Cognition 274 (2026) 106537

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Full length article

Aha! moments correspond to metacognitive prediction errors

Rachit Dubey [a,*], Mark Ho [b], Hermish Mehta [c], Thomas L. Griffiths [d]

[a] Department of Communication, University of California, Los Angeles, United States of America
[b] Department of Psychology, New York University, United States of America
[c] Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, United States of America
[d] Departments of Psychology and Computer Science, Princeton University, United States of America

ARTICLE INFO

Keywords:
Insight
Aha! moment
Problem solving
Metacognition
Prediction errors
Reinforcement learning
Monitoring and control

ABSTRACT

Psychologists have long been fascinated with understanding the nature of *Aha!* moments, moments when we transition from not knowing to suddenly realizing the solution to a problem. In this work, we present a theoretical framework that explains why we experience *Aha!* moments. Our theory posits that during problem-solving, in addition to solving the problem, people also maintain a metacognitive model of their ability to solve the problem as well as a prediction about the time it would take them to solve that problem. *Aha!* moments arise when we experience a positive error in this metacognitive prediction, i.e. when we solve a problem much faster than we expected, signaling to ourselves that we are more competent than we had realized. We posit that this metacognitive error is analogous to a positive reward prediction error thereby explaining why we feel so good after an *Aha!* moment. We provide support to our theory across two large-scale pre-registered experiments on problem solving, demonstrating a link between metacognitive prediction errors and *Aha!* moments. These results highlight the importance of metacognitive prediction errors and deepen our understanding of human metareasoning.

1. Introduction

The word “inspiration” derives from the Latin *inspiro*, meaning to excite, inflame or inspire: conjuring vivid images of artists receiving divine guidance to produce sudden and unexpected insight. While these moments may be less dramatic in real life they are nevertheless familiar, often in the form of a sudden insight — colloquially called the “*Aha!* moment” — experienced after we suddenly solve a challenging puzzle or problem. Psychology has been fascinated with understanding the nature of these *Aha!* moments for nearly a century (Maier, 1931; Mayer, 1992), in large part because of its potential to deepen our understanding of human intelligence, problem solving, and creativity (Chuderski & Jastrzębski, 2018; Cushen & Wiley, 2011; Jones, 2012; Menkel-Meadow, 2001). An impressive experimental literature has studied the phenomenology, determinants, and mechanisms of *Aha!* moments and has documented several findings about *Aha!* moments. For one, *Aha!* moments are accompanied by strong positive emotions and a feeling of joy (Danek et al., 2014; Shen et al., 2016), and are often sudden and unpredictable (Danek et al., 2018; Metcalfe, 1986b). They are also more pronounced after an impasse (Friedlander & Fine, 2018; Öllinger et al., 2008) and are influenced by subjective expectations (Danek et al., 2016). However, there remain many open questions about why *Aha!* moments have these characteristic signatures. People solve a wide variety of problems in different contexts, but only some of these tasks evoke the distinctive experience of an *Aha!* experience. For example, the experience of solving an anagram (a classic ‘insight’ problem) is profoundly different from making a cup of coffee. What features of a problem, problem-solver, and their interaction lead to *Aha!* moments? A challenge for theories of cognition is to explain *why* we experience *Aha!* moments in the first place and why these moments are so rewarding to us.

In this article, we present a computational account that provides an explanation for why we experience *Aha!* moments. Our theory connects the experience of an *Aha!* with the literature on reward learning (Sutton & Barto, 2018) and suggests that *Aha!* moments correspond to *metacognitive prediction errors* resulting from inferences about one’s own problem-solving. Specifically, we propose that while solving a problem, people maintain a metacognitive model of their ability to solve the problem as well as how long it will them take to complete it. Observations of the progress towards completing the problem then drive probabilistic updates to this metacognitive model and its predictions. Crucially, *Aha!* moments arise as a consequence of maintaining such metacognitive predictions — finishing a task faster than expected acts as a kind of “metacognitive prediction error” analogous to *reward*

* Corresponding author.
E-mail address: rdubey@ucla.edu (R. Dubey).

https://doi.org/10.1016/j.cognition.2026.106537
Received 8 September 2025; Received in revised form 24 March 2026; Accepted 25 March 2026
Available online 21 April 2026

prediction errors. An extensive literature on reinforcement learning has shown that such reward prediction errors are associated with strong positive affect (Bayer & Glimcher, 2005; Glimcher, 2011; Rutledge et al., 2014; Schultz, 2016; Schultz et al., 1997), thereby explaining why we feel so good once we experience an *Aha!* moment.

At the outset, we emphasize that our goal is not to model *how* particular problem-solving computations lead to insight, but rather, to explain *why* this sudden insight causes the positive experience of *Aha!*. This helps explain some (though not all) important findings about *Aha!* moments and clarifies the important role of metacognition in influencing this phenomenon, adding to the rich literature on computational modeling of the process of insight (Hélie & Sun, 2010; Langley & Jones, 1988; Thagard & Stewart, 2011) as well as research that relates metacognition to insight (Ishikawa et al., 2019; Metcalfe, 1986b; Metcalfe & Wiebe, 1987). We further note that our theory doesn't seek to explain all possible kinds of *Aha!* that people experience. Our aim is to develop a minimal model of metacognitive appraisal that can apply across a variety of scenarios that involve metacognitive prediction errors, one of which is the experience of an *Aha!*.

Finally, and perhaps most importantly, we note that our theory doesn't intend to claim that metacognitive errors are the *only* cause of *Aha!* moments. As noted in prior works, the affective experience of *Aha!* moments can be driven by factors such as novelty, (Caprioli et al., 2023), fluency (Moroshkina et al., 2022; Skaar & Reber, 2020), or conceptual change (Chesebrough et al., 2023). Our goal is to highlight the important and overlooked role of metacognitive prediction errors in driving *Aha!* moments. Identifying the important role of these errors is especially promising because it offers a way to connect the phenomena of the *Aha!* experience to the more general framework of error-driven learning, including reinforcement learning (Sutton & Barto, 2018) as well as to several strands of research within metacognition including, curiosity (Litman, 2005; Loewenstein, 1994; Marvin & Shohamy, 2016; Modirshanechi et al., 2022; Wade & Kidd, 2019), metacognitive monitoring (Ackerman & Thompson, 2017; Fleming & Daw, 2017; Fox & Riconscente, 2008; Metcalfe & Greene, 2007), and the rational use of limited cognitive resources (Griffiths et al., 2015; Lieder & Griffiths, 2020).

The plan of the paper is as follows. We first briefly review theoretical approaches, experimental paradigms, and empirical findings on *Aha!* moments related to our proposal. Then, we detail our computational account of *Aha!* moments and present simulations which demonstrate that the dynamics of metacognitive inferences can explain several findings about *Aha!* moments. Next, we present two pre-registered behavioral experiments evaluating whether *Aha!* moments are indeed a form of metacognitive prediction error. We close with a discussion of the significance and implications of our results.

2. Background

From the time that psychologists began to investigate the domain of problem solving, the phenomenon of insight has been of great interest to the field. Insight is supposed to occur when one solves a problem or discovers a solution by a sudden breakthrough (Kaplan & Simon, 1990). A nearly ubiquitous observation is that insight is often accompanied by an affective response in the form of the *Aha!* experience, which distinguishes insightful problem solving from gradual problem solving (Chesebrough et al., 2023; Danek et al., 2014; Hedne et al., 2016; Webb et al., 2018). Our goal here is to understand this affective response that follows an insight. In this section, we provide a brief overview of the literature on *Aha!* moments as it relates to our account.

2.1. Theoretical models of *Aha!* moments

Since the Gestalt psychologists (Duncker & Lees, 1945; Kohler, 2015) introduced insight as a component process of perception and problem solving, numerous theories have been proposed as explanations. As our focus is on the function of *Aha!* moments in the larger context of problem solving computations, we mainly review accounts that address cognitive mechanisms, although there is a large body of work studying the relationship between insight and other processes (e.g., sleep and consolidation Wagner et al., 2004) as well as its neurobiological underpinnings (Aziz-Zadeh et al., 2009; Becker & Cabeza, 2025; Becker et al., 2025; Bowden & Jung-Beeman, 2007; Kounios & Beeman, 2014; Kounios et al., 2006).

Early work on *Aha!* moments focused on how changes in task representations reliably induce feelings of insight. One such account is *representational change theory* (Knoblich et al., 1999; Ohlsson, 1984), which aims to explain the transition from impasse to insight in terms of changes to one's representation of a problem such that a solution is suddenly and immediately identifiable. Two forms of representational change that have been the subject of systematic study are *constraint relaxation*, in which the solution space is expanded to include previously excluded options, and *chunk decomposition*, in which objects in a task are individuated in a new manner that reveals more solutions (Knoblich et al., 1999). In its original formulation, representational change theory alludes to formal theories of problem solving (e.g., heuristic search Simon & Newell, 1971) but was not itself formalized. Nonetheless, there have been several attempts to formalize the idea of representational change in problem solving. For example, *criterion for satisfactory progress theory* (also known as *progress monitoring theory*) (Chu et al., 2007; MacGregor et al., 2001), builds on representational change theory and proposes that people engage in a process of search over actions within a problem representation until they are unable to satisfactorily solve a problem, at which point they begin to relax constraints on problem solutions and restart the search process.

Theories based on representational change emphasize the role that different algorithmic processes have in leading to the experience of insight (e.g., searching actions versus relaxing constraints). However, they do not explain *why* those particular algorithmic processes lead to the positive phenomenology of *Aha!* experiences. To our knowledge, only some theories have offered an explanation for the phenomenology of insight. One theory that addresses this issue is the *processing fluency account* proposed by Topolinski and Reber (2010). This account explains the positive affect associated with an *Aha!* in terms of increased fluency: Upon solving a problem, it suddenly becomes easier to understand and process, which is intrinsically rewarding. A different approach based on neural network modeling is the *combination theory of creativity* (Thagard & Stewart, 2011). This account proposes that whenever neural patterns are combined into novel combinations that have a strong relevance towards the goals of the thinker, then that combination elicits a strong positive emotion in the form of an *Aha!* experience. While these theories shed light on some aspects of *Aha!* moments, there are several things they do not fully explain. For example, neither theory can explain why an impasse results in a stronger *Aha!* (Danek et al., 2014; Friedlander & Fine, 2018; Webb et al., 2016). Further, since the production of novel and useful combinations doesn't necessarily depend on the suddenness of the combination, the account of Thagard and Stewart (2011) also cannot explain why the *Aha!* experience is sudden and unpredictable (Metcalfe, 1986b).

2.2. Experimental paradigms for studying Aha! moments

Research on *Aha!* moments has been significantly shaped by the experimental paradigms used to study them, and these have evolved over time. Indeed, as we noted, early investigations into *Aha!* moments focused on representational change. This focus was in many ways motivated by classic findings with so-called "insight problems", problems that generally require a cognitive restructuring or sudden change in the way a problem is perceived to solve (Sternberg & Davidson, 1995). For example, in the nine-dot problem, participants are shown a 3×3

array of nine dots that need to be connected, without lifting pen, using only four lines. Solving the problem requires realizing that one can draw with the pen outside of the 3 × 3 array, but participants are often unable to realize this without prompting by an experimenter. Indeed, a notable feature of classic insight problems is that they are designed to invite in ineffective encoding of problems, which makes it unclear whether feelings of insight can be attributed to seeing through deception (Metcalfe, 1986b; Sternberg & Davidson, 1995). This has motivated alternative complementary paradigms for studying *Aha!* moments such as anagrams (Ammalainen & Moroshkina, 2020; Metcalfe, 1986a; Webb et al., 2018), compound-remote association tasks (Bowden & Jung-Beeman, 2003), and magic tricks (Danek et al., 2014), which do not have the aforementioned issues related to presentation (see also Webb et al., 2016 for a review). Additionally, tasks such as anagrams can be easily administered in a trial-based fashion, allowing for more efficient collection of larger datasets.

Problems that elicit feelings of insight are often contrasted with incremental or analytic problems as controls (DeCaro et al., 2016). Analytic problems characteristically involve finding a solution through a series of predictable steps, such as reasoning deductively about a series of constraints in a situation. However, the distinction between analytic and insight problems is not entirely clear cut for several reasons. First, the analytic/insight distinction is not an intrinsic property of a problem, but rather a property of the problem solver's current solution strategy relative to the problem — e.g., once a participant knows the right representation for an insight problem, they can solve it analytically. Second, the analytic/insight distinction, when formulated in terms of processes of search versus representational change, arguably presupposes the underlying cognitive processes it attempts to explain (i.e., those proposed by representational change theory). This motivates a more theory-neutral, operational definition of insight in terms of metacognitive reports that can be elicited from participants as they are engaged in problem solving. For example, in a series of studies, Metcalfe and colleagues (Metcalfe, 1986a, 1986b; Metcalfe & Wiebe, 1987) asked participants to provide *feelings of warmth* judgments indicating how far from solving a problem they felt they were. Problems standardly classified as insight problems tended to involve low feelings of warmth up until solving the problem, whereas analytic problems (specifically, algebra problems) showed a steadier increase in feelings of warmth (Metcalfe & Wiebe, 1987).

2.3. Empirical findings about Aha! moments

Despite the multiple theoretical and experimental approaches, there is agreement on several key findings about *Aha!* moments. We now briefly review some of these findings which our theory seeks to explain.

1. **Finding 1 — *Aha!* moments are accompanied by strong positive emotions.** One of the most striking findings about *Aha!* moments is that people experience strong positive emotions, in the form of pleasure or joy, once they experience an *Aha!* (Danek et al., 2014; Gick & Lockhart, 1995; Gruber, 1995; Ishikawa et al., 2019; Shen et al., 2018, 2016). This positive emotion accompanying *Aha!* has been observed consistently inside and outside the laboratory. As one example, Danek et al. (2014), used magic tricks to elicit the experience of an *Aha!* in participants and found that positive emotions were the most prevailing aspect of the experience. In naturalistic settings, it has likewise been shown that the *Aha!* experience produces an overwhelmingly positive affective response in undergraduate mathematics students and can even help in changing the attitude of 'resistant' students towards mathematics (Liljedahl, 2005).
2. **Finding 2 — *Aha!* moments are sudden and unpredictable..** Another finding about *Aha!* moments is that the experience of an *Aha!* is often unanticipated by a problem solver and they occur when a solution is reached suddenly and unexpectedly by the learner (Metcalfe, 1986b). At the same time, when a problem is solved gradually and sequentially, then problem solvers don't experience an *Aha!* (Danek et al., 2018; Metcalfe, 1986b). As demonstrated by Metcalfe (1986), an important characteristic of non-insight problems is that problem solvers are able to accurately judge their progress towards reaching the solution whereas they are unable to do so for problems that require an insight, leading to an *Aha!* when the unanticipated insight is eventually reached.
3. **Finding 3 — *Aha!* moments are more pronounced after an impasse.** Various studies have shown that the onset of an *Aha!* often exhibits a characteristic temporal pattern: An initial period of problem solving is followed by the experience of a phase in which the problem solver feels stuck and being in a state of deadlock where they do not know how to solve the problem anymore. Upon continued exploration of the problem, eventually a new idea comes to mind which leads to a rapid completion of the problem leading to the experience of an *Aha!* (Knoblich et al., 2001; Öllinger et al., 2008, 2014). This impasse-insight sequence is another important finding about the *Aha!* experience — the emotional response following an insight is often greater if the problem solver experienced a period of impasse before the eventual resolution of the problem is reached (Danek et al., 2014; Friedlander & Fine, 2018; Webb et al., 2016).
4. **Finding 4 — *Aha!* moments are subjective and can occur in a variety of problems.** *Aha!* moments are not just a property of a particular question (Fleck & Weisberg, 2013; Webb et al., 2016). That is, the same problem can elicit different levels of an *Aha!* in different people, with one study showing that in many cases participants can even solve various classical insight problems without any *Aha!* experience (Danek et al., 2016). Further, while *Aha!* moments occur more prominently when people solve particular problems, e.g., insight problems or anagrams, they can also occur while solving non-insight problems as well (Danek et al., 2016).

2.4. Prospectus

Despite the importance of the *Aha!* experience in defining the phenomenon of insight, no comprehensive theory offers an explanation for the experience itself. While considerable progress has been made in improving our understanding of the processes *before* an insight occurs, we do not yet fully understand what happens *after* an insight occurs — why we have an *Aha!* after an insight and why insight feels the way it does. To address this challenge, we present a rational account that complements previous mechanistic theories and provides an explanation of *Aha!* moments at the computational level of analysis (Marr, 1982). Our theory formalizes the connection between metacognition and *Aha!* moments and provides an explanation for why *Aha!* moments occur.

3. A computational account of *Aha!* moments

Our theory is based on the idea that *Aha!* moments reflect prediction errors resulting from inferences about one's own problem-solving (Fig. 1). We first outline a general probabilistic framework for metacognitive inference about problem completion. This specifies how an idealized agent tracking their own progress reasons about their problem-solving ability to predict the time to complete the problem. Then, we discuss how the dynamics of this model can give rise to metacognitive prediction errors — specifically, time prediction errors — that we posit correspond to *Aha!* moments. Importantly, our account describes how observed progress, completion, and prior expectations combine to produce metacognitive prediction errors, which we use in the next section to simulate findings about *Aha!* moments documented in the literature.

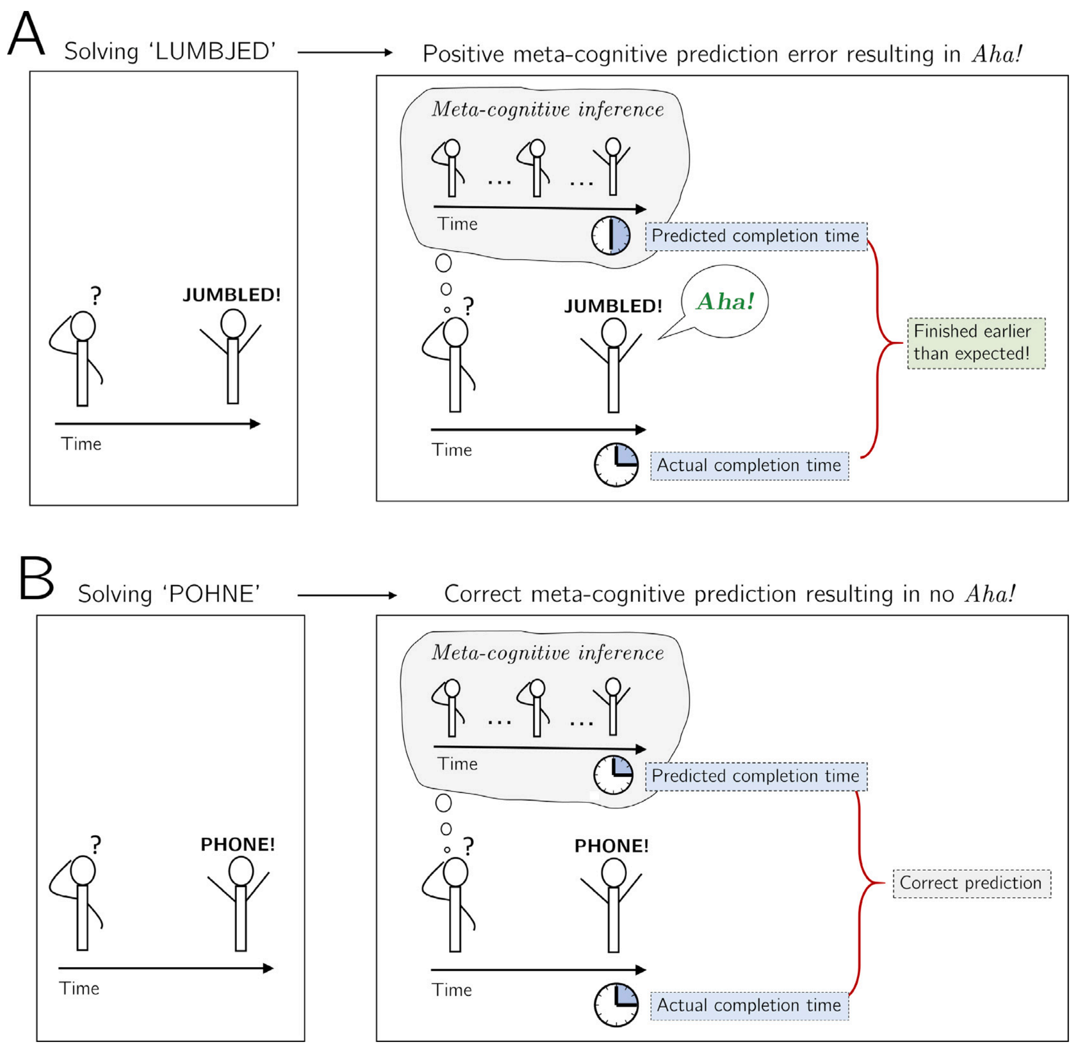

Fig. 1. Metacognitive inference and *Aha!* moments. (A) A decision maker (DM) solving the anagram 'LUMBJED'. Left panel: The DM will take some amount of time to solve the anagram. Right panel: We propose that as they are solving the task, the DM also maintains a metacognitive prediction about how long it will take them to complete the task. When the DM completes a task, their predicted completion time may differ from the actual completion time producing a *time prediction error*. Here, the DM ends up solving the anagram earlier than they expected. We propose that *Aha!* moments correspond to such positive time prediction errors. (B) A DM solving the anagram 'POHNE'. In this case, the DM's metacognitive inference about their completion time is accurate i.e. they solve the task around the same time as they predicted. Thus, the DM will not experience an *Aha!* as they did not make a metacognitive prediction error.

3.1. Metacognitive inference about problem completion

The distinguishing feature of our proposal is that key characteristics of *Aha!* moments arise from *rational inference* over a model of one's own cognition. At first glance, this seems puzzling since presumably people have privileged access to their own cognitive processing and therefore they do not need to make inferences about it. However, people's cognition is not necessarily unitary, and we often learn things about ourselves by observing our interactions with the environment much as we learn about others' cognition from their outward behavior (Cushman, 2020; Fleming & Dolan, 2012; Paulus et al., 2014; Yang et al., 2019; Yeung & Summerfield, 2012). More broadly, work on cognitive control and metacognition suggests that the brain is constantly monitoring and intervening on its own processing, which involves one subsystem engaging in non-trivial inferences about other subsystems (Ackerman & Thompson, 2017; Fleming & Daw, 2017; Schulz et al., 2021). Here, we take motivation from work in artificial intelligence and computational cognitive science which suggests that limited memory and time are a critical bottleneck in computing optimal behaviors, and so predicting *runtimes* of algorithms prior to execution is critical (Eggensperger et al., 2017; Gomes & Selman, 2001; Hutter et al., 2014; Lieder & Griffiths, 2020). In what follows, we consider a decision maker (DM) that monitors itself in order to estimate the amount of time they would take to finish a task. This requires that the DM is capable of making inferences about its own computations, as if they were learning about the cognitive processing of a separate agent.

To formalize this, let us begin by considering a DM that is solving a particular problem. At any point of time, the decision-maker will have observed n *signals of progress*, $s = (s_0, s_1, s_2, \dots, s_n)$, where each signal consists of a time point and completion amount, $s_i = (t_i, c_i)$. Here, time is continuous and completion amount is at least 0 and is equal to 1 when the task is complete. The *completion time*, t^* is the time when the DM first receives a signal of progress at which the task is fully completed i.e., the time at which $c_i = 1$. The DM uses the observed signals of progress s, to estimate the time at which they will complete the task, i.e. compute $E[t^*|s]$. We posit that the DM maintains a metacognitive model of their problem-solving ability and infers a value θ, which represents their belief about the rate at which they are likely to progress in solving the problem (see also Fig. 2). Given a model of their own ability, the DM wants to compute the amount of time they will take to finish the problem. Having observed previous signals of progress s, the distribution over the completion time can be computed

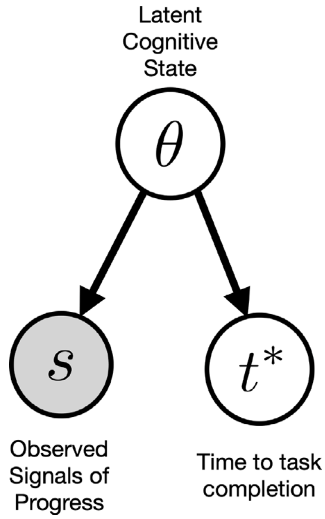

Fig. 2. Graphical representation of the metacognitive inference model. We propose the decision maker maintains a simplified, probabilistic representation of how their cognitive processing interacts with a task. They observe signals of progress (s), which include information about time and task completion amount, and use this to update predictions of completion time (t^*). These predictions are mediated by a latent cognitive state which represent the DM's belief about the completion rate (θ).

by marginalizing out the completion rate, θ:

$$p(t^* \mid s) = \int_\theta p(t^* \mid \theta) p(\theta \mid s) d\theta, \tag{1}$$

where $p(t^* \mid \theta)$ is the probability of a completion time for a fixed value of the completion rate θ, and $p(\theta \mid s)$ is the posterior distribution over θ having observed signals of progress up until the current point in time. The posterior distribution $p(\theta \mid s)$ is computed using Bayes' rule as below:

$$p(\theta \mid s) = \frac{p(s \mid \theta)p(\theta)}{\int_{\theta'} p(s \mid \theta')p(\theta')d\theta'}. \tag{2}$$

In other words, Equation (2) suggests that the DM maintains a distribution over their own possible internal states, updating it dynamically as new signals of progress are observed. This distribution is then used by the DM to compute their expected completion time by using Equation (1) and calculating $E[t^* \mid s] = \int p(t^* \mid s)t^* dt^*$.

3.2. Aha! moments as metacognitive prediction errors

In the previous section, we formulated a decision maker that dynamically reasons about their own interaction with a problem and uses that to estimate their expected completion time, $E[t^*|s]$. Importantly, because the DM is maintaining a prediction about the completion time, there can also be *errors* in their prediction. Specifically, there can be three types of time prediction errors. The first case is when there is no error, i.e. when the time estimate is accurate and the task takes exactly as long as the DM predicted (also see Fig. 1(b)). The second case is a positive error i.e. when the time estimate is greater than the actual time and the DM ends up completing a task faster than they estimated (Fig. 1(a)). Finally, negative error occurs when the task takes much longer than previously estimated. To formalize this, we define the metacognitive *time prediction error* given observed signals of progress s as:

$$\delta = E[t^* \mid s] - t^*, \tag{3}$$

where $E[t^* \mid s] = \int p(t^* \mid s)t^* dt^*$ is the expected completion time with respect to the marginalized distribution in Eq. (1).

We postulate that the metacognitive prediction error as defined in Eq. (3) is analogous to *reward prediction errors*. At a high-level, a prediction error is a difference between an actual and an expected outcome and this signal is commonly thought of as the engine of learning (Glimcher, 2011; Schultz et al., 1997). In the field of reinforcement learning, the main type of prediction error studied is reward prediction errors (Hollerman & Schultz, 1998; Pessiglione et al., 2006; Schultz, 1997). Informally, when we receive a reward that is better than the predicted reward i.e., when there is a *positive* reward prediction error (RPE), then it results in a tendency to repeat the behavior which resulted in the unexpected positive reward. Conversely, when we receive a reward that is worse than the predicted reward i.e., when we make a *negative* RPE, then there is a tendency to avoid that behavior. By contrast, when the reward is exactly as predicted, then there is no prediction error, resulting in the tendency to keep the behavior unchanged. Thus, a central component of our learning mechanism is our desire for positive RPEs, with midbrain dopamine neurons activating for positive RPEs (resulting in positive affective emotions), and avoidance of negative RPEs, with midbrain dopamine neurons showing depressed activity for negative RPEs (resulting in negative affective emotions). Conceptually, time is distinct from reward, but we assume that a decision-maker prefers to spend less time solving a particular problem. This is consistent with a number of accounts, such as those in economics (Hamermesh & Lee, 2007; Whillans et al., 2017) and theories of cognitive costs (Inbar et al., 2011; Sackett et al., 2010), and can be more generally motivated by the observation that time is a fundamentally limited resource. Thus, in the case when an insight occurs i.e., when the DM solves a problem by a sudden breakthrough, then that would result in a positive time prediction error, i.e. $E[t^* \mid s] > t^*$. We suggest that this time prediction error is akin to a positive RPE and it corresponds to the *Aha!* moment, thereby potentially explaining why we feel so good once an insight is reached. Further, because the expected completion time is derived from the DM's belief about their own problem-solving rate θ, such a prediction error can also be interpreted by the DM as evidence that their own competence was underestimated. On the other hand, if the DM makes a negative time prediction error, then $E[t^* \mid s] < t^*$ resulting in a negative RPE. We postulate that this is akin to the frustrating *Gah!* moment that people experience when they find out a task would take much longer than they expected it to take (for example when we are trying to open a door using a key and the key gets stuck or when we encounter an unexpected impasse while trying to solve a problem).

3.3. Implementation details

As per our metacognitive inference model, the DM has a latent state parameterized by θ which mediates the observed signals of progress, s (also see Fig. 2). In other words, this means that the observed data points i.e., the signals of progress s, are generated from a data-generating process with parameter θ. The goal of the DM is estimate the value of the parameter θ and use that to compute the completion time (Equation (1)). To derive quantitative predictions, we need to specify the relationship between s and θ, describing the form of the completion process. We assume that completion follows a standard model for increments in progress called the stationary Gamma process. The stationary Gamma process has been extensively used in the engineering community to model the lifetime of a system and offers the advantage of being relatively simple, very well-studied, and having various mathematical benefits (de Jonge et al., 2017; Singpurwalla, 1995; van Noortwijk et al., 2007; Wang et al., 2015). However, we note that our theory is agnostic to the exact process used to describe completion and we employ the Gamma process primarily as a demonstrative exercise. Since the mathematical details of the Gamma process are not crucial to the main thesis of the paper, we refer readers to Appendix A where we provide an in-depth background of the stationary Gamma process and how we compute completion time using the stationary Gamma process.

3.4. Summary and scope of our work

In this section, we proposed a metacognitive theory of *Aha!* moments in which we considered a decision maker that monitors its interaction with a task in order to predict the time they would take to complete the task. We suggested that the DM maintains a distribution over the completion rate θ, and uses that to compute the expected completion time (Equation (1)). Importantly, as a consequence of maintaining such predictions, the DM could also end up making errors in those predictions. We posit that any such metacognitive time prediction errors are akin to reward prediction errors. Thus, an unexpected insight results in an error similar to a positive RPE, thereby potentially explaining why an *Aha!* moment is accompanied by strong positive emotions.

Before proceeding further, it is important to clarify the scope of our work. First, we note that our framework doesn't make any claims about *where* the time prediction error comes from. During problem-solving, the DM could have a sudden insight due to a re-structuring of the solution space or due to some creative insight — we are agnostic to this process as long as it results in a prediction error. Second, it is important to note that our suggestion of *Aha!* being influenced by time prediction errors doesn't preclude that a non-metacognitive RPE is not involved in the experience of an *Aha!*. Because an *Aha!* moment occurs when a task is finished unexpectedly, this implies that the reward associated with solving the problem is also received unexpectedly. Thus, by definition, an *Aha!* moment is followed by the occurrence of a first-order RPE i.e., a prediction error that doesn't involve metacognitive monitoring. However, in our view, an explanation of the *Aha!* experience just in standard RPE terms doesn't depict a complete picture of the phenomenon — people encounter a variety of unexpected positive events in everyday life and not all of these RPE signals are followed by an *Aha!* experience. Our main proposal is that in addition to the standard prediction error, what makes an *Aha!* experience special is that it involves metacognitive inference and metacognitive surprise — in a way, by monitoring and maintaining predictions about their interaction with the problem, the onset of a sudden insight leads to the DM surprising themselves about their own problem-solving ability which results in the *Aha!* moment.

Finally, we note that metacognitive predictions during problem solving are broader than time predictions alone. People presumably also monitor solvability (Becker et al., 2024), their rate of progress towards a solution (completion rate), their felt proximity to the answer i.e., feelings of warmth (Metcalfe, 1986a), and their own general competence. Here, our focus on time prediction is not supposed to be exhaustive: time is the most tractable quantity to measure experimentally, and crucially, it serves as a common currency from which these other metacognitive quantities can be derived. An estimate of expected completion time implies an estimate of one's completion rate, an assessment of solvability, and even an inference about one's own ability. Time prediction is thus a useful organizing quantity around which richer predictions can be structured. Thus, we view time prediction error as a particularly useful and interpretable window onto the broader metacognitive inference process. In the next section, we present our model simulations that show that taking the positive time prediction error as corresponding to *Aha!* moments and taking the interactions between observed progress, prior expectations, and actual completion time can recover key findings reported in the literature.

4. Explaining findings about *Aha!* moments

Having presented our theory, we now turn towards evaluating our theory against predicting the important findings about *Aha!* moments summarized above. Providing an explanation for Finding 1 is a straightforward consequence of our theory — as discussed above, the crux of our theory is that reward prediction errors occur after the experience of an *Aha!* which explains why we feel so good after an *Aha!*. We now present a series of simulations that show how our model accounts for Findings 2–4.

4.1. Simulation 1: Solving a problem suddenly results in an Aha! moment

Aha! moments are more pronounced when a problem is solved suddenly and unpredictably by a learner as compared to when a problem is solved gradually by a learner (Finding 2). We now present a simulation to demonstrate how *Aha!* moments arise when a problem is solved suddenly by an model. For this and all following simulations, we provide our model with a sequence of data points that provides information about the amount of problem completed, c_i for various time points, t_i. These data points are used by the model to estimate the parameters underlying the Gamma process. The parameter estimates are then used to compute the expected time to finish the problem, $E[t^* \mid s]$. Upon completion of the problem, Equation (3) is used to calculate the resultant metacognitive prediction error i.e. the value of the *Aha!* experience. We further refer the reader to Appendix B where we report all simulation details.

We first present our model with a sequence of data points (completion amounts c_i and time points t_i) wherein the problem is solved gradually such that the completion amount increments gradually over time (also refer to the left panel in Fig. 3(a)). As shown in the right panel of Fig. 3(a), the time estimate (i.e. $E[t^* \mid s]$) predicted by the model quickly converges to value close to the true time estimate, t^*. This subsequently results in a metacognitive prediction error or equivalently the value of the *Aha!* experience to be close to zero when the DM finishes the task (Fig. 3(c)). This is consistent with Finding 2 about *Aha!* moments — when a DM solves a problem gradually (as shown in the above sequence), it doesn't experience an *Aha!* moment.

Next, we present our model with a sequence of data points where the completion amount increases suddenly. The left panel in Fig. 3(b) shows one such sequence wherein the completion amount makes little to no increments until the 90th time point before the completion level increases dramatically and the DM suddenly solves the problem. Here, as a consequence of the completion level making little increments over time, by the $7th$ data point the model ends up estimating a very high value for the expected time to finish the problem (also refer to the right panel in Fig. 3(b)).[1] Thus, when the DM suddenly solves the problem at the $8th$ data point (at $t^* = 100$), the DM's experience of an *Aha!* is very high as the DM's time estimate, was considerably greater than the actual time to complete. This sequence is similar to insightful problem solving wherein learners don't make progress for a while before suddenly arriving to the problem solution (Knoblich et al., 2001; Metcalfe, 1986b; Öllinger et al., 2008). We do not make any commitments about what causes the completion level to increase suddenly and the key contribution of these simulations is showing how the onset of the sudden insight results in an *Aha!* moment. In the first simulation, the observed data was consistent with predictions and the DM could accurately estimate the time to complete the problem. In the second simulation, the observed data was not sufficient to form accurate predictions and the DM erroneously estimates the problem to take a long time to complete. As a result, it was positively surprised upon the arrival of an insight. Taken together, these simulation results are consistent with Finding 2 and show why *Aha!* moments emerge when a problem is solved suddenly and unexpectedly by the DM.

4.2. Simulation 2: Aha! moments are more pronounced after an impasse

Aha! moments are more pronounced when an insight occurs after a considerable period of impasse (Finding 3). We now conduct simulations to demonstrate that our model is consistent with this finding. As before, we provide our model with a sequence of data points that provides information about the amount of problem completed, c_i for various about times, t_i. These data points are used to compute the

[1] Note that for both simulations, the model *a priori* expects the problem to finish in $t = 100$ steps.

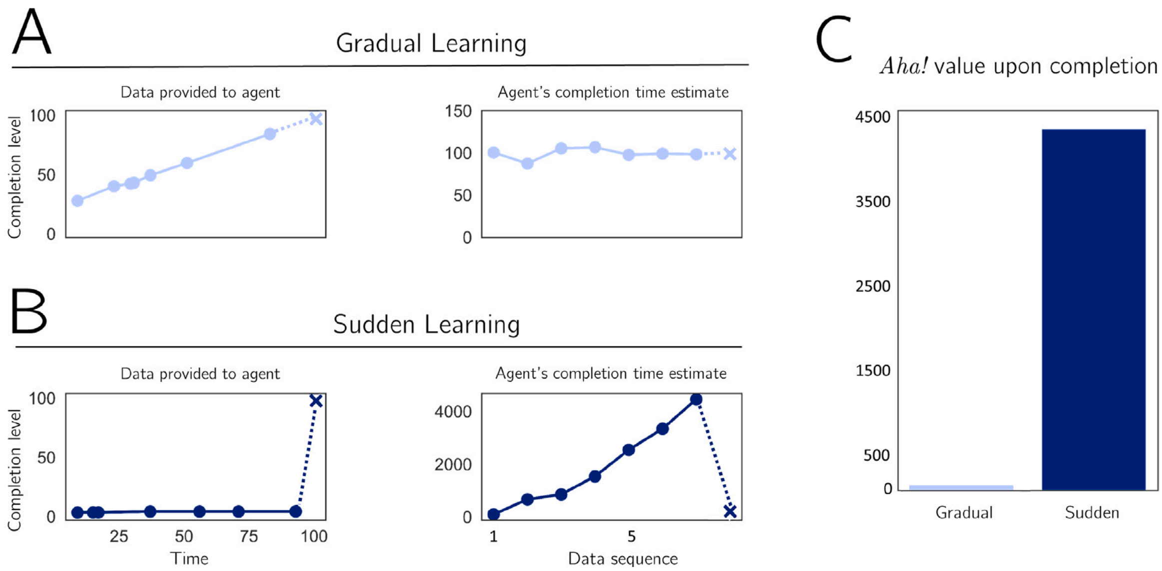

Fig. 3. Simulation 1. Solving a problem suddenly results in an *Aha!* moment. (A) Left panel: Sequence of data provided to the model. Here, the completion level increments steadily and the problem is solved gradually. Right panel: Model's completion time estimate as it observes the data sequence in the left panel. Here, the model's completion time estimate is near to the true completion time. (B) Left panel: Sequence of data provided to the model. Here, the completion level doesn't increase for a considerable period of time and the problem is solved suddenly. Right panel: Here, the model estimates that the problem will take a very long time to complete and is inaccurate about the actual time to solve the problem. (C) Graph plotting the reward prediction error experienced by the two models upon completion of the problem. When the DM solved the problem gradually, it experiences little RPE. In contrast, when the DM made little progress on the problem before suddenly solving it, it experiences high RPE.

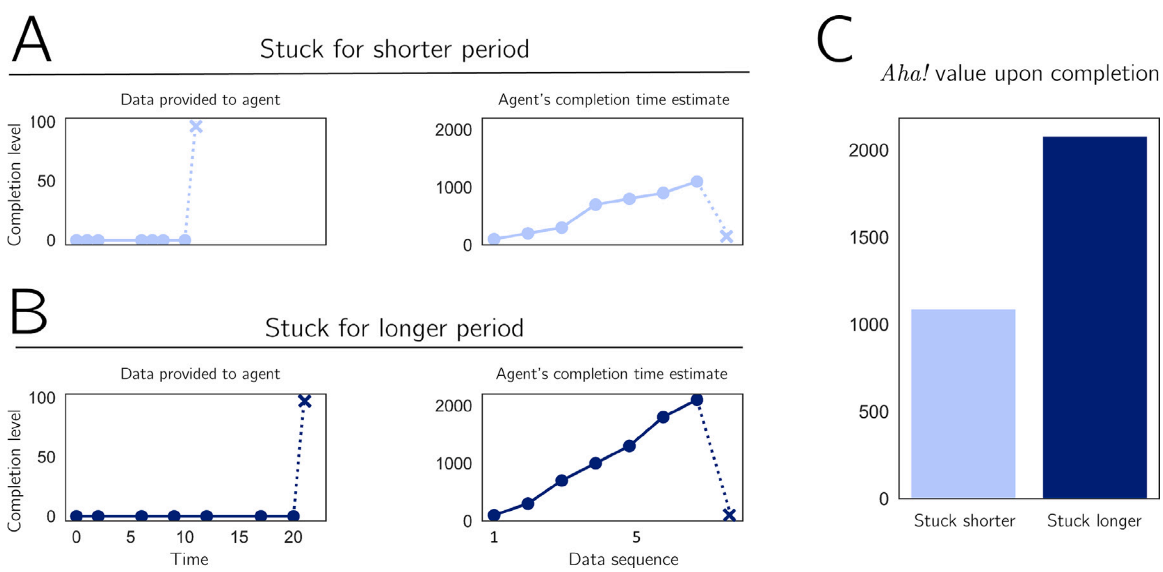

Fig. 4. Simulation 2. *Aha!* moments are more pronounced after an impasse. (A) Left panel: Sequence of data provided to the model. Here, the DM is stuck on a problem before suddenly solving it. Right panel: Model's completion time estimate as it observes the data sequence in the left panel. Here, the model's completion time estimate is quite far away from the actual time to complete the problem. (B) Left panel: In comparison to (a), here, the DM is stuck on the problem for a longer period of time before suddenly solving it. Right panel: In comparison to (a), the model estimates that the problem will take an even longer time to complete and is more inaccurate about the actual time to solve the problem. (C) Graph plotting the reward prediction error experienced by the two models upon completion of the problem and showing that the length of the period of impasse influences the *Aha!* experience.

expected time to finish the problem which in turn is used to compute the resultant prediction error (Equation (3)).

We first present our model with a sequence of 7 data points where the completion level makes no increments over time before the completion level jumps to 100 at the 8th data point ($t^* = 11$, also refer to the left panel in Fig. 4(a)). This results in the DM estimating a very high value for $E[t^* \mid s]$ by the 7th data point (right panel, Fig. 4(a)). Therefore upon the onset of the 8th data point, the DM experiences a high positive RPE when it solves the problem at $t = 11$ (Fig. 4(c)).

Next, to demonstrate how the length of the period of impasse influences the *Aha!* experience, we present our model with a sequence of data points where the completion level doesn't increase for a longer period of time. Specifically, like before, we present our model with a sequence of 7 data points where the completion level makes no increments over time before the completion level jumps to 100 at the 8th data point. However, in comparison to the previous sequence where the impasse ends at $t = 10$, here, the DM doesn't increment its completion level until $t = 20$ (left panel, Fig. 4(b)). In this case, our

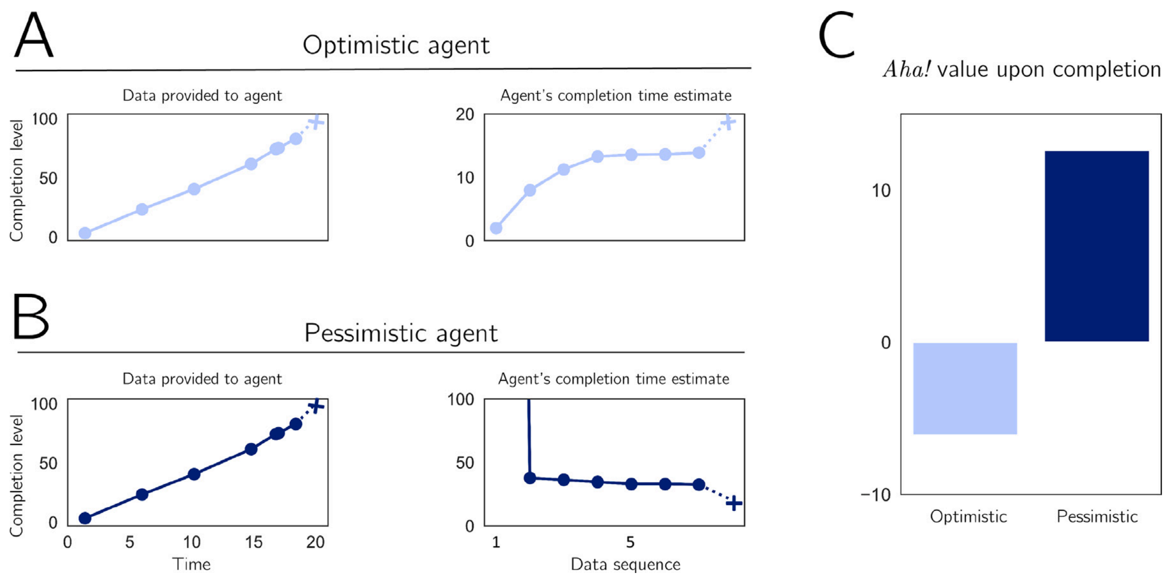

Fig. 5. Simulation 3. *Aha!* moments are influenced by prior expectations. (A) Left panel: Sequence of data provided to the model. Here, the completion level increments constantly and the problem is solved gradually. Right panel: Here, the DM is an optimistic DM that *a priori* believes the problem will take a short period of time to complete. As a result, it estimates the completion time estimate to be lower than the actual time to complete the problem. (B) Left panel: Same sequence of data that was provided to the previous model. Right panel: In comparison to (a), the DM is pessimistic and *a priori* assumes that the problem will take a very long time to complete. As a result, it estimates that the problem will take a longer time to complete compared to the actual completion time. (C) Graph plotting the reward prediction error experienced by the two models upon completion of the problem. Here, the optimistic DM experiences a negative RPE and the pessimistic DM experiences a positive RPE.

model ends up estimating an even higher value for $E[t^* \mid s]$ (right panel, Fig. 4(b)). Thus, in this setting, at the onset of the 8th data point (at $t = 21$), the DM experiences an even higher RPE and a more pronounced *Aha!* moment (Fig. 4(c)). This is consistent with Finding 3 and according to our model, the longer a DM has been stuck on a problem (before an insight), the longer it estimates the completion time to be, resulting in a higher experienced *Aha!* moment when the insight is eventually produced.

4.3. Simulation 3: Aha! moments are influenced by prior expectations

While *Aha!* moments are more prominent for particular problems such as insight problems or anagrams, they can also occur while solving non-insight problems as well (Finding 4). Our theory suggests that metacognitive prediction errors lead to similar *Aha!* effects regardless of the problem type i.e., if people get stuck or mis-predict completion times while solving non-insight problems, then they would experience an *Aha!* moment upon unexpectedly solving them. This helps explain why *Aha!* moments can occur in a wide variety of problems (Finding 4). Our theory also suggests that the same problem can generate different levels of an *Aha!* depending on one's prior expectations. We demonstrate this by conducting simulations with two different DMs that varied in their prior expectations about the time to complete the problem.

We first consider an optimistic DM which *a priori* expects the problem to take a very short amount of time to complete ($E[t^* \mid s_0] = 0$, see Appendix B for the specific parameter values). We then present the model with a sequence of $n = 8$ data points where the completion level, c_t increases linearly with time (with some added noise) such that the problem is completed at $t^* = 21$ (left panel, Fig. 5(a)). Here, because the DM only receives a few data points and it started with a very low estimate for the time to complete the problem, it incorrectly estimates that the problem would finish in a shorter amount of time (right panel, Fig. 5(a)). This results in a negative RPE for the DM when the completion level reaches 100 (Fig. 5(c)).

We next consider a pessimistic DM which *a priori* expects the problem to take a very large amount of time to complete ($E[t^* \mid s_0] = 2000$, see Appendix B for the specific parameter values). We then present the pessimistic DM with the same sequence of data points presented to the optimistic DM. Here, because the DM started with such a high estimate for the time to complete the problem and it receives only a few data points, it incorrectly estimates that the problem would finish in a larger amount of time (right panel, Fig. 5(b)). As a result, in contrast to the optimistic DM, the pessimistic DM experiences a positive RPE for this problem when the completion level reaches 100 (Fig. 5(c)). These results are consistent with Finding 4 and show that the depending on the prior expectations, the *same* problem can result in a different experience of an *Aha!* moment to different DMs.[2]

4.4. Additional predictions

In addition to capturing the above findings, our framework makes some more subtle predictions about *Aha!* moments. We now present simulations that generate these predictions and demonstrate how our framework can be extended to capture *Aha!* moments across a wider range of scenarios.

Simulation 4: *Aha!* moments can occur even when the DM has limited access to the signals of progress. So far, in our simulations, we have assumed that the DM can often access the rate at which they are making progress in solving the problem. This assumption is consistent with prior literature on metacognition and problem-solving (Metcalfe, 1986a; Metcalfe & Wiebe, 1987), and we also explicitly test this assumption in our own experiment later in the paper. Nevertheless, our theory also applies to cases where the DM has very limited access to the rate at which they are making progress. To demonstrate this, we repeat Simulation 1 but only provide our model with two random data points which are far apart from each other. When the problem is

[2] Here, the prior influences the time prediction error because the amount of data points presented to the DMs is considerably low. If the two DMs were presented with more data points, then they would have eventually converged to the right time estimate and would not have experienced a reward prediction error.

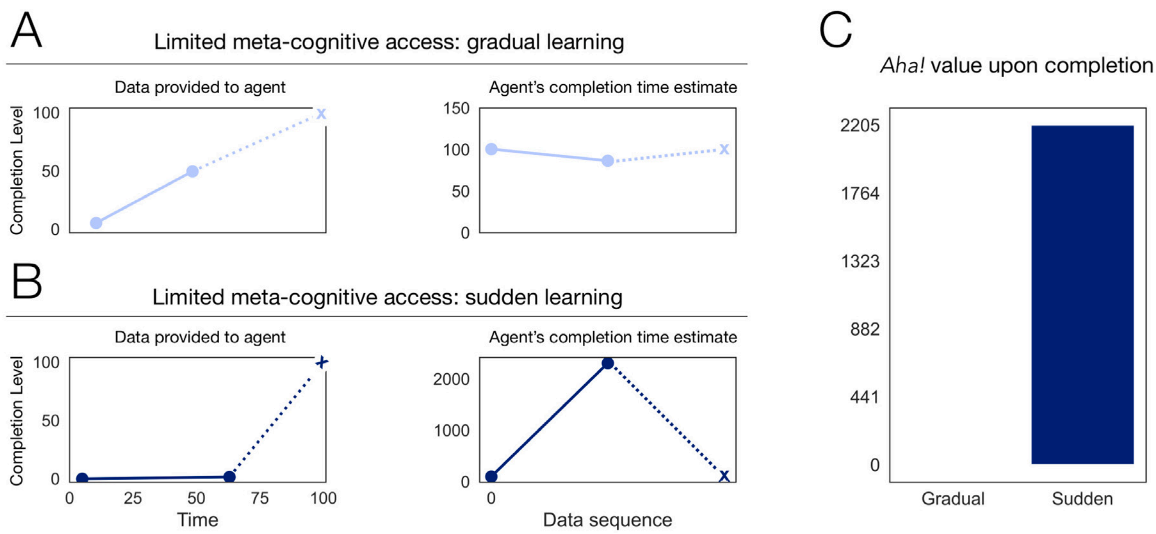

Fig. 6. Simulation 4. *Aha!* moments can also occur when the signals of progress are sparse. (A) Left panel: Data provided to the model. Here, the problem is solved gradually but the model is provided sparse feedback about the rate of progress. Right panel: The DM's completion time estimate is close to the true completion time. (B) Left panel: Data provided to the model. Here, the completion level doesn't increase for a long time before increasing suddenly. The model is again provided sparse feedback about the rate of progress. Right panel: The DM predicts that the problem will take a long time to solve. (C) Graph plotting the reward prediction error experienced by the two DMs upon completion of the problem.

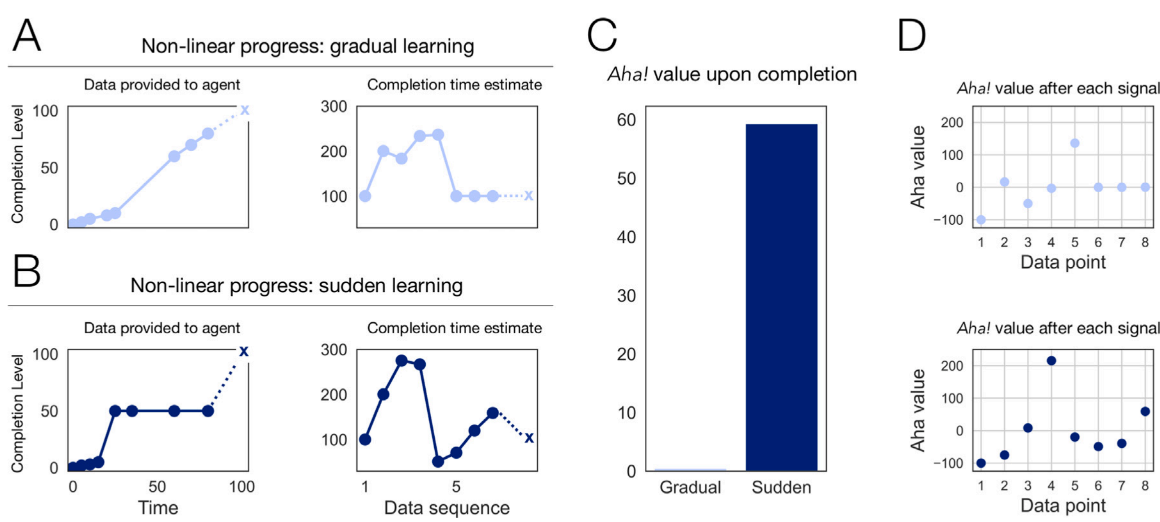

Fig. 7. Simulations 5 and 6. (A) Left panel: Data provided to the model. Here, the problem is solved incrementally but in a non-linear manner. Right panel: The model's completion time estimate. (B) Left panel: Data provided to the model. Here, the completion level increases in short bursts during problem-solving. Right panel: The DM predicts that the problem will take a long time to solve. (C) Graph plotting the reward prediction error experienced by the two DMs upon completion of the problem. (D) Graphs plotting the reward prediction error experienced by the two DMs after each signal of progress.

solved gradually but the DM has only sparse feedback about the rate of progress (left panel in Fig. 6(a)), the time estimate predicted by the DM is still close to the true time estimate (right panel in Fig. 6(a)). Subsequently, the DM experiences a low *Aha!* value upon finishing the task (Fig. 6(c)). We next present our model with a sequence of data points where the completion level makes little to no increments for a long time before increasing suddenly (left panel in Fig. 6(b)). Here, even though the DM only receives sparse feedback, it ends up estimating a very high value for the expected time to finish the problem (right panel in Fig. 6(b)) and subsequently, experiences a very large *Aha!* moment when it suddenly completes the problem (Fig. 6(c). In summary, even when a DM doesn't have perfect access to the rate of progress, our theory predicts that a DM will not experience an *Aha!* moment when a problem is solved gradually and will experience a high *Aha!* moment when a problem is solved suddenly.

Simulation 5: Simulating scenarios that fall in between the extreme cases. Until now, our simulations have only considered cases where the DM makes either linear progress or no progress. We now simulate scenarios which fall in between these extremes. We first present our model with data points where the completion level increases constantly but at different rates at different time points (left panel in Fig. 7(a)). The DM's completion time estimate is inaccurate during the middle of problem-solving but it is more accurate towards the end of the task (right panel in Fig. 7(a)). Thus, the DM experiences a low *Aha!* value upon finishing the task (Fig. 7(c)). Next, we provide our model with a sequence of data points where the completion level increases suddenly at two different time points (left panel in Fig. 7(b)). In the beginning, the completion level rises slowly before increasing suddenly (at the 20th time point). However, the DM gets stuck for a considerable time period before having a sudden breakthrough again. Because of this impasse, the DM incorrectly predicts the completion time (right panel in Fig. 7(b)), and experiences a large *Aha!* moment upon finishing the problem (Fig. 7(c)).

Simulation 6: Aha! moments can also occur during problem solving. As per Equation (3), *Aha!* moments occur when the DM solves a problem earlier than expected. This requires the DM to successfully solve the problem in order to experience an *Aha!* moment. Our theory can be easily extended to account for cases where the DM may experience an *Aha!* moment *during* problem-solving. Specifically, we can define the metacognitive time prediction error as follows:

$$\delta = E[t^* \mid s_0, \ldots, s_n] - E[t^* \mid s_0, \ldots, s_{n+1}], \tag{4}$$

where $E[t^* \mid s_0, \ldots, s_n]$ is the expected completion time when the DM has observed n signals of progress, s and $E[t^* \mid s_0, \ldots, s_{n+1}]$ is the updated completion time when the DM has observed a new signal of progress i.e., when they have observed $n+1$ signals of progress. Thus, according to Eq. (4), the DM will experience *Aha!* moments whenever they make unexpected progress during problem solving (and not just when they unexpectedly complete the problem). Fig. 7(d) plots the experienced *Aha!* value for the two DMs from Simulation 4b after each progress signal. The first DM experiences a high *Aha!* moment after the 5th data point since the rate of progress at this time point is much faster than the previous rate of progress. Following this, the DM progresses at a similar rate and doesn't experience any *Aha!* moments during the rest of the task. The second DM experiences a very large *Aha!* moment at the 4th data point (bottom panel, Fig. 7(d)) as the completion level rises suddenly at this time point. Following this, the DM gets stuck for a considerable time-period and hence experiences negative reward prediction errors from the 5th data point till the 7th data point (analogous to experiencing a *Gah!* moment). At the 80th time point, the DM completes the problem via a sudden breakthrough and experiences a large *Aha!* moment at the end.

4.5. Summary

The simulations presented in this section tested and found support for three key predictions of our model as well as illustrated three more nuanced predictions. Simulation 1 showed that the experience of an *Aha!* is stronger when a problem is solved suddenly compared to when a problem is solved gradually and incrementally. Simulation 2 demonstrated that *Aha!* moments are stronger after a period of an impasse. Simulation 3 showed that *Aha!* moments are not just a property of a particular task and are also influenced by prior expectations of the learner. Simulation 4 showed that *Aha!* moments can occur even when the signals of progress are sparse. Finally, Simulations 5 and 6 presented scenarios where the DM makes non-linear progress and showed how *Aha!* moments can also occur during the course of problem-solving. These results also demonstrate the advantages of our modeling choice, in particular the use of a stochastic process to estimate the completion time. While we could have employed a simpler model to estimate the completion time e.g., by assuming that the underlying data generation process was drawn from a function (instead of a stochastic process) and then doing extrapolation, it would not have been sufficient to explain all of the properties associated with *Aha!* moments. For instance, while an extrapolation model might have yielded similar results for Simulation 1, that model wouldn't produce the results of Simulation 2. This is because an extrapolation model's time prediction wouldn't be affected by the length of an impasse — a simple extrapolation would predict an infinite completion time as soon as the DM encounters an impasse, regardless of the length of the impasse period. While our work is agnostic to the exact computations underlying how people solve the problem of estimating completion time, we believe that our proposed model is useful in understanding when and why errors in those estimates can occur and how those errors relate to phenomena like *Aha!* moments. In the remainder of paper, we present a series of experiments that further test our central hypothesis that metacognitive prediction errors play an important role in driving *Aha!* moments.

5. Experiment 1: Is the experience of an *Aha!* predicted by time prediction errors?

In the previous section we showed that our theory is consistent with some important findings about *Aha!* moments. In addition to being consistent with these findings, our theory also suggests that the experience of an *Aha!* should be predicted by the time prediction error made by a learner. This prediction is unique to our theory — because we suggest that *Aha!* moments arise due to metacognitive time prediction errors, this also implies that any prediction error made regarding completion time should be reflected in the subjective experience of *Aha!*.

We now present a large-scale behavioral experiment desired to test this prediction. Participants solved a series of anagrams and we recorded their judgments of how long it would take them to solve each anagram and their subsequent time prediction errors as well as their rating of the extent they experienced an *Aha!* moment. Based on our theory, we hypothesized that participants' *Aha!* ratings should be highly correlated with their time estimation errors. To the best of our knowledge, our behavioral experiment is the largest experiment to date in the study of *Aha!* moments. Prior to collecting the data, we pre-registered the study on OSF. The pre-registration included the data collection protocol, stimuli, and the data analysis plan (link removed to preserve anonymity).

5.1. Participants

We recruited 1230 participants from Amazon Mechanical Turk (AMT). They earned $1.60 for participation in a study that took around 7–8 min to complete. Prior to data collection, we aimed to obtain a minimum of 30 ratings for each anagram in our dataset and we collected data from AMT until we reached this minimum level for each anagram in our dataset. We obtained a minimum of 30 ratings for each anagram after applying the exclusion criteria outlined in the results section.

5.2. Stimuli

For the purposes of our experiment, we chose to use anagrams as stimuli as various studies have highlighted they are effective in eliciting and measuring insight (Ammalainen & Moroshkina, 2020; Aziz-Zadeh et al., 2009; Laukkonen et al., 2020; Metcalfe, 1986b; Novick & Sherman, 2003; Webb et al., 2018). We first created a dataset consisting of 100 anagrams. To create this dataset, we manually selected 50 words and then created 100 anagrams using these words (refer to the OSF link for the 50 words used and the resulting dataset). These 50 words were of varying length such that 10 were 3 letter words, 10 were 4 letter words, and so on. We further ensured that each of these 50 words only had a single unique grammatically correct configuration i.e. no more than one grammatically correct word could be formed from the given letters. We then used these 50 words to create our dataset of 100 anagrams. To do this, we scrambled each word randomly using two different methods. In the first method, the resultant scrambled word had the same first letter as the original word and the number of edits required to construct the original word from the scrambled word was exactly two. In the second method, we ensured that the resultant scrambled word did not have the same first letter as the original word and the number of edits required to construct the original word from the scrambled word was exactly three. For example, if the original word was 'PHONE', then one example of the resultant anagram using the first methodology would be 'POHNE' and using the second methodology

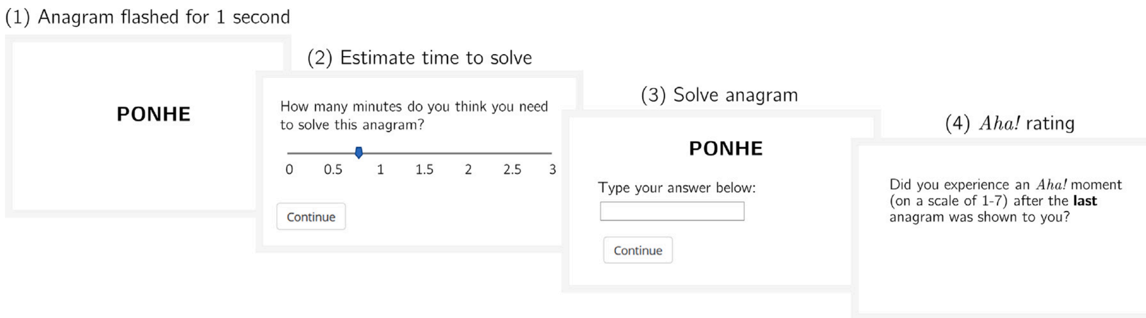

Fig. 8. Procedure for Experiment 1. Participants were asked to solve five anagrams sequentially. Each anagram was first flashed for 0.5 s after which participants were asked to provide an estimate of the time they think it would take them to solve the anagram. After providing the estimate, participants were given unlimited time to solve the anagram after which they were asked to provide their ratings of an *Aha!*.

would be ‘NHOEP’. In this way, repeating this scrambling procedure for each of the 50 words resulted in 100 anagrams in the end (2 anagrams for each word using the two different methods). Note that the anagram created using the first method was easier than the anagram created using the second method.

5.3. Procedure

At the beginning of the experiment, participants were presented with instructions about the task which consisted of informing the participants that they have to solve 5 anagrams, showing examples of what an anagram is, and informing participants about the procedure of the experiment. To solve an anagram, participants were required to construct a word from the presented anagram by using all of the given letters. For example, for the anagram, ‘PNHOE’, the correct answer was ‘PHONE’. To ensure that the participants understood the instructions, a brief quiz was presented after the instructions. If a participant answered the quiz incorrectly, then they were shown the instructions again and presented with the quiz again. This procedure was conducted repeatedly until they gave the correct quiz answer.

After the instructions, the main experiment began in which participants were asked to solve a total of 5 anagrams one after another (also refer to Fig. 8 for the procedure). Before solving the anagram, each anagram was very briefly flashed to the participants for 0.5 s. After the brief display of the anagram, participants were asked to estimate how long they think it would take them to solve the presented anagram. Participants provided their time estimate using a slider which was a continuous scale starting from 0 s and ending at 3 min. After the participants submitted their time estimate, they were then given unlimited time to solve the anagram. Participants could enter their answer to the anagram in a text box (that appeared in the page). If the provided answer was correct, then the participants were taken to the next page. If the answer was incorrect, then the participants were informed that their answer was incorrect and were asked to try again. At any time, if the participant felt that the anagram was too tough for them or if they wanted to skip to the next anagram, they could simply enter ‘next’ in the text box to move to the next screen. After the participants were done solving the anagram (either by providing the correct answer or by typing ‘next’), they were taken to the next page where they were asked to provide their subjective rating of an *Aha!* moment on a scale of 1–7 (where 1 indicated ‘no *Aha!*’ and 7 indicated ‘very high *Aha!*’). This procedure was then repeated for the remaining 4 anagrams i.e. participants were briefly flashed the anagram for 0.5 s, were asked to estimate time it would take to solve the anagram, then were given unlimited time to solve the anagram, and then they finally provide their *Aha!* rating after solving the anagram. After the participants were finished with the experiment, they were de-briefed with the motivation of the task and were paid for their participation. Participants were paid regardless of whether they solved the anagrams correctly and this was also communicated to them in the instructions.

To ensure that each participant would face roughly the same levels of task difficulty, out of the 5 anagrams presented to the participants, one anagram was a 3 letter word (randomly picked from the data-set), one anagram was a 4 letter word, one anagram was a 5 letter word, one anagram was a 6 letter word, and one anagram was a 7 letter word. We further ensured consistency of task difficulty across participants by ensuring that out of the 5 randomly selected anagrams, 2 (or 3) anagrams were anagrams constructed using the first scramble method and 3 (or 2) anagrams were anagrams that were constructed using the second scramble method. The choice of whether 2 or 3 anagrams were picked using the first scramble method (and equivalently whether 3 or 2 anagrams were picked using the second scramble method) was determined randomly. Finally, the order of the 5 anagrams presented to the participants was randomized (i.e. the first anagram presented to the participants could be of any length and could be constructed using any of the two scramble methods and so forth for the remaining four anagrams).

5.4. Results

For the results that follow, we used the following pre-registered criteria for exclusion of data. First, for all participants, we did not analyze the anagrams which the participants couldn’t solve (recall that participants in our experiment could choose to skip to the next anagram by typing ‘next’ if the anagram was too tough for them). This led to the exclusion of 592 datapoints (out of the total 6150 datapoints). Next, for all participants, we did not analyze the anagrams for which the participants took more than 1.5 times the standard deviation of the mean response time of the hardest anagram in the dataset to solve. This lead to the exclusion of a further 25 datapoints.

In order to test our prediction, we computed the mean time prediction error of the participants for each anagram in our dataset. We computed this by using participants’ time estimation to complete the anagram as well as the total time they took to complete the anagram (time prediction error = response time - estimated time). This calculated time prediction error and the *Aha!* rating served as the two key variables of interest for our analysis. Fig. 9 plots the mean time prediction error with the mean *Aha!* rating for the 100 anagrams and shows that as per our hypothesis, *Aha!* ratings have a high correlation with time prediction error ($r(98) = 0.68, p < 0.001$). An interesting finding we observe is that all time prediction errors were positive on average. We believe this arose from the fact that on average, people were more likely to overestimate their completion time for anagrams and finished the anagrams earlier than expected (which perhaps explain why even simple anagrams like HTA have a mean *Aha!* rating > 3).

We next ran a mixed effects linear model with time prediction error and anagram length as fixed effects and anagram ID and subject ID as random effects. As shown in Table 1, time prediction error is

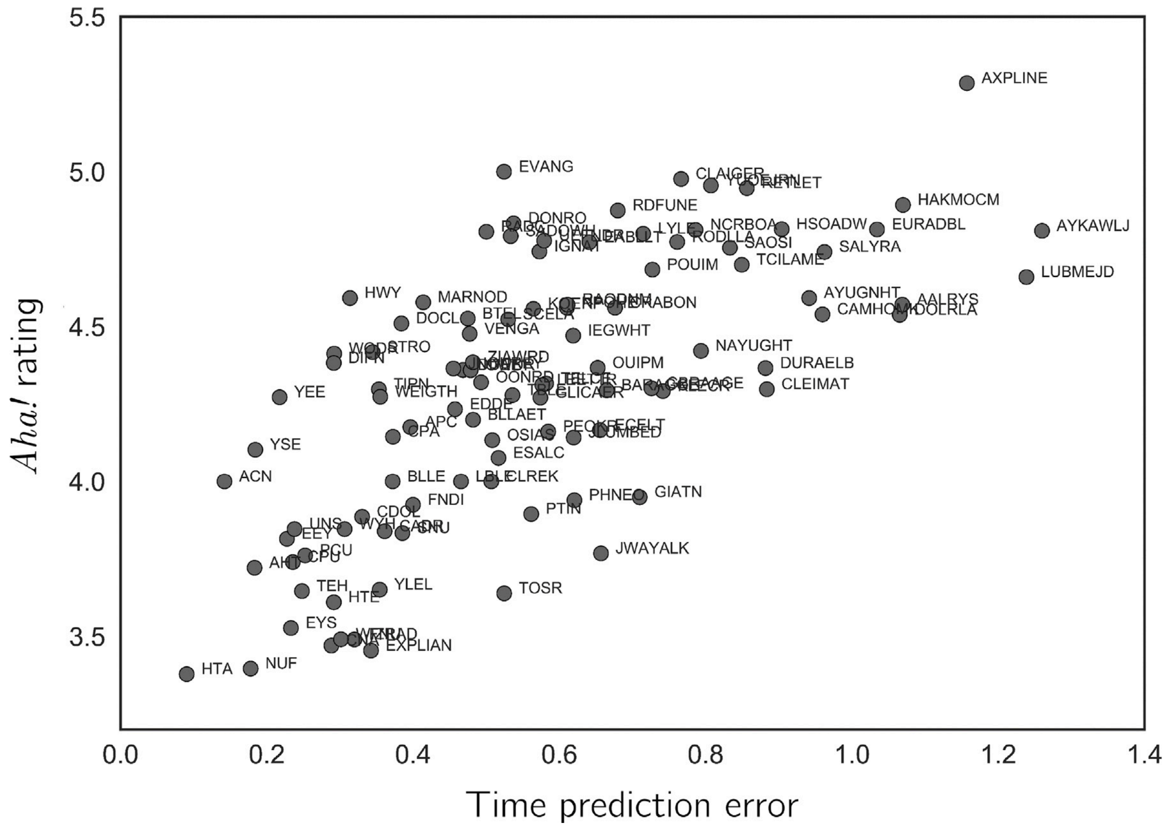

Fig. 9. Results of Experiment 1. ***Aha!*** **ratings are correlated with time prediction errors.** The horizontal axis plots the mean time prediction error and the vertical axis plots the mean *Aha!* rating for the various anagrams in our dataset. As per our prediction, anagrams for which participants made greater time prediction errors also elicited a greater *Aha!*.

Table 1
Experiment 1: Results from the mixed linear effect model.

	Coef.	Std. Error	z	P> |z|	[0.025	0.975]
Intercept	2.445	0.112	21.929	0.000	2.227	2.664
Time prediction error	0.311	0.038	8.272	0.000	0.237	0.384
Anagram length	0.135	0.021	6.490	0.000	0.094	0.176
Anagram ID	0.038	0.009				
Subject ID	2.026	0.077				

positive and significant, confirming that it is a significant predictor of participants' *Aha!* rating.

We next considered whether response time or estimated time correlated with *Aha!* ratings and found that both had a high correlation with *Aha!* ratings ($r(98) = 0.59, p < 0.001$ for response time and $r(98) = 0.74, p < 0.001$ for estimated time). Thus, one concern with Experiment 1 is that *Aha!* ratings might be instead driven by perceptions of difficulty i.e., participants might be providing higher *Aha!* ratings to anagrams which they thought were difficult to solve (which presumably have greater response times and/or have greater initial time estimates). To address this, we conducted an exploratory analysis where we analyzed the more difficult datapoints in our dataset. To do this, we analyzed the points which had a response time greater than the median response time (= 15 seconds) and found that time prediction error had a moderate correlation with *Aha!* ratings ($r(98) = 0.26, p < 0.01$) but both response time ($r(98) = -0.13, p = 0.18$) and time estimate ($r(98) = 0.19, p = 0.06$) were not correlated with *Aha!* ratings. This suggests that time prediction errors might still predict *Aha!* ratings even for anagrams that are perceived as difficult and/or that take a long time to solve.

5.5. Discussion

Results from Experiment 1 showed that participants' experience of an *Aha!* after they solved an anagram was highly correlated with the metacognitive error they made about the time to solve that anagram. This provides support for a key prediction of our theory. However, we also found that *Aha!* ratings were correlated with response time and estimated completion time. This raises the concern that *Aha!* ratings could be driven by perceptions of difficulty. We address this concern in Experiment 2.

6. Experiment 2: Do people monitor their progress during problem solving?

In presenting our theory, we assumed that people have metacognitive access to the rate at which they are solving a task. While Experiment 1 showed the existence of a correlation between *Aha!* moments and metacognitive prediction errors, it didn't show whether people monitor their progress in the first place. Here, we explicitly test this assumption. We also control for task difficulty and response time thereby addressing the potential confound in Experiment 1.

Participants were tasked to solve various problems, each of which were equivalent in difficulty. As participants solved a given problem, they were probed every 10 s and were asked to estimate how close they thought they were towards solving the problem. We hypothesized that people's closeness ratings would exhibit more sudden changes over time for problems that elicited higher *Aha!* ratings compared to the ones that generated lower *Aha!* ratings. Prior to collecting the data, we pre-registered the study on OSF. The pre-registration included the data collection protocol, stimuli, and the data analysis plan (link removed to preserve anonymity).

6.1. Participants

We recruited 502 participants from Prolific. Participants earned \$1.00 for participation in a study that took around 3–4 min to complete.

6.2. Stimuli

To control for task difficulty, we used two different kinds of problems: anagrams and analytical word problems. Prior research has shown that anagrams typically generate higher *Aha!* ratings than analytical problems (Metcalfe & Wiebe, 1987; Webb et al., 2018). Since we aimed to select anagrams that were similar in difficulty to the analytical problems, we hypothesized that the difference in the *Aha!* ratings between the two couldn't be explained by difficulty. This allows us to test whether this difference would be instead explained by differences between participants' metacognitive process and subsequent prediction errors.

We first created a dataset of 30 anagrams and 30 analytical word problems, where the 30 anagrams were taken from the previously described 100 anagram dataset (Experiment 1) and the 30 analytical problems were adapted from previous studies on problem-solving (Gilhooly & Murphy, 2005; Laukkonen et al., 2021; Metcalfe & Wiebe, 1987; Schooler et al., 1993; Webb et al., 2018). From this dataset, we aimed to select anagrams and analytical problems that were perceived to be equally difficult. To that end, we conducted a large-scale norming study where we collected participants' perceived difficulty and response time for each problem in our dataset (we recruited $n = 850$ subjects and collected at least 50 ratings for each problem). From the norming study, we selected five anagrams and five analytical problems for Experiment 2. The anagrams were: ['EURADBL', 'MATCILE', 'AYKAWLJ','LUBMEJD', 'NITGA']. The analytical problems were: ['Eric is trying to get in shape. He wants to do this by climbing stairs. He starts on the third floor, climbs up five stories, down seven, up six, down three, up four, down four, up five, down three, and then up two again. What floor is he on now?', 'Ben spent \$42 for shoes. This was \$14 less than what he spent for a shirt and twice more expensive than the tie. The tie was \$20 cheaper than the jeans. How much was the jeans?','Andrew is twice the age of Brian, Charles is three times older than Brian, and the sum of their ages is 60 years. How old is Charles?','Erin had 24 marbles and Evan had 3 marbles. Erin gave some of her marbles to Evan. Now Erin has exactly double the number of marbles that Evan has. How many marbles did Erin give to Evan?','Lebrun, Lenoir, and Leblanc are, not necessarily in that order, the accountant, warehouseman, and traveling salesman of a firm. The salesman is unmarried. Both Lebrun and Lenoir are married. Lebrun is not the accountant. What job does Lenoir have? Accountant, warehouseman, or salesman?']. These problems were equivalent in difficulty ($M = 3.2, SD = 0.18$), with a one-way ANOVA revealing that the ten problem were not significantly different from each other, $F(9397) = 0.59, MSE = 1.68, p = 0.81$. Further, the five anagrams were rated to be similar in difficulty to the five analytical problems, $t(630) = 0.42, p = 0.7$. We refer the reader to Appendix D for further details about the norming study and the stimuli used.

6.3. Procedure

At the beginning of the experiment, participants were instructed that they had to solve one anagram and one analytical word problem, and were also provided examples of sample anagrams and word problems. After the instructions, the main experiment began in which they were given two problems to solve one after another (one randomly selected anagram and one randomly selected analytical problem). For each problem, participants were displayed the problem text and were provided a text box below the problem description. They were asked to provide their answer in the text box, which would then allow them to move to the next screen. At any time, if a participant decided they were unable to solve the given problem, they could move to the next problem by simply typing 'next' in the provided text box.

As participants were solving the problem, they were prompted every 10 s and were asked to provide their ratings of closeness on a scale of 1 to 100 ('How close are you towards solving this problem?'). This prompt was shown every 10 s until the participants provided their respective answer in the provided text box. After participants provided their answer, they were asked to provide their subjective rating of an *Aha!* moment on a scale of 1–7 for the previous problem. Participants were given unlimited time to solve a problem and were paid regardless of their performance (and were instructed about this as well).

Table 2
Experiment 2: Results from the mixed linear effect model.

	Coef.	Std. Error	z	P> |z|	[0.025	0.975]
Intercept	2.98	0.31	9.72	0.00	2.38	3.58
Maximum increment	0.02	0.004	5.09	0.00	0.01	0.03
Response time	0.00	0.003	0.397	0.69	−0.004	0.01
Subject ID	1.74					

6.4. Results

For all participants, we excluded the problems for which they didn't provide the correct solution. Our final analysis was performed over 313 datapoints for anagrams (out of the total 502 datapoints) and 331 datapoints for analytical problems (again out of the total 502 datapoints).

First, as expected, we found that anagrams generated higher *Aha!* ratings ($M = 5.27, SD = 1.59$) compared to the analytical problems ($M = 3.46, SD = 1.59$), $t(642) = 13.8$, $p < 0.01, d = 1.09$ (Fig. 10(a)). We then analyzed people's closeness ratings during the course of problem solving for both anagrams and analytical problems. As shown in Fig. 10(b), there was a considerable difference in the metacognitive process during problem solving for the two kinds of problems. Closeness ratings increased slowly for anagrams before exhibiting a sudden jump (which happened right before the problem was solved). For the analytical problems, the closeness ratings increased more gradually. For instance, 20 s before the solution time, participants' average closeness rating for anagrams was 19.36 which was significantly lower than the closeness rating for analytical problems (= 36.54), $t(365) = 6.506$, $p < 0.01, d = 0.651$. Next, for each problem that each participant solved, we calculated the *maximum* increment in the closeness rating i.e., we computed the maximum "jump" the participants made while solving a problem. This computed maximum increment served as the metacognitive prediction error made during problem solving. Consistent with our hypothesis, the maximum increment made while solving anagrams was significantly greater than the maximum increment for analytical problems, $t(642) = 13.41$, $p < 0.01, d = 1.15$ (Fig. 10(a)). Further, maximum increment was correlated with *Aha!* ratings, $r(644) = 0.22, p < 0.01$, but participants' response time was not correlated with *Aha!* ratings, $r(644) = -0.04, p = 0.37$, which successfully eliminates the potential confound from Experiment 1. We then ran a mixed effects linear model with maximum increment and response time as fixed effects and subject ID as random effect and found that maximum increment is positive and significant but response time is not significant (Table 2). This confirms that the metacognitive prediction error, as measured by the maximum "jump" during problem solving, is a significant predictor of *Aha!* rating while controlling for response time. To complete our analysis, we calculated the *average* increment in the closeness rating after excluding the maximum increment and found that the average increment for analytical problems was greater than the average increment for anagrams, $t(642) = 5.18$, $p < 0.01$, $d = 0.34$. This is consistent with the results in Fig. 10(a) and suggests that people make constant increments while solving analytical problems but they progress more slowly while solving anagrams (before exhibiting a sudden jump).

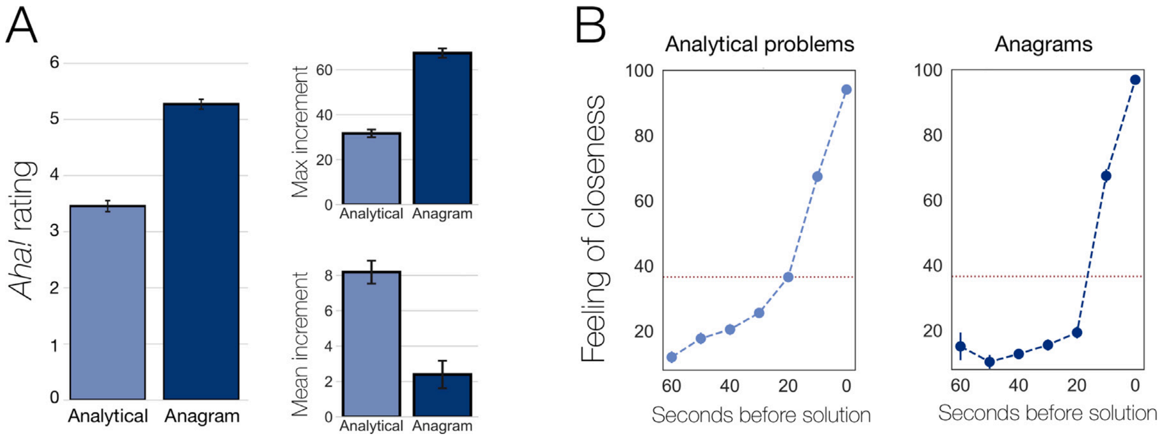

Fig. 10. Results of Experiment 2. Prediction errors during the metacognitive monitoring process predict *Aha!* moments. (A) Left panel: The vertical axis plots the mean *Aha!* rating for the analytical problems and the anagrams. For this and all following graphs, errors bars show the standard error of the mean. Right panel: (Top) Graph plots the mean *maximum* increment in the closeness ratings during the course of problem solving. The maximum increment made while solving anagrams is higher than that of analytical problems. (Bottom) Graph plots the mean *average* increment in the closeness rating after excluding the maximum increment. The average increment made during solving analytical problems is greater than that of anagrams. (B) Participants' mean closeness ratings during the course of problem solving for anagrams and analytical problems. Anagrams are solved suddenly whereas analytical problems are solved more incrementally.

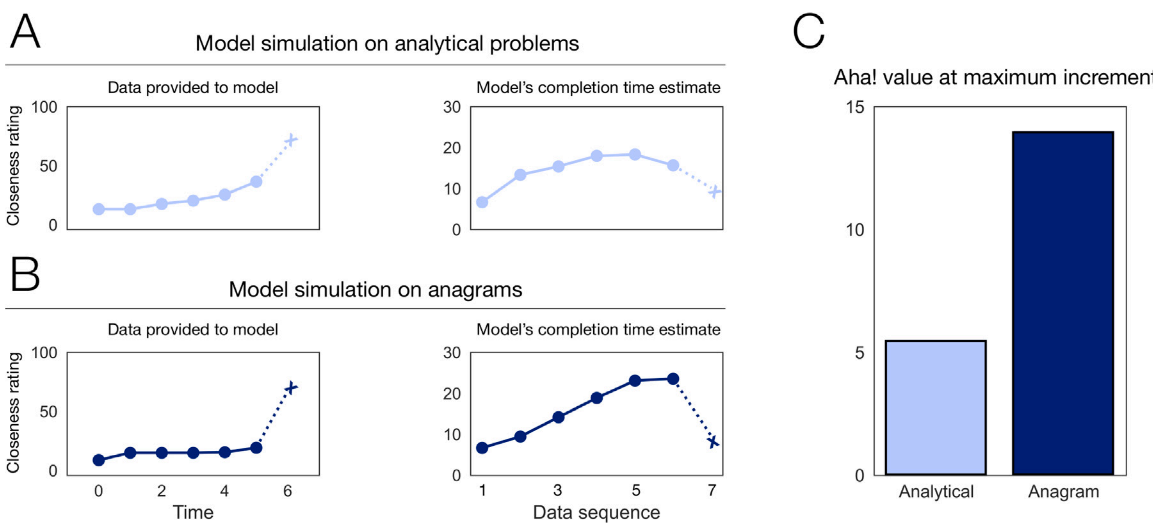

Fig. 11. Results of Experiment 2. Model simulation on participants' data. (A) Left panel: Participants' data from Fig. 10(b) which was provided to the model. Right panel: The DM's completion time estimate for the analytical problems. (B) Left panel: Participants' closeness ratings for the anagrams which was provided to the model. Right panel: The DM's completion time estimate for the anagrams is higher compared to the analytical problems. (C) The model predicts a higher *Aha!* rating for the anagrams compared to the analytical problems at the time point of the maximum increment.

6.5. Model simulation

To further test our theory, we conducted simulations using data from Experiment 2. We provided our model with participants' closeness ratings for both anagrams and analytical problems (from Fig. 10(b)). This sequence of data points provided information about the amount of problem completed, c_i for various time points, t_i.[3] As before, we used these data points to compute the expected time to finish the problem, $E[t^*|s]$. We then used Equation (4) to compute the value of the *Aha!* experience at the time point where the maximum increment in completion level occurred. Results in Fig. 11 show that based on the provided participant data, our model predicts a significantly higher *Aha!* rating for the anagrams compared to the analytical problems (also refer to Appendix D for individual-level model validation).

6.6. Discussion

Experiment 2 found that people indeed monitor their progress during the course of problem-solving. This was consistent with prior empirical research (Metcalfe, 1986a; Metcalfe & Wiebe, 1987) and also confirmed a key assumption of our theory. Importantly, we found that the prediction errors made by the participants during this metacognitive monitoring successfully predicted their feelings of an *Aha!* experience, even after controlling for response time and task difficulty. Model simulations on participants' data provided further support to our theory. Together, these results suggest that the metacognitive prediction error made during problem solving might be an important predictor of the *Aha!* moment. To test the generality of our theory, we conducted another experiment using classic "insight" problems. Consistent with prior research (Webb et al., 2018), insight problems

[3] Note that we made the data monotonic prior to model fitting as the Gamma process requires the data to be monotonically increasing. However, our main results remain the same even if we use a non-increasing stochastic process.

elicited higher *Aha!* ratings than analytical problems but lower than anagrams. Crucially, in line with our theory, the maximum "jump" made by participants while solving insight problems was higher than analytical problems but lower than anagrams. Thus, metacognitive prediction error predicted *Aha!* ratings for insight problems as well. We refer the reader to Appendix E for more details about the study design and results.

7. General discussion

The proverbial *Aha!* moment, though rare, is an essential part of the scientific experience. Throughout our scientific journey, we have all presumably experienced moments of extraordinary joy and elation in which the solution to a difficult problem suddenly appeared in the mind. In this article, we presented a theoretical framework that explains why an *Aha!* moment feels the way it does by suggesting that these moments correspond to metacognitive time prediction errors and that they emerge from inferences about one's own problem solving ability. In a sense, an *Aha!* moment is akin to an *intrinsic* reward that arises because we end up positively surprising ourselves about our own capabilities. Across two large-scale behavioral experiments, we find support for our theory as they demonstrate the existence of a possible link between metacognitive prediction errors and *Aha!* moments. While the idea that the experience of *Aha!* involves metacognition has been noted previously in several influential works (Metcalfe, 1986a; Metcalfe & Wiebe, 1987), our theory formalizes this connection and highlights the important role that meta-cognitive prediction errors play in driving this phenomenon (see also: Becker et al., 2024; Moroshkina et al., 2024).

In the remainder of the paper, we explore how our theory relates to the Eureka heuristic theory, discuss how metacognitive errors might interact with other metacognitive inferences, consider potential shortcomings of our work, and outline questions for future research.

7.1. Connections with the Eureka heuristic theory

An important finding, which we have not yet discussed, is that *Aha!* moments are often accompanied by a strong sense of certainty (Danek & Wiley, 2017; Laukkonen et al., 2020; Salvi et al., 2016; Webb et al., 2016). Why is it that solutions accompanying *Aha!* moments feel correct and true? According to the Eureka heuristics theory, *Aha!* moments signal the extent to which past knowledge and the current context align with the newly generated idea, which is then used as an aid for quick and efficient decision-making (Laukkonen, 2019; Laukkonen et al., 2020, 2018). The feeling of insight thus provides an adaptive benefit by signaling which ideas one should pay attention to and which should they ignore (Laukkonen et al., 2020; Valueva et al., 2016). We believe that our proposal is in line with this theory. Just as reward prediction errors are used by the mind to compute the value of an action (which then aids subsequent decision-making and learning), metacognitive errors (in the form of *Aha!* moments) could be the building block upon which the metacognitive system evaluates the worth of an idea. In other words, while metacognitive errors by themselves would not invoke a sense of certainty, they could be used as a signal by the DM when computing the value of an idea. Thus, complementing the Eureka heuristic theory, our theory suggests that the onset of a strong positive metacognitive error could prompt and bias the DM to have greater confidence in the solution, thereby explaining why *Aha!* moments not only feel pleasurable but also true.

7.2. Interaction with other metacognitive inferences

Given that time is an important cognitive resource, we expect that an estimate of completion time could be vital for cognitive control and it will be interesting to study its interaction with various other metacognitive inferences. For example, an estimate of completion time along with other cues such as fluency, confidence etc., could be used by the decision maker (DM) to compute the difficulty or solvability of a problem (Ackerman & Thompson, 2017; Aguilar-Lleyda et al., 2020; Lauterman & Ackerman, 2019; Topolinski et al., 2016).[4] This judgment can subsequently influence the learner's curiosity (Baranes et al., 2014; Berlyne, 1954; Dubey & Griffiths, 2020; Gruaz et al., 2025), study-time allocation strategies (Metcalfe & Kornell, 2005; Son & Metcalfe, 2000; Undorf & Ackerman, 2017), and even beliefs about their competence (Ladd & Price, 1986; Lucca et al., 2020; Rouault & Fleming, 2020). Thus, whenever the DM experiences a prolonged impasse during problem solving, they would also estimate that the problem is difficult to solve. If the DM manages to solve this difficult problem via an unexpected sudden insight, then that could result in a feeling of 'accomplishment' and having a sense of 'pride' (as the DM is more competent than they had predicted). In a study of undergraduate students taking a foundational mathematics course, Liljedahl (2005) showed that *Aha!* moments also resulted in improved confidence and a positive change in the belief and attitudes of the students towards mathematics. We find this potential connection to be very exciting — perhaps feelings of 'accomplishment' and 'pride' are also a form of metacognitive errors induced when the DM realizes they are smarter or more capable than they previously thought.

This suggests that metacognitive errors might be a reflection of learning about yourself, and feelings of 'pride' might be signaling that we are better than we previously thought and are a learning signal for our minds to update models of ourselves (Savinova & Korovkin, 2022). Indeed, this also explains why *Aha!* moments can sometimes elicit a sense of 'drive' and have an energizing effect on problem solving behavior (Danek & Wiley, 2017; Mercier et al., 2025; Yu et al., 2024). We believe that inducing *Aha!* moments can be a potential pathways towards improving the confidence and self-belief of students and can even be used to develop curiosity and perseverance. Studying this question especially in classroom settings could be a promising direction for future research.

7.3. Limitations and future directions

Limitations of experiments. Our experiments were conducted on anagrams and insight problems because previous research showed their effectiveness in studying *Aha!* moments (Ammalainen & Moroshkina, 2020; Aziz-Zadeh et al., 2009; Laukkonen et al., 2020; Metcalfe, 1986b; Novick & Sherman, 2003; Webb et al., 2018). However, these problems are presumably limited in the aspects of *Aha!* moments they can generate and our findings should be investigated across a wider range of stimuli. In Experiment 2, based on prior research, we used warmth ratings to capture metacognitive patterns during the course of problem-solving (Laukkonen & Tangen, 2017; Metcalfe & Wiebe, 1987). However, warmth ratings were collected only every 10 s, which is not ideal since an insight could occur within that time window. In the future, it might be valuable to use a more continuous measure, e.g. a dynamometer (Laukkonen et al., 2021), to better capture the metacognitive process during problem solving. In general, it is important to investigate the generalizability of our findings using different types of materials and cues in order to develop a greater understanding of *Aha!* moments. Two particularly promising paradigms would be compound remote associate (CRA) problems (Bowden & Jung-Beeman, 2003) and magic tricks (Danek et al., 2014). CRA problems require finding a single word linking three seemingly unrelated cues, while magic tricks involve suddenly grasping a hidden mechanism after a period of puzzlement. Both of these problems are characterized by flat progress trajectories followed by sudden resolution. Our model predicts that this pattern should produce large metacognitive prediction errors

[4] Judgment of solvability could even be inferred directly from the belief about completion rate θ i.e., the metacognitive inference model (Fig. 2).

and correspondingly strong *Aha!* experiences. Future work testing our framework with these and other paradigms would help establish the breadth of the metacognitive prediction error account.

Time prediction as one facet of metacognition. An important limitation of our work is that our framework operationalizes metacognitive prediction error specifically in terms of time. As noted previously, metacognitive monitoring during problem solving likely encompasses a richer set of predictions including solvability (Becker et al., 2024), completion rate, and feelings of warmth (Metcalfe & Wiebe, 1987), among others. Our focus on time prediction captures only one facet of this broader system, and it remains an open question how errors in these other metacognitive quantities contribute to the *Aha!* experience, either independently or in combination with time prediction errors. Future work that jointly models multiple metacognitive predictions and examines their respective contributions to *Aha!* moments would be a valuable extension of the present framework.

***Aha!* without metacognitive errors.** An important future question would be investigating whether *Aha!* moments can occur even when a learner makes little to no metacognitive error. For instance, it is plausible that one could experience an *Aha!* moment because they discovered a surprising or a novel solution to a problem (without making any metacognitive prediction errors). This might especially be the case for perceptual problems such as Mooney faces or illusions such as the Duck–Rabbit illusion. Future work should examine variations of *Aha!* moments across different learners and tasks to better understand the boundary conditions of our theory.

***Aha!* when not actively monitoring.** In our simulations, we only considered cases where the DM was actively monitoring themselves during the course of problem solving. However, *Aha!* moments can also occur unexpectedly about things long forgotten or about problems that one isn't actively working upon (Gilhooly, 2016; Gilhooly et al., 2013). Although we have not formally considered such scenarios, we believe that our theory might still be applicable to these situations. For instance, consider a DM that believes that they could never solve a particular problem and they give up trying to solve it. In this case, it might be plausible that the DM develops a sub-conscious or unconscious metacognitive expectation about the completion time (for e.g., in the extreme case, the DM might think they could never solve the problem within their lifetime). If the DM then has a sudden realization and unexpectedly solves the problem, then that would induce a reward prediction error (since the DM unexpectedly knows the answer to the question) as well as a metacognitive error (since the DM solved the problem against their sub-conscious expectation), which would lead to an *Aha!* moment. Future work could consider formalizing this intuition further and describe how metacognitive errors might arise even when someone is not explicitly working on a problem. Relatedly, while our framework focuses on the metacognitive prediction error that accompanies sudden problem resolution, it does not directly address the role of mental restructuring in producing that resolution. We see restructuring-based accounts and our own as complementary with the former explaining how insight occurs, the latter explaining why it feels the way it does. Future work examining both processes jointly would be a valuable extension.

Connections with the planning fallacy. It will also be intriguing to study the connection of our work with the planning fallacy, a phenomenon where people *underestimate* the time it will take to complete a task (Buehler et al., 1994; Kruger & Evans, 2004; Sanna et al., 2005). This is in contrast with the scenarios considered in this paper (Simulations 1–6), where the decision-maker overestimated the time to complete a task. Why is it that people often predict shorter completion times for many real-world projects (such as writing a research paper) but overestimate completion times for tasks such as solving a puzzle or an anagram (especially during an impasse)? Studying this question computationally through a more detailed model of time prediction might provide valuable insights about the two phenomena.

Metacognitive errors and learning about the self. The capacity to accurately predict future events enables adaptive agents to interact effectively with their environments. A rich literature in neuroscience and psychology has shown that prediction errors are a central part of achieving this capacity as they are used to drive *learning* and make future predictions more accurate (Daw & Tobler, 2014; Niv & Schoenbaum, 2008; Pessiglione et al., 2006; Rescorla & Wagner, 1972; Schultz & Dickinson, 2000; Watkins & Dayan, 1992). Having shown that the notion of metacognitive prediction errors is consistent with empirical findings on *Aha!* moments, a promising area for future research would be investigating whether metacognitive prediction errors also offer similar learning benefits. Perhaps metacognitive errors are a mechanism via which adaptive agents come to develop a more accurate model of their internal selves (Frömer et al., 2021). In this vein, an exciting area for future research would be studying how people learn about themselves and the possible role metacognitive errors play in learning about the self (see also the hypercorrection effect e.g., Butler et al., 2011, 2008; Butterfield & Metcalfe, 2001; Kulhavy et al., 1976).

7.4. Conclusion

The unique *Aha!* moment has captured the attention of psychologists for over hundred years. Here, we have taken a step towards explaining why a sudden discovery feels so good to us: it reflects completion of a cognitive task earlier than we had predicted based on our experience so far and provides a positive surprise about our own competence. Our work outlines the importance of these metacognitive prediction errors and demonstrates how these errors relate to key phenomena associated with *Aha!* moments. We hope that the framework developed here provides a useful foundation for future work and encourages further computational and empirical investigations about metacognitive errors.

CRediT authorship contribution statement

Rachit Dubey: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Mark Ho:** Writing – review & editing, Methodology, Investigation, Conceptualization. **Hermish Mehta:** Formal analysis, Conceptualization. **Thomas L. Griffiths:** Writing – review & editing, Supervision, Methodology, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A

A.1. Modeling problem completion under the stationary Gamma process

As per our meta-cognitive inference model (Fig. 2), the observed data points are generated from an unknown process and the agent's goal is to infer the latent parameters of the process. Here, we assume that the observed data points are generated from a stationary Gamma process and thus, the goal of the DM is to estimate the parameters underlying the Gamma process.

Intuitively, a stochastic process describes the movement of a random variable through time. For example, the random variable could be the price of a stock or the position of a particle moving through a fluid. In the context of problem solving, a stochastic process describes how the completion amount changes its value over time i.e., it describes how the amount of the problem that the DM completes changes as a function

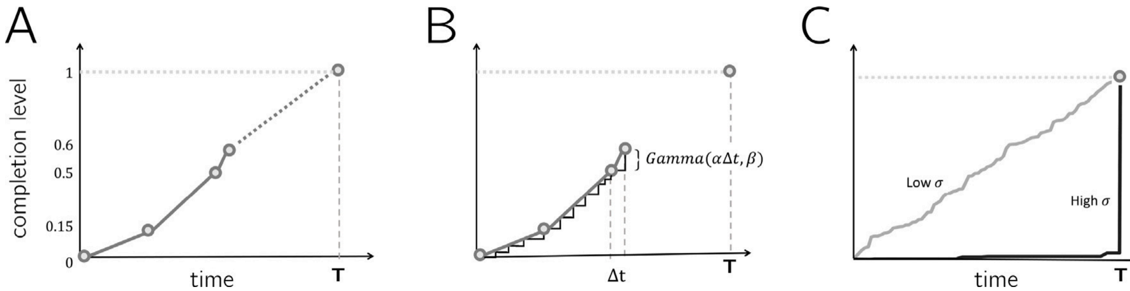

Fig. 12. A Gamma process model of time estimation (A) While solving a problem, the DM makes some increments at various time points towards completing the problem i.e., finishing 100% of the problem. (B) We model this process under a Gamma process framework where the increment made by the DM between successive time-points Δt, is determined by a gamma distribution with scale parameter equal to $\alpha\Delta t$ and rate parameter β. Given this knowledge, the goal of the DM is to estimate the time at which it would complete 100% of the problem. (C) Sample paths of a stationary Gamma process with low σ and high σ values respectively (where $\sigma = \sqrt{\alpha}/\beta$).

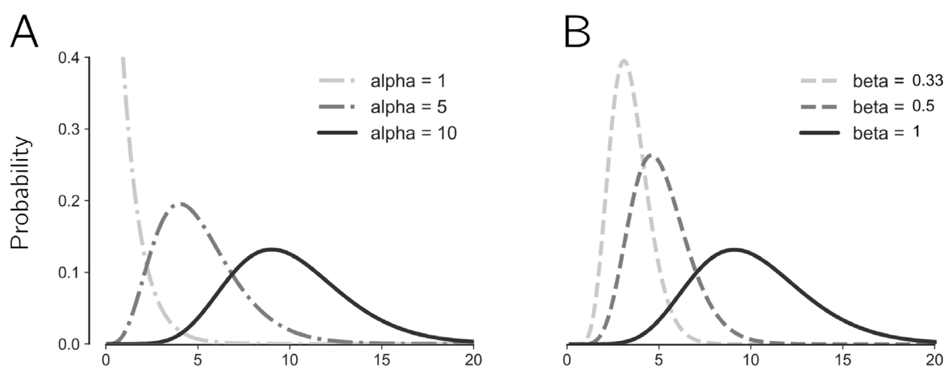

Fig. 13. The role of the parameters, α and β in determining the shape of the Gamma distribution (A) How the probability distribution function (PDF) of the Gamma distribution changes when α changes and β is kept fixed ($\beta = 1$). The shape parameter α represents the number of events that one is waiting to occur. For a fixed value of β, a larger α implies one is waiting for more events to occur and thus the graph is shifted to the right. (B) How the PDF of the Gamma distribution changes when β changes and α is kept fixed ($\alpha = 10$). The scale parameter β represents the mean waiting time for an event. For a fixed value of α, increasing β implies that the mean waiting time increases and the graph again shifts to the right. Note that, although both the parameters have an effect on the shape of the PDF, a change in β shows a sharper change compared to a change in α.

of time (refer to Fig. 12(a)). Depending on the stochastic process we choose, the random variable under consideration can move in different ways over time. For our purposes, we model the problem completion process using the stationary Gamma process.

The Gamma distribution. At the heart of the Gamma process lies the Gamma distribution — in a Gamma process, the increments that a random variable makes over time are Gamma distributed. Before describing the process, we first briefly review the Gamma distribution. If X is a continuous random variable, then the probability distribution function (PDF) of the Gamma distribution is given below:

$$f(x) = \begin{cases} \frac{1}{\beta^{\alpha}\Gamma(\alpha)} x^{\alpha-1} e^{-\frac{x}{\beta}}, & \text{if } x \geq 0 \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

Here Γ is the gamma function given by the equation $\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}dt$. Further, α is the shape parameter and β is the scale parameter with both α and β being greater than 0. Here, α can be considered as the number of events one is waiting for and β can be considered as the mean waiting time until the first event arrives. When the value of α is fixed and the value of β increases, the PDF shifts to the right because the waiting times lengthen. Similarly, when the value of β is fixed and the value of α increases, the PDF also shifts to the right (also refer to Fig. 13 that shows the effect of changing α and β on the shape of the Gamma distribution).

The stationary Gamma process. Recall that a stochastic process describes how the value of a random variable would change through time. In a stationary Gamma process, the amount that the value of a random variable changes (within a given time interval) is drawn from a Gamma distribution. Formally, let $c(t)$ be the completion amount at the time point, t and let $\Delta c(t)$ be the change in the DM's completion amount between two successive time points t and $t + \Delta t$. Then, according to the stationary Gamma process, $\Delta c(t) \sim \mathrm{Ga}(\alpha\Delta t, \beta)$ (refer to Fig. 12(b)). That is, the amount of progress that the DM makes between time points, t and $t + \Delta t$ is drawn from a Gamma distribution with shape parameter $\alpha\Delta t$ and scale parameter β.

The stationary Gamma process has two key properties. First, given that the Gamma distribution is always positive, the stationary Gamma process is also positive and is monotonically increasing.[5] A second key property of the stationary Gamma process is that the increments that the random variable makes are time-independent. As an illustration, consider two different time periods: time period 1 consisting of the time points t_1 and $t_1 + \Delta t$, and time period 2 consisting of the time points t_2 and $t_2 + \Delta t$. Then, in a stationary Gamma process, the increments made by the random variable in these two different time periods will be drawn from the same Gamma distribution with shape parameter $\alpha\Delta t$ and scale parameter β i.e. the amount of increment only depends on the length of the time interval.

Relevance to the problem completion process. As described above, in the stationary Gamma process increments are made according

[5] This property essentially implies that the progress towards solving the problem cannot be negative as a function of time. We chose to model problem completion using a monotonically increasing process primarily for mathematical convenience, however note that it is relatively straightforward to model problem completion under a non-increasing stochastic process as well (such as a Brownian motion with drift).

to a Gamma distribution with the parameters, α and β. These two parameters have straightforward parallels to the process of problem completion. In the Gamma process, α represents the number of events that one is waiting for and β represents the mean waiting time for an event. In the context of problem completion, α can be interpreted as the 'unit' of signals of progress. Similarly, β can be interpreted as the mean waiting time for such progress signals. To provide an intuition, consider a DM solving the anagram 'crgyin'. Here, the DM would solve the anagram by making several Gamma-distributed increments, where the unit of progress i.e. α could be identifying syllables in the anagram (for example identifying '-ing') and β would be analogous to the mean waiting time for the identification of a syllable.

In addition to the above similarities, we chose the stationary Gamma process as it is well-suited to modeling a wide variety of problem completion processes, including problem solving through insight. To illustrate this, consider a Gamma process described by a Gamma distribution with parameters α and β. The standard deviation of the Gamma distribution is given as $\sigma = \sqrt{\alpha}/\beta$. When the value of σ is low, then according to the stationary Gamma process, the random variable will increment gradually over time. This makes it ideal to model solving a problem in which progress is made gradually and incrementally over time. Fig. 12(c) shows one such sample path of the stationary Gamma process when the σ value is low. On the other hand, when the value of σ is high, then the value of the random variable can undergo very small or very large increments in a time period as the standard deviation is very high. This makes it ideal to model the process of sudden jumps that can occur during problem-solving. Fig. 12(c) shows one such sample path of the stationary Gamma process when the σ value is high resulting in a little to no progress over long periods of time and then a sudden increase in the completion level in a short period of time. Further, because the process is time-independent (stationarity property), it means that these large jumps can occur anytime during problem solving making it suitable to model problem solving via insight (as the insight can potentially occur anytime during problem solving and is not necessarily dependent on the time). The stationary Gamma process also offers the benefit of being relatively simple, very well-studied, and having mathematical properties that make it easy to analyze in the context of understanding the process of problem solving.

A.1.1. Computing expected completion time

Having described the stationary Gamma process, we now detail how we compute the expected completion time of a problem. Suppose that $\mathbf{c} = [c(0), c(1), .., c(t)]$ is the observed amount of problem that the DM has completed at various time points $\mathbf{t} = [t_0, t_1, .., t_t]$. If the DM assumes that the problem completion process follows a stationary Gamma process, then the goal is to use the observed data, $\mathbf{c}$ to infer the unknown parameters α and β and then compute the expected completion time. This can be done by computing the cumulative distribution function (CDF) of the completion level as below:

$$P(c(T) \leq 1 | c(0), \dots, c(t)) = \int_0^{1-c(t)} Gamma(\alpha(T-t), \beta) P(\alpha, \beta | c(0), \dots, c(t)) d\alpha d\beta. \tag{6}$$

Here, the completion level at time T is being inferred by marginalizing out the problem solving parameters, α and β which we infer from the observed data, $\mathbf{c}$. From the above equation of the CDF, the exact probability density function (PDF) and the expected time to complete the problem can be subsequently calculated. However, solving Eq. (8) and computing the exact PDF is too complex to perform. For mathematical convenience, we solve this integral iteratively by breaking the above equation into two terms as shown below:

$$P(\alpha, \beta | \mathbf{c}) \propto P(\mathbf{c} | \alpha, \beta) P(\alpha, \beta), \tag{7}$$

$$P(c(T) < 1) = \int_0^1 Gamma(\alpha T, \beta) dT. \tag{8}$$

Eq. (9) computes the parameters α and β using Bayesian inference by utilizing the history of completion levels, $\mathbf{c}$. Eq. (10) computes the CDF of the completion level assuming that the α and β parameters are already known. To solve Eq. (9), we perform Bayesian inference using the method of conjugate priors. However, for Gamma processes, the conjugate prior distribution only considers the random effect of β and assumes that β follows a Gamma distribution, when α is given. Thus, for simplicity purposes, we infer the parameter β assuming a constant value for α. To solve Eq. (10), we use an approximation approach that uses a form of Birnbaum–Saunders (BS) distribution highlighted in Wang et al. (2015). Following this, we obtain the expected completion time as below:

$$E[T|\mathbf{c}] = \frac{E[\beta|\mathbf{c}]}{\alpha} + \frac{1}{2\alpha}. \tag{9}$$

Here, α is a constant and $E[\beta|\mathbf{c}]$ is given below:

$$E[\beta|\mathbf{c}] = \frac{\alpha_0 + \alpha_0(t_i - t_0)}{\beta_0 + c(i) - c(0)}, \tag{10}$$

where α_0 and β_0 are the prior value of the two parameters α and β, $c(i)$ is the amount of problem completed after the arrival of the ith data point, and t_i is the time at which this data point arrived. Thus, as the DM solves a problem and obtains new data points, we use Eqs. (8) and (9) to compute the expected completion time. At the arrival of a new data point, first $E[\beta|\mathbf{c}]$ is calculated using Eq. (9) and then this estimate is plugged into Eq. (8) to obtain the expected completion time.

Appendix B

B.1. Data sequence for Simulation 1

To test our model's prediction of an *Aha!* experience when a problem solved is gradually, we provide our model with a sequence of $n = 8$ data points where the completion level, $c(t)$ is drawn from a Gamma process with the parameters $\alpha = 1$ and $\beta = 1$. This results in the completion level increasingly linearly with time with a low standard deviation. To test our model's prediction of an *Aha!* experience when a problem solved is suddenly, we provide our model with a sequence of $n = 8$ data points where the completion level is drawn from a gamma process with the parameters $\alpha = 10000$ and $\beta = 0.0001$. This means that the Gamma process has a high standard deviation which results in very sudden jumps in the increase of completion level with time.

B.2. Parameter values for Simulation 3

To simulate an optimistic DM, we set the parameter values of Eqs. (8) and (9) as follows: $\alpha = 50, \alpha_0 = 50, \beta_0 = 50$. This results in the DM a priori estimating completion time to be 0. To simulate a pessimistic DM, we set the parameter values of Eqs. (8) and (9) as follows: $\alpha = 0.05, \alpha_0 = 0.05, \beta_0 = 0.05$. This results in the DM a priori estimating the completion time to be 2000.

Appendix C

In this section, we examine whether our model can be used to understand individual-level patterns of the participants in Experiment 2.

C.1. Parameter sweep analysis

To validate the gamma process model at the individual level, we computed model-predicted *Aha!* scores for each participant using their closeness rating trajectories and correlated these with self-reported *Aha!* ratings. Here, instead of fitting parameters to individual data (which faces potential overfitting issues), we performed a sweep over

Table 3
Top parameter combinations from the sweep, ranked by correlation strength. r = Pearson correlation between model-predicted and self-reported *Aha!* ratings across all participants. t = t-statistic for the difference in predicted *Aha!* scores between anagram and analytical problems.

α	α_0	β_0	r	p	t	p_t
0.1	0.1	50.0	0.166	0.0001	7.947	<.001
0.1	0.1	25.0	0.139	0.0007	6.395	<.001
1.0	5.0	50.0	0.130	0.0016	7.043	<.001
0.5	1.0	50.0	0.128	0.0019	3.981	<.001
0.1	0.5	50.0	0.126	0.0022	7.767	<.001
5.0	25.0	50.0	0.125	0.0024	6.629	<.001
10.0	50.0	50.0	0.124	0.0026	6.572	<.001
0.1	0.1	20.0	0.123	0.0027	5.905	<.001

$\alpha, \alpha_0, \beta_0 \in \{0.1, 0.5, 1, 5, 10, 20, 25, 50\}$ (yielding 512 parameter combinations) and investigated which parameter combinations best explained participant data. Critically, we assumed fixed priors across problem types (rather than tailoring prior expectations to analytical vs. anagram problems separately) to avoid the risk of explaining away individual differences purely via prior assumptions. The parameter sweep therefore reveals which values of α, α_0, and β_0 explain participant data given a common prior.

For each combination, we computed (1) the Pearson correlation between model-predicted and self-reported *Aha!* ratings across all participants, and (2) an independent samples t-test comparing predicted *Aha!* scores between analytical and anagram problems.

C.2. Results

Of the 512 parameter combinations, 6.6% yielded a significant correlation ($p < .05$) and 10.7% yielded a significant difference between predicted *Aha!* scores for analytical and anagram problems ($p < .05$). The mean correlation across all combinations was modest ($r = 0.038 \pm 0.036$), indicating that significant fits are concentrated in a specific region of parameter space rather than broadly across all combinations. Table 3 shows the top parameter combinations ranked by correlation strength.

Two consistent patterns emerge across significant parameter combinations. First, $\beta_0 \geq 25$ is required in virtually all significant fits. Importantly, this is not an arbitrary constraint — high β_0 ensures that the model's prior expected completion time is reasonable. Specifically, with $\alpha = 0.1, \alpha_0 = 0.1, \beta_0 = 50$, the prior expected completion time at $t = 0$ is ~2 steps (~40 s), which is reasonable for laboratory problem-solving tasks such as anagrams or analytical problems. Lower values of β_0 produce implausibly pessimistic priors (e.g., $\beta_0 = 1$ yields an expected completion time of ~100 steps).

Second, significant fits consistently satisfy $\alpha \leq \alpha_0$, meaning that prior strength dominates over the time-elapsed term in the update equation (Eq. (10)). This implies that model updates are driven primarily by *progress made* (the denominator term $c(i) - c(0)$) rather than *time elapsed* (the numerator term $\alpha(t_i - t_0)$). In other words, how close one feels to solution matters more than how long one has been working, which is psychologically interpretable.

C.3. Individual-level correlation

Using the best-fitting parameters ($\alpha = 0.1, \alpha_0 = 0.1, \beta_0 = 50$), the model significantly predicted individual self-reported *Aha!* ratings (combined: $r = 0.166$, $p < .001$), and correctly predicted higher *Aha!* scores for anagram than analytical problems ($t = 7.95$, $p < .001$). Although the correlation is modest, it is achieved without fitting any parameters to individual data, making it a conservative estimate of the model's predictive validity. The concentration of significant fits within a theoretically motivated parameter region further supports the interpretability of the model constraints identified above.

Table 4
Norming data for the 30 anagrams. The 5 anagrams used in Exp 2 are highlighted.

Anagram	Difficulty	RT	% solved
AHSDOW	1.56	11.75	0.98
IEGWHT	1.59	14.80	0.95
ERLKC	1.69	17.73	0.96
AALRYS	1.95	26.01	0.83
SAOSI	2.18	21.83	0.94
TINAG	2.25	30.34	0.77
DONRO	2.32	32.40	0.72
NAYTUGH	2.46	44.32	0.83
HAKMOCM	2.49	30.88	0.90
AYKWALJ	2.50	22.81	0.88
AXPLINE	2.51	33.22	0.78
ALIPNXE	2.62	35.75	0.81
EGANV	2.68	29.08	0.82
HKAMOMC	2.81	28.38	0.90
EURADBL	3.03	41.10	0.69
NCRBOA	3.06	56.35	0.64
RDUNEF	3.08	58.97	0.72
NITGA	3.17	61.70	0.65
PMUOI	3.25	60.03	0.58
MATCILE	3.28	66.35	0.71
AYKAWLJ	3.39	54.67	0.70
UERADBL	3.42	49.96	0.62
NCBROA	3.45	76.36	0.62
LUBMEJD	3.46	52.44	0.80
EGNAV	3.48	52.85	0.75
MARNOD	3.53	56.47	0.71
HSOAWD	3.73	80.44	0.66
CRLAIGE	3.76	76.95	0.57
DRUNEF	3.81	64.43	0.83
ANYTUGH	3.94	66.63	0.63

C.4. Robustness of aggregate-level analysis

For the averaged data analysis (Model Simulation, Experiment 2 in the main text), we used $\alpha = \alpha_0 = \beta_0 = 50$. In contrast to the individual-level analysis, we found that the qualitative result at the aggregate level i.e., anagram problems produce higher predicted *Aha!* scores than analytical problems, is robust across various parameter combinations tested. This robustness arises because the signal in the averaged trajectories is strong enough that the model correctly captures the qualitative difference between problem types regardless of prior assumptions. The sensitivity to parameters observed at the individual level therefore reflects the inherent noisiness of single-participant closeness trajectories.

Appendix D

We now provide details about the norming study conducted to select the stimuli for Experiment 2.

D.1. Participants

We recruited 850 participants from Prolific with the goal of obtaining at least 50 ratings for each problem in our dataset. Participants earned $1.60 for participation in a study that took around 6–7 minutes to complete.

D.2. Stimuli

The stimuli consisted of 30 anagrams and 30 analytical word problems, where the 30 anagrams were taken from the 100 anagram dataset in Experiment 1 and the 30 analytical problems were adapted from previous studies on problem-solving (Gilhooly & Murphy, 2005; Laukkonen et al., 2021; Metcalfe & Wiebe, 1987; Schooler et al., 1993; Webb et al., 2018).

Table 5
Norming data for the 30 analytical problems. The 5 problems used in Exp 2 are highlighted.

Analytical word problem	Difficulty	RT	% solved
Lana has 2 bags with 2 marbles in each bag. Markus has 2 bags with 3 marbles in each bag. How many more marbles does Markus have?	1.32	24.16	0.93
Lilah's band had practiced 24 songs. At a performance, they played 7 songs in their first set. In their second set, they played 8 songs. How many songs did they play for their third and final set?	1.77	30.76	0.98
Carrie grew three inches taller last year, and five inches taller this year. How many inches taller did she grow in the last two years?	1.80	25.59	0.91
Mary won't eat fish or spinach, Sally won't eat fish or green beans, Steve won't eat shrimp or potatoes, Alice won't eat beef or tomatoes, and Jim won't eat fish or tomatoes. If you are willing to host them for a small party, which item from the following list can you serve: green beans, creamed codfish, tomatoes, celery, and roast beef?	1.91	54.92	0.93
A farmer has 19 sheep on his land. One day, a big storm hits, and seven sheep run away. The next day, four sheep return and find their way home. The next day, another big storm hits, and six sheep run away. How many sheep does the farmer have left?	1.97	37.64	0.95
Lauren's chicken laid an average of six eggs per week. Lauren sold those eggs for $3 per dozen. How much money did she collect in four weeks if she sold all her eggs?	2.07	43.84	0.93
Ricky has a magic money box. Every day the box doubles the number of coins placed inside of it. Ricky put in 3 pennies on Monday. He looked inside his box on Friday night. How many pennies did Ricky see?	2.27	44.71	0.83
The book store is very busy today. There are 25 children listening to a story. 35 people are shopping for books. 18 people are at the checkout counter. How many people are at the bookstore?	2.30	35.12	0.83
Before Gary injured his arm, he was able to type 9 words per minute on his phone. After he injured his arm, he had to use his left arm for a while, and he could only type 6 words per minute on his phone. What is the difference between the number of words he could type in 5 min before and after he broke his arm?	2.33	46.42	0.84
The book store is very busy today. There are 25 children listening to a story narrated by 1 person. 35 people are shopping for books. 20 people are studying. 18 people are at the checkout counter. There are 5 cashiers at the counter. How many people are at the bookstore in total?	2.36	65.76	0.78
Before Gary injured his arm, he was able to type 9 words per minute on his phone. After he injured his arm, he had to use his left arm for a while, and he could only type 6 words per minute on his phone. What is the difference between the number of words he could type in 8 min before and after he broke his arm?	2.38	57.46	0.81
Gary had 100 dollars. He went to the grocery store and bought three pints of ice cream for 8 dollars each. Then, he went to the farmer's market and bought four dozen eggs for 5 dollars per dozen. Finally, he went to a book store and bought three books for 3 dollars each. How much money does he have left?	2.40	60.97	0.86
Ben spent $40 for shoes. This was $14 more than what he spent for a shirt. The tie was $10 cheaper than the shirt. How much was the tie?	2.50	42.77	0.88
Next week I am going to have lunch with my friend, visit the new art gallery, go to the Social Security office, and have my teeth checked at the dentist. My friend cannot meet me on Wednesday; the Social Security office is closed weekends; the dentist has office hours only on Tuesday, Friday, and Saturday; the art gallery is closed Tuesday, Thursday, and weekends. On what single day can I do everything I have planned?	2.57	73.03	0.82
Eric is trying to get in shape. He wants to do this by climbing stairs. He starts on the fourth floor, climbs up five stories, down seven, up six, down three, up four, and then up two again. What floor is he on now?	2.76	51.92	0.87
Given a source of unlimited water and four containers of different capacities — 14, 20, 24, and 12 liters — how would you obtain exactly 50 liters of water?	2.84	85.74	0.92
Ben spent $42 for shoes. This was $14 less than what he spent for a shirt and twice more expensive than the tie. The tie was $20 cheaper than the jeans. How much was the jeans?	2.86	74.85	0.59
Lebrun, Lenoir, and Leblanc are, not necessarily in that order, the accountant, warehouseman, and traveling salesman of a firm. The salesman is unmarried. Both Lebrun and Lenoir are married. Lebrun is not the accountant. What job does Lenoir have? Accountant, warehouseman, or salesman?	3.04	79.72	0.74
Erin had 24 marbles and Evan had 3 marbles. Erin gave some of her marbles to Evan. Now Erin has exactly double the number of marbles that Evan has. How many marbles did Erin give to Evan?	3.06	82.50	0.65
Lauren had nine chickens. Each chicken laid an average of six eggs per week. Lauren sold those eggs for $3 per dozen. How much money did she collect in four weeks if she sold all her eggs?	3.10	88.31	0.55
Eric is trying to get in shape. He wants to do this by climbing stairs. He starts on the third floor, climbs up five stories, down seven, up six, down three, up four, down four, up five, down three, and then up two again. What floor is he on now?	3.15	61.35	0.83
Erin had 40 marbles, Ben had 8 marbles, and Chris had 3 marbles. Erin gave 16 marbles to Ben. Ben then gave some of his marbles to Chris. Now Ben has exactly double the number of marbles that Chris has. How many marbles did Ben give to Chris?	3.50	94.64	0.65

(*continued on next page*)

Table 5 (*continued*).

Andrew is twice the age of Brian, Charles is three times older than Brian, and the sum of their ages is 60 years. How old is Charles?	3.50	69.40	0.81
Carrie grew three inches taller last year, and five inches taller this year. Lana was four inches taller than Carrie two years ago. Lana then grew five inches taller last year, and seven inches taller this year. How many inches taller is Lana compared to Carrie now?	3.58	104.92	0.53
What day follows the day before yesterday if two days from now will be Sunday?	3.82	61.75	0.56
Lebrun, Lenoir, Ledroi, and Leblanc are, not necessarily in that order, the accountant, cashier, warehouseman, and salesman of a firm. The salesman has never been married. Lebrun, Lenoir, and Ledroi are married. Ledroi is the cashier. Lebrun is not the accountant. What job does Lenoir have? Accountant, cashier, warehouseman, or salesman?	3.88	135.63	0.72
Three spies, suspected as double agents, speak as follows when questioned: Albert: “Bertie is a mole.” Bertie: “Cedric is a mole.” Cedric: “Bertie is lying.” Assuming that moles lie, other agents tell the truth, and there is just one mole among the three, who is the mole? Albert, Bertie, or Cedric?	4.23	91.93	0.90
Andrew is half the age of Brian, Brian is three times older than Charles and the sum of their ages is 44 years. How old is Charles?	4.42	176.68	0.48
Smith is a butcher and president of the street storekeepers’ committee, which also includes the grocer, the baker, and the pharmacist. They all sit around a table. Smith sits on Jones’ left. Molly sits at the grocer’s right. Bailey, who faces Jones, is not the pharmacist. Who is Molly? The butcher, grocer, baker, or the pharmacist?	4.77	167.93	0.62
The police were convinced that either A, B, C, or D had committed a crime. Each of the suspects, in turn, made a statement.	5.29	92.90	0.69

D.3. Procedure

Participants were given four problems to solve one after another (two randomly selected anagrams and two randomly selected analytical problems). For each problem, participants were displayed the problem text and were provided a text box below the problem description. They were asked to provide their answer to the problem in the text box, which would then allow them to move to the next screen. At any time, if a participant decided they were unable to solve the given problem, they could move to the next screen by simply typing ‘next’ in the provided text box. After participants had typed their answer (or typed ‘next’), they were asked to provide their rating of *difficulty* for the previously shown problem (on a scale of 1–5). Note that participants were given unlimited time to solve a problem and were paid regardless of their performance.

D.4. Results

Tables 4 and 5 show the mean difficulty, mean response time, and the percentage of participants that correctly solved the different anagrams and analytical problems. Our collected dataset was quite varied as we obtained a range of difficulty ratings. We then selected five anagrams and five analytical problems for Exp 2 (highlighted in gray in the respective tables). These problems were equivalent in difficulty ($M = 3.2, SD = 0.18$), with a one-way ANOVA revealing that the ten problem were not significantly different from each other, $F(9397) = 0.59, MSE = 1.68, p = 0.81$. Further, the five anagrams were similar in difficulty to the five analytical problems, $t(630) = 0.42, p = 0.7$. Note that the five analytical problems had a slightly larger response time than the five anagrams. In general, we found that anagrams tended to have lower response times than the analytical problems of equivalent difficulty. However, this was not an issue for Experiment 2 since our goal was to address the potential confound of items with *larger* response time driving *Aha!* moments. That is, this is in the opposite direction of the potential confound which we aim to address.

Appendix E

In this section, we report a replication of Experiment 2 using classic insight problems. This experiment allows us to test the generality of our theory and evaluate whether our framework can be applied to problems outside of anagrams and analytical word problems.

Table 6
Insight problems used for the experiment.

Insight problem
Mr. Hardy slipped and fell off a sixty-foot ladder onto the concrete floor below. However, he did not injure himself in any way. How is this possible?
Joe Fan has no psychic powers but he can tell you the score of any football game before it starts. How?
Our basketball team won 72–49, and yet not one man scored as much as a single point. How is that possible?
A murderer is condemned to death. He has to choose among three rooms. The first is full of raging fires, the second is full of assassins with loaded guns, and the third is full of lions that have not eaten in 3 years. Which room is safest for him?
Alex and Casey are blood relatives of Bobbie. However, Alex and Casey are not blood relatives at all. How is this possible?
A man is reading a book when the lights go off but even although the room is pitch dark the man goes on reading. How?

E.1. Participants

We recruited 298 participants from Prolific. Participants earned \$1.00 for participation in a study that took around 3–4 min to complete.

E.2. Stimuli

We used six popular insight problems, shown in Table 6.

E.3. Procedure

The procedure was similar to Experiment 2 except that participants were given two randomly selected insight problems to solve (instead of one anagram and one analytical problem). As before, participants were prompted every ten seconds and were asked to provide their ratings of closeness on a scale of 1 to 100 (‘How close are you towards solving this problem?’). This prompt was shown every ten seconds until the participants provided their respective answer in the provided text box. After participants provided their answer, they were asked to provide their subjective rating of an *Aha!* moment on a scale of 1–7 for the previous problem.

E.4. Results

In Experiment 2, as reported in the main text, we had excluded the problems for which participants had provided the incorrect answer.

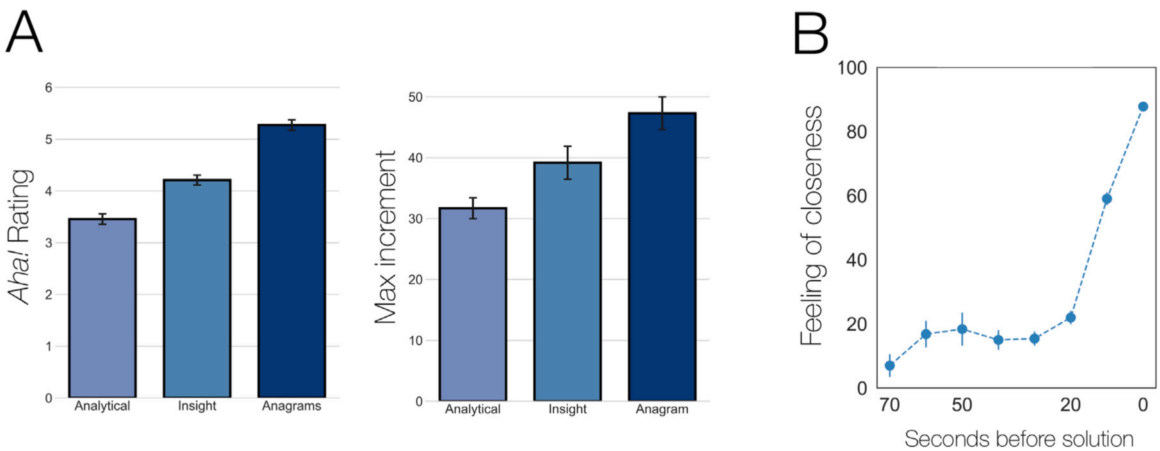

Fig. 14. Prediction errors during the meta-cognitive monitoring process predict *Aha!* ratings for insight problems as well. (A) Left panel: The y-axis plots the mean *Aha!* rating for the analytical problems and the anagrams. Insight problems have larger *Aha!* than analytical problems but lower than anagrams. Right panel: The mean *maximum* increment in the closeness ratings during the course of problem solving. The maximum increment made while solving insights is higher than that of analytical problems but lower than anagrams. (B) Participants' mean closeness ratings during the course of problem solving for insight problems. Similar to anagrams, insight problems are solved suddenly.

However, insight problems elicited a more diverse range of solutions, many of which were potentially correct. Thus, for this analysis, we used a different exclusion criteria. For all participants, we only included those problems for which they provided a warmth rating of at least 75 on the last trial i.e., we included those datapoints for which participants' confidence in the final trial was relatively high. Our final analysis was performed over 410 datapoints (out of the total 596 datapoints).

First, as shown in Fig. 14(a), insight problems generated higher *Aha!* compared to the analytical problems, $t(721) = 5.21, p < 0.01$, but lower than the anagrams, $t(739) = 7.48, p < 0.01$. This finding is consistent with prior research (Webb et al., 2018). Next, similar to anagrams, we see that closeness ratings during the course of problem solving show a sudden jump for insight problems (Fig. 14(b)). Importantly, in line with our theory, the maximum increment in the closeness rating while solving insight problems was greater than that of analytical problems, $t(721) = 2.17, p < 0.05$, but lower than that of anagrams, $t(739) = 2.10, p < 0.05$ (Fig. 14(a)). Thus, metacognitive prediction, as measured by the maximum "jump" during problem solving, is a significant predictor of *Aha!* rating for insight problems as well.

Data availability

We will be sharing data upon acceptance of the publication. We have deanonymized our code base and database to share with reviewers.

References

Ackerman, Rakefet, & Thompson, Valerie A. (2017). Meta-reasoning: Monitoring and control of thinking and reasoning. *Trends in COgnitive SCiences*, *21*(8), 607–617.

Aguilar-Lleyda, David, Lemarchand, Maxime, & de Gardelle, Vincent (2020). Confidence as a priority signal. *Psychological Science*, *31*(9), 1084–1096.

Ammalainen, Artur, & Moroshkina, Nadezhda (2020). The effect of true and false unreportable hints on anagram problem solving, restructuring, and the aha!-experience. *Journal of Cognitive Psychology*, *32*, 1–15.

Aziz-Zadeh, Lisa, Kaplan, Jonas T., & Iacoboni, Marco (2009). "Aha!": The neural correlates of verbal insight solutions. *Human Brain Mapping*, *30*(3), 908–916.

Baranes, Adrien F., Oudeyer, Pierre-Yves, & Gottlieb, Jacqueline (2014). The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration. *Frontiers in Neuroscience*, *8*, 317.

Bayer, Hannah M., & Glimcher, Paul W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, *47*(1), 129–141.

Becker, Maxi, & Cabeza, Roberto (2025). The neural basis of the insight memory advantage. *Trends in Cognitive Sciences*, *29*(3), 255–268.

Becker, Maxi, Sommer, Tobias, & Cabeza, Roberto (2025). Insight predicts subsequent memory via cortical representational change and hippocampal activity. *Nature Communications*, *16*(1), 4341.

Becker, Maxi, Wang, Xinhao, & Cabeza, Roberto (2024). Surprise!—Clarifying the link between insight and prediction error. *Psychonomic Bulletin & Review*, *31*(6), 2714–2723.

Berlyne, Daniel Ellis (1954). A theory of human curiosity. *British Psychological Society*, *45*, 180–191.

Bowden, Edward M., & Jung-Beeman, M. (2003). Aha! insight experience correlates with solution activation in the right hemisphere. *"Psychonomic Bulletin & Review"*, *10*(3), 730–737.

Bowden, Edward M., & Jung-Beeman, Mark (2007). Methods for investigating the neural components of insight. *Methods*, *42*(1), 87–99.

Buehler, Roger, Griffin, Dale, & Ross, Michael (1994). Exploring the "planning fallacy": Why people underestimate their task completion times. *Journal of Personality and Social Psychology*, *67*(3), 366.

Butler, Andrew C., Fazio, Lisa K., & Marsh, Elizabeth J. (2011). The hypercorrection effect persists over a week, but high-confidence errors return. *Psychonomic Bulletin & Review*, *18*(6), 1238–1244.

Butler, Andrew C., Karpicke, Jeffrey D., & Roediger III, Henry L. (2008). Correcting a metacognitive error: feedback increases retention of low-confidence correct responses. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *34*(4), 918.

Butterfield, Brady, & Metcalfe, Janet (2001). Errors committed with high confidence are hypercorrected. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *27*(6), 1491.

Caprioli, Sara, Fuchs, Christoph, & Van den Bergh, Bram (2023). On breaking functional fixedness: How the aha! moment enhances perceived product creativity and product appeal. *Journal of Consumer Research*, *50*(1), 48–69.

Chesebrough, Christine, Chrysikou, Evangelia G, Holyoak, Keith J, Zhang, Fengqing, & Kounios, John (2023). Conceptual change induced by analogical reasoning sparks aha moments. *Creativity Research Journal*, *35*(3), 499–521.

Chu, Yun, Dewald, Andrew D., & Chronicle, Edward P. (2007). Theory driven hints in the cheap necklace problem: A preliminary investigation. *The Journal of Problem Solving*, *1*(2), 4.

Chuderski, Adam, & Jastrzębski, Jan (2018). Much ado about aha!: Insight problem solving is strongly related to working memory capacity and reasoning ability. *Journal of Experimental Psychology: General*, *147*(2), 257.

Cushen, Patrick J., & Wiley, Jennifer (2011). Aha! voila! eureka! bilingualism and insightful problem solving. *Learning and Individual Differences*, *21*(4), 458–462.

Cushman, Fiery (2020). Rationalization is rational. *Behavioral and Brain Sciences*, *43*, E28.

Danek, Amory H, Fraps, Thomas, von Müller, Albrecht, Grothe, Benedikt, & Öllinger, Michael (2014). It's a kind of magic—what self-reports can reveal about the phenomenology of insight problem solving. *Frontiers in Psychology*, *5*, 1408.

Danek, Amory H., & Wiley, Jennifer (2017). What about false insights? Deconstructing the aha! experience along its multiple dimensions for correct and incorrect solutions separately. *Frontiers in Psychology*, *7*, 2077.

Danek, Amory H., Wiley, Jennifer, & Öllinger, Michael (2016). Solving classical insight problems without aha! experience: 9 dot, 8 coin, and matchstick arithmetic problems. *The Journal of Problem Solving*, *9*(1), 4.

Danek, Amory H., Williams, Joshua, & Wiley, Jennifer (2018). Closing the gap: connecting sudden representational change to the subjective aha! experience in insightful problem solving. *Psychological Research*, 1–9.

Daw, Nathaniel D., & Tobler, Philippe N. (2014). Value learning through reinforcement: the basics of dopamine and reinforcement learning. In *Neuroeconomics* (pp. 283–298). San Diego: Elsevier.

de Jonge, Bram, Teunter, Ruud, & Tinga, Tiedo (2017). The influence of practical factors on the benefits of condition-based maintenance over time-based maintenance. *Reliability Engineering & System Safety*, *158*, 21–30.

DeCaro, Marci S., Wieth, Mareike B., et al. (2016). When higher working memory capacity hinders insight. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *42*(1), 39.

Dubey, Rachit, & Griffiths, Thomas L. (2020). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, *127*(3), 455–476.

Duncker, Karl, & Lees, Lynne S. (1945). On problem-solving.. *Psychological Monographs*, *58*(5).

Eggensperger, Katharina, Lindauer, Marius, & Hutter, Frank (2017). Neural networks for predicting algorithm runtime distributions. arXiv preprint arXiv:1709.07615.

Fleck, Jessica I., & Weisberg, Robert W. (2013). Insight versus analysis: Evidence for diverse methods in problem solving. *Journal of Cognitive Psychology*, *25*(4), 436–463.

Fleming, Stephen M., & Daw, Nathaniel D. (2017). Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review*, *124*(1), 91–114.

Fleming, Stephen M., & Dolan, Raymond J. (2012). The neural basis of metacognitive ability. *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)*, *367*(1594), 1338–1349.

Fox, Emily, & Riconscente, Michelle (2008). Metacognition and self-regulation in james, piaget, and vygotsky. *Educational Psychology Review*, *20*(4), 373–389.

Friedlander, Kathryn J., & Fine, Philip A. (2018). "The penny drops": Investigating insight through the medium of cryptic crosswords. *Frontiers in Psychology*, *9*, 904.

Frömer, Romy, Nassar, Matthew R, Bruckner, Rasmus, Stürmer, Birgit, Sommer, Werner, & Yeung, Nick (2021). Response-based outcome predictions and confidence regulate feedback processing and learning. *Elife*, *10*.

Gick, Mary L., & Lockhart, Robert S. (1995). Cognitive and affective components of insight. In *The nature of insight* (pp. 197–228). Cambridge, MA: The MIT Press.

Gilhooly, Kenneth J. (2016). Incubation and intuition in creative problem solving. *Frontiers in Psychology*, *7*, 1076.

Gilhooly, Ken J., Georgiou, George, & Devery, Ultan (2013). Incubation and creativity: Do something different. *Thinking & Reasoning*, *19*(2), 137–149.

Gilhooly, Kenneth J., & Murphy, P. (2005). Differentiating insight from non-insight problems. *Thinking & Reasoning*, *11*(3), 279–302.

Glimcher, Paul W. (2011). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, *108*(3), 15647–15654.

Gomes, Carla P., & Selman, Bart (2001). Algorithm portfolios. *Artificial Intelligence*, *126*(1–2), 43–62.

Griffiths, Thomas L., Lieder, Falk, & Goodman, Noah D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, *7*(2), 217–229.

Gruaz, Lucas, Modirshanechi, Alireza, Becker, Sophia, & Brea, Johanni (2025). Merits of curiosity: a simulation study. *Open Mind*, *9*, 1037–1065.

Gruber, Howard E. (1995). Insight and affect in the history of science. In *The nature of insight* (pp. 397–431). Cambridge, MA: The MIT Press.

Hamermesh, Daniel S., & Lee, Jungmin (2007). Stressed out on four continents: Time crunch or yuppie kvetch? *The Review of Economics and Statistics*, *89*(2), 374–383.

Hedne, Mikael R., Norman, Elisabeth, & Metcalfe, Janet (2016). Intuitive feelings of warmth and confidence in insight and noninsight problem solving of magic tricks. *Frontiers in Psychology*, *7*, 1314.

Hélie, Sebastien, & Sun, Ron (2010). Incubation, insight, and creative problem solving: a unified theory and a connectionist model. *Psychological Review*, *117*(3), 994.

Hollerman, Jeffrey R., & Schultz, Wolfram (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, *1*(4), 304–309.

Hutter, Frank, Xu, Lin, Hoos, Holger H., & Leyton-Brown, Kevin (2014). Algorithm runtime prediction: Methods & evaluation. *Artificial Intelligence*, *206*, 79–111.

Inbar, Yoel, Botti, Simona, & Hanko, Karlene (2011). Decision speed and choice regret: When haste feels like waste. *Journal of Experimental Social Psychology*, *47*(3), 533–540.

Ishikawa, Tetsuo, Toshima, Mayumi, & Mogi, Ken (2019). How and when? Metacognition and solution timing characterize an "aha" experience of object recognition in hidden figures. *Frontiers in Psychology*, *10*, 1023.

Jones, David (2012). *The Aha! Moment: A Scientist's Take on Creativity*. Baltimore,MD: JHU Press.

Kaplan, Craig A., & Simon, Herbert A. (1990). In search of insight. *Cognitive Psychology*, *22*(3), 374–419.

Knoblich, Günther, Ohlsson, Stellan, Haider, Hilde, & Rhenius, Detlef (1999). Constraint relaxation and chunk decomposition in insight problem solving. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *25*(6), 1534–1555.

Knoblich, Günther, Ohlsson, Stellan, & Raney, Gary E. (2001). An eye movement study of insight problem solving. *Memory & Cognition*, *29*(7), 1000–1009.

Kohler, Wolfgang (2015). *The task of Gestalt psychology*. Princeton University Press.

Kounios, John, & Beeman, Mark (2014). The cognitive neuroscience of insight. *Annual Review of Psychology*, *65*(1), 71–93.

Kounios, John, Frymiare, Jennifer L, Bowden, Edward M, Fleck, Jessica I, Subramaniam, Karuna, Parrish, Todd B, & Jung-Beeman, Mark (2006). The prepared mind: Neural activity prior to problem presentation predicts subsequent solution by sudden insight. *Psychological Science*, *17*(10), 882–890.

Kruger, Justin, & Evans, Matt (2004). If you don't want to be late, enumerate: Unpacking reduces the planning fallacy. *Journal of Experimental Social Psychology*, *40*(5), 586–598.

Kulhavy, Raymond W., Yekovich, Frank R., & Dyer, James W. (1976). Feedback and response confidence. *Journal of Educational Psychology*, *68*(5), 522.

Ladd, Gary W., & Price, Joseph M. (1986). Promoting children's cognitive and social competence: The relation between parents' perceptions of task difficulty and children's perceived and actual competence. *Child Development*, 446–460.

Langley, Pat, & Jones, Randolph (1988). A computational model of scientific insight. *The Nature of Creativity: Contemporary Psychological Perspectives*, *177*, 177–201.

Laukkonen, Ruben E. (2019). *The phenomenology of truth: The psychological functions of the insight experience*. (Doctoral Dissertation).

Laukkonen, Ruben E, Ingledew, Daniel J, Grimmer, Hilary J, Schooler, Jonathan W, & Tangen, Jason M (2021). Getting a grip on insight: real-time and embodied aha experiences predict correct solutions. *Cognition and Emotion*, *35*(5), 918–935.

Laukkonen, Ruben E, Kaveladze, Benjamin T, Tangen, Jason M, & Schooler, Jonathan W (2020). The dark side of eureka: Artificially induced aha moments make facts feel true. *Cognition*, *196*, 104–122.

Laukkonen, Ruben E., & Tangen, Jason M. (2017). Can observing a necker cube make you more insightful? *Consciousness and Cognition*, *48*, 198–211.

Laukkonen, Ruben E, Webb, Margaret E, Salvi, Carola, Tangen, Jason M, & Schooler, Jonathan (2018). Eureka heuristic: How feelings of insight signal the quality of a new idea. PsyArXiv Preprint.

Lauterman, Tirza, & Ackerman, Rakefet (2019). Initial judgment of solvability in non-verbal problems–a predictor of solving processes. *Metacognition and Learning*, *14*(3), 365–383.

Lieder, Falk, & Griffiths, Thomas L. (2020). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, E1.

Liljedahl, Peter G. (2005). Mathematical discovery and affect: The effect of aha! experiences on undergraduate mathematics students. *International Journal of Mathematical Education in Science and Technology*, *36*(2–3), 219–234.

Litman, Jordan (2005). Curiosity and the pleasures of learning: Wanting and liking new information. *Cognition & Emotion*, *19*(6), 793–814.

Loewenstein, George (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, *116*(1), 75–98.

Lucca, Kelsey, Horton, Rachel, & Sommerville, Jessica A. (2020). Infants rationally decide when and how to deploy effort. *Nature Human Behaviour*, *4*(4), 372–379.

MacGregor, James N., Ormerod, Thomas C., & Chronicle, Edward P. (2001). Information processing and insight: a process model of performance on the nine-dot and related problems. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *27*(1), 176.

Maier, Norman R. F. (1931). Reasoning and learning. *Psychological Review*, *38*(4), 332–346.

Marr, David (1982). Vision: A computational approach. San Francisco, CA: Freeman.

Marvin, Caroline B., & Shohamy, Daphna (2016). Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *Journal of Experimental Psychology: General*, *145*(3), 266.

Mayer, R. E. (1992). *Thinking, Problem Solving, Cognition*. New York, NY: W. H. Freeman.

Menkel-Meadow, Carrie (2001). Aha–is creativity possible in problem solving and teachable in legal education. *Harv. Negot. L. Rev.*, *6*, 97–144.

Mercier, Marius, Garsmeur, Alexis, & Mercier, Hugo (2025). The appeal of insight: Why riddles and whodunits captivate us. *Psychology of Aesthetics, Creativity, and the Arts*.

Metcalfe, Janet (1986a). Feeling of knowing in memory and problem solving.. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *12*(2), 288–294.

Metcalfe, Janet (1986b). Premonitions of insight predict impending error. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *12*(4), 623–634.

Metcalfe, Janet, & Greene, Matthew Jason (2007). Metacognition of agency. *Journal of Experimental Psychology: General*, *136*(2), 184–199.

Metcalfe, Janet, & Kornell, Nate (2005). A region of proximal learning model of study time allocation. *Journal of Memory and Language*, *52*(4), 463–477.

Metcalfe, Janet, & Wiebe, David (1987). Intuition in insight and noninsight problem solving. *Memory & Cognition*, *15*(3), 238–246.

Modirshanechi, Alireza, Brea, Johanni, & Gerstner, Wulfram (2022). A taxonomy of surprise definitions. *Journal of Mathematical Psychology*, *110*, Article 102712.

Moroshkina, Nadezhda V, Pavliuchik, Elena I, Ammalainen, Artur V, Gershkovich, Valeria A, & Lvova, Olga V (2024). The aha! experience is associated with a drop in the perceived difficulty of the problem. *Frontiers in Psychology*, *15*, Article 1314531.

Moroshkina, Nadezhda V, Savina, Alina I, Ammalainen, Artur V, Gershkovich, Valeria A, Zverev, Ilia V, & Lvova, Olga V (2022). How difficult was it? Metacognitive judgments about problems and their solutions after the aha moment. *Frontiers in Psychology*, *13*, Article 911904.

Niv, Yael, & Schoenbaum, Geoffrey (2008). Dialogues on prediction errors. *Trends in Cognitive Sciences*, *12*(7), 265–272.

Novick, Laura R., & Sherman, Steven J. (2003). On the nature of insight solutions: Evidence from skill differences in anagram solution. *The Quarterly Journal of Experimental Psychology Section A*, *56*(2), 351–382.

Ohlsson, Stellan (1984). Restructuring revisited: II. An information processing theory of restructuring and insight. *Scandinavian Journal of Psychology*, *25*(2), 117–129.

Öllinger, Michael, Jones, Gary, & Knoblich, Günther (2008). Investigating the effect of mental set on insight problem solving. *Experimental Psychology*, *55*(4), 269–282.

Öllinger, Michael, Jones, Gary, & Knoblich, Günther (2014). The dynamics of search, impasse, and representational change provide a coherent explanation of difficulty in the nine-dot problem. *Psychological Research*, *78*(2), 266–275.

Paulus, Markus, Tsalas, Nike, Proust, Joelle, & Sodian, Beate (2014). Metacognitive monitoring of oneself and others: Developmental changes during childhood and adolescence. *Journal of Experimental Child Psychology*, *122*, 153–165.

Pessiglione, Mathias, Seymour, Ben, Flandin, Guillaume, Dolan, Raymond J, & Frith, Chris D (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*(7106), 1042–1045.

Rescorla, Robert A., & Wagner, Allan R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Current Research and Theory*, 64–99.

Rouault, Marion, & Fleming, Stephen M. (2020). Formation of global self-beliefs in the human brain. *Proceedings of the National Academy of Sciences*, *117*(44), 27268–27276.

Rutledge, Robb B, Skandali, Nikolina, Dayan, Peter, & Dolan, Raymond J (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences*, *111*(33), 12252–12257.

Sackett, Aaron M, Meyvis, Tom, Nelson, Leif D, Converse, Benjamin A, & Sackett, Anna L (2010). You're having fun when time flies: The hedonic consequences of subjective time progression. *Psychological Science*, *21*(1), 111–117.

Salvi, Carola, Bricolo, Emanuela, Kounios, John, Bowden, Edward, & Beeman, Mark (2016). Insight solutions are correct more often than analytic solutions. *Thinking & Reasoning*, *22*(4), 443–460.

Sanna, Lawrence J, Parks, Craig D, Chang, Edward C, & Carter, Seth E (2005). The hourglass is half full or half empty: Temporal framing and the group planning fallacy. *Group Dynamics: Theory, Research, and Practice*, *9*(3), 173–188.

Savinova, Anna, & Korovkin, Sergei (2022). Surprise! why insightful solution is pleasurable. *Journal of Intelligence*, *10*(4), 98.

Schooler, Jonathan W., Ohlsson, Stellan, & Brooks, Kevin (1993). Thoughts beyond words: When language overshadows insight. *Journal of Experimental Psychology: General*, *122*(2), 166.

Schultz, Wolfram (1997). Dopamine neurons and their role in reward mechanisms. *Current Opinion in Neurobiology*, *7*(2), 191–197.

Schultz, Wolfram (2016). Dopamine reward prediction error coding. *Dialogues in Clinical Neuroscience*, *18*(1), 23–32.

Schultz, Wolfram, Dayan, Peter, & Montague, P. Read (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599.

Schultz, Wolfram, & Dickinson, Anthony (2000). Neuronal coding of prediction errors. *Annual Review of Neuroscience*, *23*(1), 473–500.

Schulz, Lion, Fleming, Stephen M., & Dayan, Peter (2021). Metacognitive computations for information search: Confidence in control. *BioRxiv*, http://dx.doi.org/10.1101/2021.03.01.433342.

Shen, Wangbing, Tong, Yu, Yuan, Yuan, Zhan, Huijia, Liu, Chang, Luo, Jing, & Cai, Houde (2018). Feeling the insight: uncovering somatic markers of the "aha" experience. *Applied Psychophysiology and Biofeedback*, *43*(1), 13–21.

Shen, Wangbing, Yuan, Yuan, Liu, Chang, & Luo, Jing (2016). In search of the 'aha!'experience: Elucidating the emotionality of insight problem-solving. *British Journal of Psychology*, *107*(2), 281–298.

Simon, Herbert A., & Newell, Allen (1971). Human problem solving: The state of the theory in 1970. *American Psychologist*, *26*(2), 145.

Singpurwalla, Nozer D. (1995). Survival in dynamic environments. *Statistical Science*, *10*, 86–103.

Skaar, Øystein O., & Reber, Rolf (2020). The phenomenology of aha-experiences. *Motivation Science*, *6*(1), 49.

Son, Lisa K., & Metcalfe, Janet (2000). Metacognitive and control strategies in study-time allocation. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *26*(1), 204–221.

Sternberg, Robert J., & Davidson, Janet E. (1995). *The nature of insight*. The MIT Press.

Sutton, Richard S., & Barto, Andrew G. (2018). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Thagard, Paul, & Stewart, Terrence C. (2011). The aha! experience: Creativity through emergent binding in neural networks. *Cognitive Science*, *35*(1), 1–33.

Topolinski, Sascha, Bakhtiari, Giti, & Erle, Thorsten M. (2016). Can I cut the gordian tnok? The impact of pronounceability, actual solvability, and length on intuitive problem assessments of anagrams. *Cognition*, *146*, 439–452.

Topolinski, Sascha, & Reber, Rolf (2010). Gaining insight into the "aha" experience. *Current Directions in Psychological Science*, *19*(6), 402–405.

Undorf, Monika, & Ackerman, Rakefet (2017). The puzzle of study time allocation for the most challenging items. *Psychonomic Bulletin & Review*, *24*(6), 2003–2011.

Valueva, Ekaterina, Lapteva, Ekaterina, & Ushakov, Dmitry (2016). Aha-cueing in problem solving. *Learning and Individual Differences*, *52*, 204–208.

van Noortwijk, Jan M, van der Weide, Johannes AM, Kallen, Maarten-Jan, & Pandey, Mahesh D (2007). Gamma processes and peaks-over-threshold distributions for time-dependent reliability. *Reliability Engineering & System Safety*, *92*(12), 1651–1658.

Wade, Shirlene, & Kidd, Celeste (2019). The role of prior knowledge and curiosity in learning. *Psychonomic Bulletin & Review*, *26*(4), 1377–1387.

Wagner, Ullrich, Gais, Steffen, Haider, Hilde, Born, Jan, et al. (2004). Sleep inspires insight. *Nature*, *427*(6972), 352–355.

Wang, Haowei, Xu, Tingxue, & Mi, Qiaoli (2015). Lifetime prediction based on Gamma processes from accelerated degradation data. *Chinese Journal of Aeronautics*, *28*(1), 172–179.

Watkins, Christopher J. C. H., & Dayan, Peter (1992). Q-learning. *Machine Learning*, *8*(3–4), 279–292.

Webb, Margaret E., Little, Daniel R., & Cropper, Simon J. (2016). Insight is not in the problem: Investigating insight in problem solving across task types. *Frontiers in Psychology*, *7*, 1424.

Webb, Margaret E., Little, Daniel R., & Cropper, Simon J. (2018). Once more with feeling: Normative data for the aha experience in insight and noninsight problems. *Behavior Research Methods*, *50*(5), 2035–2056.

Whillans, Ashley V, Dunn, Elizabeth W, Smeets, Paul, Bekkers, Rene, & Norton, Michael I (2017). Buying time promotes happiness. *Proceedings of the National Academy of Sciences*, *114*(32), 8523–8527.

Yang, Scott Cheng-Hsin, Vong, Wai Keen, Yu, Yue, & Shafto, Patrick (2019). A unifying computational framework for teaching and active learning. *Topics in Cognitive Science*, *11*(2), 316–337.

Yeung, Nick, & Summerfield, Christopher (2012). Metacognition in human decision-making: confidence and error monitoring. *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)*, *367*(1594), 1310–1321.

Yu, Yuhua, Salvi, Carola, Becker, Maxi, & Beeman, Mark (2024). Solving problems with an aha! increases risk preference. *Thinking & Reasoning*, *30*(3), 509–530.