihajibashi, S., Foldal, M., Stavrinou, M., Halvorsen, J. H., ... Espeseth, T. (2018). Ocular signatures of proactive versus reactive cognitive control in young adults. *Cognitive, Affective, & Behavioral Neuroscience*, 18(5), 1049-1063
- Mäki-Marttunen, V., Hagen, T., & Espeseth, T. (2019a). Proactive and reactive modes of cognitive control can operate independently and simultaneously. *Acta Psychologica*, 199, 102891
- Mäki-Marttunen, V., Hagen, T., & Espeseth, T. (2019b). Task context load induces reactive cognitive control: An fMRI study on cortical and brain stem activity. *Cognitive, Affective, & Behavioral Neuroscience*, 19(4), 945-965. <https://doi.org/10.3758/s13415-019-00691-6>.
- Parker, S. K., Bindl, U. K., & Strauss, K. (2010). Making things happen: A model of proactive motivation. *Journal of Management*, 36(4), 827-856
- Parker, S. K., Wang, Y., & Liao, J. (2019). When is proactivity wise? A review of factors that influence the individual outcomes of proactive behavior. *Annual Review of Organizational Psychology and Organizational Behavior*, 6, 221-248
- Redick, T. S. (2014). Cognitive control in context: Working memory capacity and proactive control. *Acta Psychologica*, 145, 1-9
- Reynolds, J. R., Braver, T. S., Brown, J. W., & Van der Stigchel, S. (2006). Computational and neural mechanisms of task switching. *Neurocomputing*, 69(10-12), 1332-1336
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593-1599
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461-464
- Seibert, S. E., Crant, J. M., & Kraimer, M. L. (1999). Proactive personality and career success. *Journal of Applied Psychology*, 84(3), 416
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron*, 79(2), 217-240
- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a rational and mechanistic account of mental effort. *Annual Review of Neuroscience*, 40, 99-124.
- Siebert, J., & Kunz, R. (2016). Developing and validating the multidimensional proactive decision-making scale. *European Journal of Operational Research*, 249(3), 864-877
- Tornau, K., & Frese, M. (2013). Construct clean-up in proactivity research: A meta-analysis on the nomological net of work-related proactivity concepts and their incremental validities. *Applied Psychology*, 62(1), 44-96
- van der Meer, M., Kurth-Nelson, Z., & Redish, A. D. (2012). Information processing in decision-making systems. *The Neuroscientist*, 18(4), 342-359
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38(4), 599-637

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.