UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

A General Architecture for Modeling the Dynamics of Goal-Directed Motivation and Decision-Making

Ballard, T.; Neal, A.; Farrell, S.; Lloyd, E.; Lim, J.; Heathcote, A.

DOI
10.1037/rev0000324

Publication date
2022

Document Version
Final published version

Published in
Psychological Review

License
Article 25fa Dutch Copyright Act (https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment)

Link to publication

Citation for published version (APA):
Ballard, T., Neal, A., Farrell, S., Lloyd, E., Lim, J., & Heathcote, A. (2022). A General Architecture for Modeling the Dynamics of Goal-Directed Motivation and Decision-Making. *Psychological Review*, *129*(1), 146-174. https://doi.org/10.1037/rev0000324

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.
UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:16 Oct 2026

AMERICAN PSYCHOLOGICAL ASSOCIATION

Psychological Review

© 2021 American Psychological Association
ISSN: 0033-295X

2022, Vol. 129, No. 1, 146–174
https://doi.org/10.1037/rev0000324

A General Architecture for Modeling the Dynamics of Goal-Directed Motivation and Decision-Making

Timothy Ballard[1], Andrew Neal[1], Simon Farrell[2], Erin Lloyd[1], Jonathan Lim[1], and Andrew Heathcote[3, 4]

[1] School of Psychology, The University of Queensland
[2] School of Psychological Science, University of Western Australia
[3] School of Psychological Sciences, University of Newcastle
[4] School of Psychology, University of Amsterdam

We present a unified model of the dynamics of goal-directed motivation and decision-making. The model—referred to as the GOAL architecture—provides a quantitative framework for integrating theories of goal pursuit and for relating their predictions to different types of data. The GOAL architecture proposes that motivation changes over time according to three gradients that capture the effects of the distance to the goal (i.e., the progress remaining), the time to the deadline, and the rate of progress required to achieve the goal. This enables the integration and comparison of six theoretical perspectives that make different predictions about how these dynamics unfold when pursuing approach and avoidance goals. Hierarchical Bayesian modeling was used to analyze data from three experiments which manipulate distance to goal, time to deadline, and goal type (approach vs. avoidance), and data from the naturalistic context of professional basketball. The results show that people rely on the distance and rate gradients, and to a lesser degree the time gradient, when making resource allocation decisions during goal pursuit, although the relative influence of the gradients depends on the goal type. We also demonstrate how the GOAL architecture can be used to answer questions about the influence of goal importance. Our findings suggest that goal pursuit unfolds in a complex manner that cannot be accounted for by any one previous theoretical perspective, but that is well-characterized by our unified framework. This research highlights the importance of theoretical integration for understanding motivation and decision-making during goal pursuit.

Keywords: goal pursuit, motivation, decision-making, computational modeling, approach/avoidance

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Goals play a central role in the regulation of human behavior. To survive and thrive in a complex, dynamic environment, people need to constantly make decisions about which goals to prioritize, how much to invest in pursuit of those goals, and when to divert their attention toward other goals. Such decisions are challenging because a goal's motivational value—the extent to which one is driven to allocate resources toward it—changes over time. For example, motivation may change as the amount of work required to achieve the goal changes, as the deadline looms, or as other goals become more urgent. Moreover, decisions made during goal pursuit have downstream effects because they shape the environment in which future decisions are made. For example, a day trader who has daily growth targets for a series of different portfolios may choose to work on a portfolio with an easier growth target early in the day. However, in doing so, the growth targets for the neglected portfolios become more difficult to achieve because the time available to work on those portfolios is reduced.

There is a long history of research on goal pursuit dating back to Lewin's (1935) field theory and Hull's (1938) goal gradient hypotheses. There are now many theories that seek to describe how the motivational value of a goal changes during goal pursuit. They focus on a range of factors, including distance to goal (i.e., the amount of progress needed to reach it; Bonezzi et al., 2011; Carver & Scheier, 1990; Heath et al., 1999), time to deadline (Steel & König, 2006), expectancy of goal attainment (Vroom, 1964), difficulty (Brehm & Self, 1989; Wright, 2008), and whether the goal involves approaching a desired state versus avoiding an undesired state (Carver & Scheier, 1998). In many cases, these theories make predictions that are in direct conflict with one another, yet there have been relatively few attempts to directly compare their predictions and identify the elements of the theories that need to be integrated to explain how people regulate time, effort, and other resources during goal pursuit.

The aim of this article is to introduce a unified quantitative framework that can be used to integrate, test, and compare many different theoretical perspectives on goal-directed motivation and

This article was published Online First September 27, 2021.
Timothy Ballard https://orcid.org/0000-0001-8875-4541
Andrew Neal https://orcid.org/0000-0002-5344-9376
Simon Farrell https://orcid.org/0000-0001-7452-8789
Erin Lloyd https://orcid.org/0000-0003-3817-9491
Jonathan Lim https://orcid.org/0000-0002-8469-6572
This project was supported by ARC Discovery Early Career Research Award DE180101340 (awarded to Timothy Ballard), and ARC Discovery Projects DP150102658 (awarded to Andrew Neal), DP200100655 and DP210100313 (awarded to Andrew Heathcote), and DP190102881 (awarded to Simon Farrell, Andrew Neal, and Timothy Ballard). We thank Olivia Grzazek and Carla Magi-Prowse for assisting with data collection. The supplementary material on the open science framework (Ballard et al., 2021; https://osf.io/cjynd/) contains all the data, files, and code used to conduct the experiments and run all the analyses presented in this article.

Correspondence concerning this article should be addressed to Timothy Ballard, School of Psychology, The University of Queensland, Level 3, McElwain Building, St Lucia QLD 4072, Australia. Email: t.ballard@uq.edu.au

decision-making in a rigorous manner. We apply the framework to data from a series of laboratory studies and a field setting in which people have to prioritize the allocation of time and resources while pursuing different types of goals. Our results show that the dynamics of goal pursuit cannot be explained by any one previously existing theory alone. They can only be explained through a framework that integrates these perspectives. This unifying framework then allows us illuminate and resolves areas of theoretical conflict in the goal pursuit literature and, in doing so, provides novel insights into longstanding research questions.

Overview of GOAL Architecture

Our GOAL (goal-oriented action linking) framework is a "linking" architecture that provides a quantitative infrastructure for integrating previously disparate theories of goal pursuit, and for relating their predictions to different types of data. In the sections that follow, we identify six different theoretical perspectives regarding the mechanisms responsible for changes in a goal's motivational value when pursuing approach versus avoidance goals. The GOAL architecture integrates these six perspectives by instantiating each as a special case of a more general model, allowing these perspectives to be directly compared, and conflicts between them to be resolved in a systematic way.

The GOAL architecture proposes that the motivational value of a goal changes in response to changes along three *gradients*. The *distance* gradient reflects the component of a goal's motivational value that changes according to the amount of progress needed to reach the goal (e.g., the amount of portfolio growth needed to meet the day trader's target). The *discrepancy* perspective suggests that motivation is higher when the person is further from the goal (a positive distance gradient). The *proximity* perspective, by contrast, suggests that motivation is higher when the person is closer to the goal (a negative distance gradient). The *time* gradient reflects the component of a goal's motivational value that changes according to the amount of time remaining before a goal's deadline (e.g., the number of trading hours left in the day). As explained below, previous research indicates that motivation typically increases as deadlines approach, suggesting that the time gradient is negative in most cases. The *rate of progress* gradient—accounts for interactions between distance and time. It reflects the component of a goal's motivational value that changes according to the rate of progress required to achieve the goal (e.g., the amount of growth per hour that must be sustained to reach the growth target). The *expectancy* perspective suggests that motivation is higher when goals are easier to achieve. The *difficulty* perspective suggests that motivation is higher when goals are harder to achieve, while the *achievability* perspective suggests that motivation is higher at moderate levels of difficulty.

In the next section, we review each theoretical perspective and describe how the three gradients are formalized within the GOAL architecture. We then apply the architecture to data from two different domains. In Study 1, we apply it to data from an experimental task in which people make a series of prioritization decisions while pursuing approach or avoidance goals with varying distances and deadlines. In Study 2, we apply the architecture to naturalistic data from the context of professional basketball. In both studies, we use hierarchical Bayesian modeling to examine each gradient's influence when pursuing approach versus avoidance goals. This analysis allows us to determine whether any one theoretical perspective is sufficient to account for the full complexity of the ways in which goals influence motivation and decision-making, or whether multiple perspectives need to be integrated. We then examine a series of more focused research questions that address specific areas of conflict between the different perspectives. In Study 2, we also show how the GOAL architecture can be used to examine the impact of goal importance on the relative influence of the three gradients.

The Distance Gradient

A goal is an internal representation of a desired or undesired state that serves as a reference point for evaluating behavior (Austin & Vancouver, 1996; Carver & Scheier, 1998; Neal et al., 2017). The distance to the goal refers to the amount of progress needed to reach the goal. For example, if a day trader has the goal of growing their portfolio by $1 m by the end of the day, and has achieved $400 k growth, they are $600 k from the goal. Differing perspectives regarding the relationship between distance and motivation can be classified into two categories: the *proximity* perspective and the *discrepancy* perspective.

The Proximity Perspective

The proximity perspective assumes that motivation derives from nearness to the goal. It dates back to the goal gradient hypothesis (Hull, 1938; Lewin, 1935), which assumes that the reinforcement provided by the goal is conditioned most strongly to the behavior that immediately preceded the goal being reached, and progressively weakens with increasing distance to the goal (Hull, 1932). This causes the goal to become more salient as it draws nearer, increasing motivation to either approach or avoid it. Early tests of the goal gradient hypothesis were almost exclusively conducted with rats, with distance being operationalized as the literal physical distance between the animal and the desired or undesired stimulus (e.g., Anderson, 1933; Berg & Weisman, 1942; Weiss, 1960). More recent work with humans has conceptualized distance in more general terms—as the amount of any type of progress required to achieve the goal—and supported the goal gradient hypothesis in a variety of domains (e.g., Cryder et al., 2013; Kivetz et al., 2006).

Several mathematical models have been used to capture these effects (e.g., Bonezzi et al., 2011; Kivetz et al., 2006). In particular, a number of theorists have argued that the relationship between distance to goal and motivation follows the *s*-shaped value function from prospect theory (e.g., Heath et al., 1999; Koop & Johnson, 2012; Wang & Johnson, 2012). The closer the reference point (i.e., goal), the greater the psychological impact of each unit of progress that is made, and consequently, the more one is motivated to approach (if below the reference point) or avoid (if above it). It follows that people should prefer to prioritize whichever goal is closest to completion (or failure in the case of avoidance). These arguments suggest that the distance gradient should be negative. Returning to the day trader example, according to the proximity perspective, the trader should become more motivated to work toward the growth target as the gap between the current portfolio value and the target value narrows. When juggling multiple portfolios, they should prioritize whichever target is closer to being reached.

Most research examining the effect of goal distance has focused on approach goals rather than avoidance goals. Consider the case of a day trader working to prevent their losses from exceeding $1 m in

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

a given day. How does their motivation to avoid the $1 m threshold change as the gap narrows? According to the goal gradient hypothesis, the increase in motivation as the gap narrows should be more pronounced for avoidance goals. However, prospect theory models of goal pursuit make the opposite prediction. According to these models, the value function is steeper below the reference point than above the reference point. When the outcome variable is below the reference point, the person is pursuing an approach goal. When it is above the reference point, the person is pursuing an avoidance goal. This means that the distance-motivation relationship should be stronger when pursuing approach goals.

The Discrepancy Perspective

The discrepancy perspective assumes that people are motivated by the existence of a discrepancy between one's goal and current state. This perspective can be traced back to perceptual control theory (Powers, 1973, 1978) and the theories of work motivation (e.g., Carver & Scheier, 1990; Jagacinski & Flach, 2003; Klein, 1989; Lord & Levy, 1994; Vancouver, 2005) and well-being (Higgins, 1997; Trew, 2011) that it inspired. These theories assume that a discrepancy alerts the individual of the need to apply effort, which in turn motivates them to take actions that correct the discrepancy. It follows that there should be a positive relationship between distance to a goal and the amount of effort or attention directed toward it, because longer distances signal a larger discrepancy to be corrected (e.g., DeShon & Gillespie, 2005; Unsworth et al., 2009). In our running example, if the day trader has an approach goal for a particular portfolio, the discrepancy perspective predicts that the trader will be less motivated to work on the portfolio as the gap between the current and target values decreases. When juggling multiple portfolios, they should prioritize the portfolio with the largest gap.

The positive distance gradient predicted by the discrepancy perspective is supported by studies of approach goal pursuit demonstrating that the magnitude of the discrepancy is positively related to the amount of effort expended toward a goal (Campion & Lord, 1982; Kernan & Lord, 1990) and the likelihood of choosing to prioritize it (Schmidt & DeShon, 2007; Schmidt & Dolis, 2009; Schmidt et al., 2009). For example, the discrepancy perspective accounts for the balance matching heuristic observed in the behavioral economics literature, in which individuals managing multiple debts tend to make larger repayments to debts with higher balances (Gathergood et al., 2019).

Although perceptual control theory does not explicitly address avoidance goals, it has been used to argue that, when faced with an avoidance goal, small discrepancies are motivating because they make threat of reaching the undesired state salient (e.g., Ballard, Yeo, Loft, et al., 2016; Carver & Scheier, 1998). It follows that motivation to avoid an undesired state should increase as it gets closer to being realized, which is the same prediction the proximity perspective makes regarding avoidance goals. This prediction is supported by a small number of studies showing that people managing avoidance goals generally prioritize whichever goal has the smallest discrepancy (Ballard et al., 2018; Ballard, Yeo, Loft, & Farrell, 2016; Ballard, Yeo, Neal, et al., 2016).

Modeling the Distance Gradient

The discrepancy and proximity perspectives make contrasting predictions regarding the nature of the distance gradient in the approach context, but the same prediction in the avoidance context. The following form of the distance gradient within the GOAL architecture is sufficiently flexible to capture both perspectives:

$$DG_{ij} = \begin{cases} D_{ij}^{\delta}, & \text{if } \delta >= 0 \\ (1 - D_{ij}^{-\delta}), & \text{if } \delta < 0 \end{cases} \quad (1)$$

where DG_{ij} is the level of the distance gradient for goal i at time j. D_{ij} represents the distance to goal i at time j. D_{ij} is scaled to be between 0 and 1, where 1 represents the maximum possible distance from the goal. The δ parameter controls the shape of the distance gradient, and is constrained between -1 and 1. When $\delta < 0$ the distance gradient is negative, so motivation decreases with distance to the goal, consistent with the proximity perspective. When $\delta > 0$, the distance gradient is positive, consistent with the discrepancy perspective. When δ is equal to either -1 or 1, the distance gradient is linear. When δ falls between these values, the effect of distance to goal on motivation exhibits diminishing returns.[1]

The predictions of proximity and discrepancy perspectives, as formalized by the GOAL architecture, are summarized in the first two panels in the top row of Figure 1. The figure highlights that these perspectives assume that motivational value is determined solely by distance to goal and is unaffected by time to deadline. In order to achieve a more complete understanding of the dynamics of goal pursuit, we need to distinguish between how motivation responds to changes in the amount of *progress* remaining, and how it responds to changes in the amount of *time* remaining. In the next section, we address the second gradient—the time gradient—which captures the effects that changes in the time remaining have on a goal's motivational value.

The Time Gradient

There is a large body of work that has attempted to understand the effects of time on motivation by studying intertemporal choice: Choice between outcomes of different values and that occur at different future times (for reviews see: Doyle, 2013; Green & Myerson, 2004; Lempert & Phelps, 2016; Loewenstein et al., 2008; Read, 2004; Rung & Madden, 2018; Scholten & Read, 2010). This research has found that people will often prefer smaller more immediate rewards to larger rewards available later, a phenomenon variously referred to as temporal, time, or delay discounting.

Although temporal discounting has typically been studied in paradigms involving static, one-shot decisions, it may also be influential in the context of goal pursuit. Longer deadlines often mean that more time must pass before the rewards associated with goal achievement can be obtained. This may lead people to discount the perceived value of goal achievement. In other words, the reward associated with achieving a goal (or the punishment associated with goal failure) may be more salient, and therefore more motivating, when the deadline determining whether the goal has been achieved or failed is nearer. These arguments suggest that the time gradient

[1] Although other shapes are possible, such as convex or nonmonotonic effects of distance, we found the relatively simple formulation in Equation 1 was sufficient, and importantly that simplicity aided model identifiability. In more complex cases (e.g., Bonezzi et al., 2011), mixtures of gradients of the form in Equation 1 with positive and negative parameters might be considered.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Figure 1

Summary of Competing Predictions Regarding the Distance, Time, and Rate of Progress Gradients

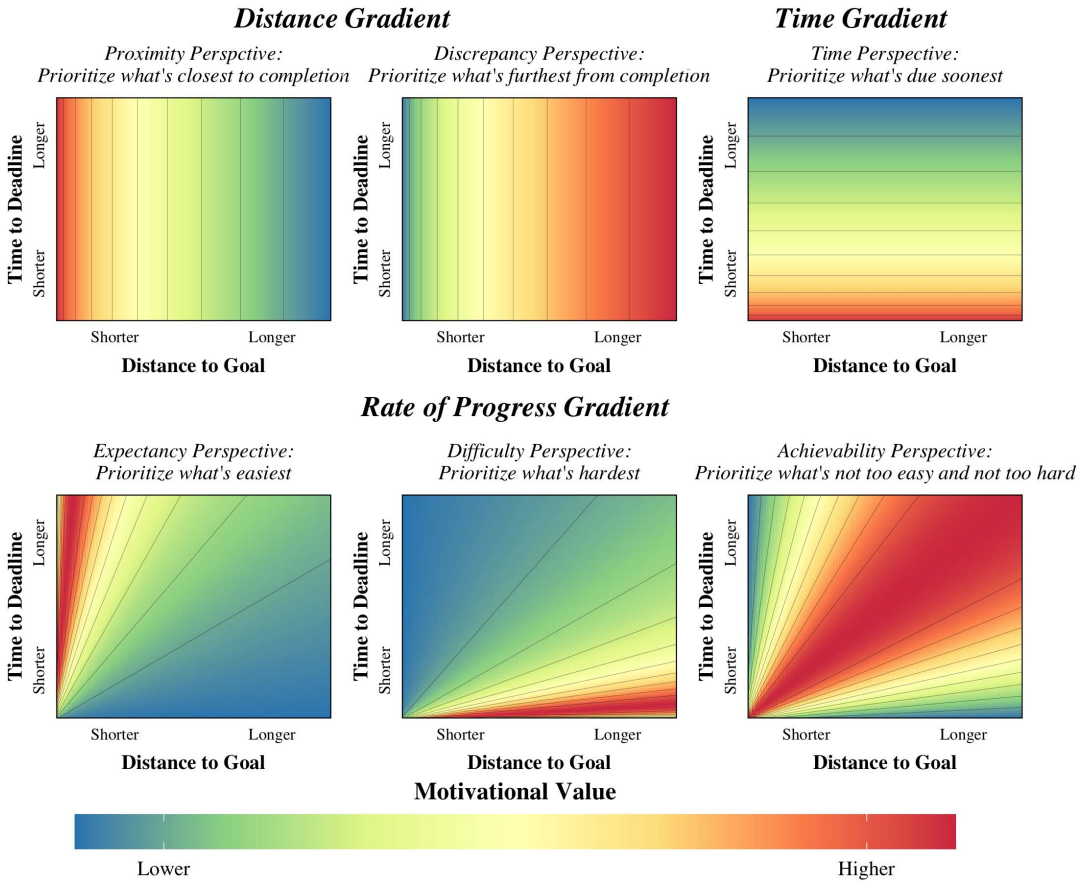

Note. In each panel, the *x*-axis represents the distance to a hypothetical goal and the *y*-axis represents the time to the deadline for that goal. The color of the surface represents the motivational value of the goal given the combination of distance to goal and time to deadline. Note that the predictions and captions apply to approach goals. For avoidance goals, the discrepancy, expectancy, and difficulty perspectives predict the opposite pattern to what is shown in the figure. Predictions for each perspective were generated using the following parameter values (with unlisted parameters set to zero): Proximity, $w_1 = 1$, $\delta = -0.5$; Discrepancy, $w_1 = 1$, $\delta = 0.5$; Time, $w_2 = 1$, $\tau = -0.5$; Expectancy, $w_3 = 1$, $\tau = 0.005$; Difficulty, $w_3 = 1$, $\alpha = 0.995$; Achievability, $w_3 = 1$, $\alpha = 0.5$. See the online article for the color version of this figure.

should be negative. All else being equal, the day trader in our example should become more motivated to work on a portfolio as the deadline for its growth target draws nearer. Consequently, they should prioritize whichever deadline is due soonest. This prediction parallels the proximity perspective regarding the distance gradient, with the difference being that proximity perspective assumes that motivation increases as the amount of progress required to achieve the goal decreases whereas the prediction regarding the time gradient is that motivation increases as the amount of time available before the deadline decreases.

Research has also shown that rates of temporal discounting depend on whether the outcome is desirable or undesirable (the so-called "sign effect"; Read, 2004). For example, people tend to discount the pleasure of gaining a monetary reward more strongly than the pain of losing money (Gonçalves & Silva, 2015; Murphy et al., 2001; Read et al., 2017; Thaler, 1981). This evidence suggests that the time gradient should be more influential for approach goals, where a desirable state is being pursued, compared with avoidance goals, where an undesirable state is being avoided.

Modeling the Time Gradient

The GOAL architecture assumes that the time gradient has the same functional form as the distance gradient, which is both parsimonious and makes the two gradients directly comparable:

$$TG_{ij} = \begin{cases} T_{ij}^{\tau}, & \text{if } \tau >= 0 \\ (1 - T_{ij}^{-\tau}), & \text{if } \tau < 0 \end{cases} \quad (2)$$

where TG_{ij} is the level of the time gradient for goal i at time j. T_{ij} is time remaining before the deadline for goal i at time j. The τ parameter controls that shape of the time gradient, and would be expected to be negative, consistent with existing models of intertemporal choice and motivation (e.g., Ainslie, 1975; Steel & König, 2006). As illustrated

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

in the top-right panel of Figure 1, the literature reviewed predicts that this gradient is negative, with motivation increasing as the deadline for goal achievement becomes nearer.

The Rate of Progress Gradient

The distance and time gradients represent the independent effects of distance to goal and time to deadline on a goal's motivational value. But are these effects sometimes interdependent? For example, the day trader attempting to increase their portfolio value to a desired level may expend more effort as the deadline looms if the gap between the portfolio's current and desired value is small enough that the goal is sufficiently easy to achieve. However, if this gap is large enough that the desired value is too difficult to achieve, the trader may instead reduce their effort as the deadline draws nearer in order to focus on more achievable goals.

These arguments are consistent with a range of theories suggesting that motivation is influenced by the difficulty of goal achievement in ways that cannot be represented by a simple additive combination of separate distance and time gradients. Although these accounts may not explicitly address the effects of distance to goal and time to deadline, a relationship between goal difficulty and motivation suggests that people combine information about distance and time because both determine the difficulty of goal achievement. For example, the goal of increasing a portfolio value by \$1 m when there are 4 hr remaining to do so is easier to achieve than attempting to generate the same increase in value in 1 hr, despite the fact that the amount of progress that needs to be made is the same. Similarly, an increase of \$1 m in 4 hr is easier to achieve than an increase of \$5 m in 4 hr, despite the fact that both goals have the same deadline.

We propose that the interaction between time and distance can be conceptualized in terms of the *rate of progress* that is required to reach the goal. That is, goal difficulty is proportional to the amount progress that must be sustained per unit of time remaining in order to achieve the goal—or in other words, the ratio of the amount of progress and time remaining. In order to capture this relationship, we propose a third gradient: The rate of progress gradient. Despite the association between goal difficulty and motivation being widely studied, there are conflicting views regarding its nature. In the following subsections, we describe three different perspectives, and describe how they are captured by the GOAL architecture's rate of progress gradient.

The Expectancy Perspective

Expectancy has played a central role in theories of motivation and decision-making (Bandura, 1977; Vroom, 1964). Within the motivation literature, it has been defined as the belief that effort will result in a desired outcome (e.g., reaching an approach goal or not reaching an avoidance goal). Expectancy is assumed to be positively related to effort (Bandura, 1977; Locke & Latham, 1990; Steel & König, 2006; Vroom, 1964). Indeed, some have argued that the increase in effort often observed as the goal becomes closer to being achieved (the proximity perspective) is actually due to expectancy (Liberman & Förster, 2008). As the goal draws nearer, its expectancy increases, which in turn increases the motivation to pursue it. Expectancy is also an important element in theories of economic decision-making (e.g., Luce & Raiffa, 1957; von Neumann & Morgenstern, 1947), where it has been defined as the subjective likelihood of an outcome occurring. They assume that the expectancy of a favorable outcome increases the likelihood of selecting a particular course of action. These motivational and decision-making theories, which we refer to collectively as the expectancy perspective, suggest that the easier it is to achieve a desired state (in the approach context) or avoid an undesired state (in the avoidance), the higher the goal's motivational value.

The Difficulty Perspective

The expectancy perspective treats the allocation of effort as a strategic process, whereby the person chooses courses of action that they evaluate as being more likely to pay off. However, an opposing view is that people are motivated by goals they find challenging. Inzlicht et al. (2018) describe a so-called "effort paradox," in which effort exertion, while sometimes perceived as costly, can also be a source of value. In many cases, the exertion of effort itself is intrinsically rewarding (Csikszentmihalyi, 2014). Alternatively, effort may be rewarding because it serves a charitable or collective cause (Olivola, 2011; Olivola & Shafir, 2013) or leads individuals to be more satisfied with the product of their efforts (Arkes et al., 1994; Inglis et al., 1997; Kameda et al., 2002; Loewenstein & Issacharoff, 1994; Mochon et al., 2012; Norton et al., 2012). Information processing models (Hendy et al., 1997; Hockey, 1997; Humphreys & Revelle, 1984; Kool et al., 2010; Westbrook & Braver, 2015) suggest that people regulate their effort expenditure based on the demands of the situation. As such, people work harder when the goal is more difficult (Bryan & Locke, 1967; Latham & Locke, 1975; Palada et al., 2018; Peters et al., 1984). According to this view, which we refer to as the difficulty perspective, a goal's motivational value is positively related to the difficulty of reaching a desired state or avoiding an undesired state.

The Achievability Perspective

Although the expectancy and difficulty perspectives are well established, it is possible that neither is sufficient to account for the complex dynamics of motivation during goal pursuit. For example, if a task is too easy, people may divert resources to other tasks where time and effort are needed more (Carver & Scheier, 1998; Louro et al., 2007). On the other hand, if a goal is too hard to achieve, people may give up (Schmidt & Dolis, 2009; Vancouver et al., 2010). These behaviors violate the expectancy and difficulty perspectives, respectively, highlighting the need to consider a third perspective regarding the nature of the rate of progress gradient.

The third perspective, which we refer to as achievability, proposes that the influence of difficulty and expectancy on motivation is nonmonotonic. It dates back to Atkinson's work on achievement motivation (e.g., Atkinson & Birch, 1970), which proposes an inverted-*u*-shaped relationship, where motivation is the highest when difficulty and expectancy are moderate and the lowest when either difficulty is very low (and therefore expectancy is very high) or very high (and therefore expectancy is very low). Others have proposed a discontinuous relationship, where the intensity of motivation increases with difficulty up to the point where the individual decides that the goal cannot be achieved or that it is not worth the effort. When difficulty exceeds this level, motivation drops to zero (Brehm & Self, 1989; Kukla, 1972; Wright, 2008). The inverted-u and discontinuous models are

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

difficult to differentiate empirically, but there is broad support for a nonmonotonic relationship between difficulty or expectancy and motivation (e.g., Biner, 1987; Brehm et al., 1983; Karabenick & Youssef, 1968; Vancouver, 2008; Wright et al., 1986). For example, Louro et al. (2007) found that dieters striving for weight loss goals allocated more effort to the goal when they had a moderate expectancy of achieving it than when they had either a low or a high expectancy. These arguments suggest that motivation should be the highest when the rate of progress required to reach a desired or undesired state is moderate, and lower when the rate required is too low or too high.

Modeling the Rate of Progress Gradient

The model of the rate gradient in the GOAL architecture must be sufficiently flexible to account for a nonmonotonic relationship between the required rate of progress and motivation. The specific functional form we assume places the rate gradient on a commensurate scale to the distance and time gradients so that all three can be combined additively:

$$RG_{ij} = \beta\left[\frac{\alpha T_{ij}}{D_{ij}} + \frac{(1-\alpha)D_{ij}}{T_{ij}}\right]^{-1}, \quad (3)$$

where RG_{ij} is the level of the rate gradient for goal i at time j. The β value ensures that the rate gradient is on same scale as the distance and time gradients, and is a direct function of the α parameter: $\beta = 2(1-\alpha)\sqrt{\frac{\alpha}{1-\alpha}}$. The α parameter takes on a value of between 0 and 1 and controls the gradient's shape. As α approaches one, the gradient increases linearly with the required rate of progress. As α approaches zero, the gradient increases linearly with the inverse of the required rate (T/D). As α moves closer to 0.5, the rate gradient becomes nonmonotonic and higher at moderate rates. The α parameter, therefore, controls the correspondence between the rate gradient and the expectancy, difficulty, and achievability perspectives. In doing so, it provides a way of accounting for Inzlicht et al.'s (2018) effort paradox. When the rate gradient aligns with the expectancy perspective, goals that require more effort to achieve are demotivating. When the rate gradient aligns with the difficulty perspective, goals that require more effort to achieve are motivating.

The α parameter also captures the influence of skill level. A more skilled individual may be expected to have a higher α value in the approach context and a lower α value in the avoidance context, as these values indicate they are more motivated by more challenging goals. A less skilled individual may be expected to have lower α value in the approach context and a higher α value in the avoidance context, as these values indicate they are more motivated by easier goals.

The competing predictions regarding the expectancy, difficulty, and the achievability perspective are illustrated in the bottom row of Figure 1. As can be seen, all three perspectives assume that motivation changes with the required rate of progress. Under the expectancy perspective motivation to pursue an approach goal is the highest when there is a lot of time remaining and only a small distance to cover before the goal is reached. Under the difficulty perspective motivation to pursue an approach goal is the highest when there is little time remaining and a lot of distance to cover. Both perspectives predict the opposite relationships for avoidance goals (not shown in the figure). The achievability perspective assumes that motivation is the highest when the time remaining before the deadline and the distance to the goal have similar relative values (e.g., are both relatively high or both relatively low), which makes the goal moderately difficult to achieve.

Putting It All Together

To summarize, the GOAL architecture is an integrative framework for modeling the dynamics of motivation during goal pursuit. Its integrative nature allows for more specific theories to be instantiated as special cases, which facilitates their evaluation and comparison. The architecture assumes that changes in motivation are attributable to changes in position along three gradients. The distance gradient represents changes in motivation attributable to the distance to the goal (i.e., the amount of progress needed before the goal is reached). The time gradient represents changes in motivation attributable to changes in the amount of time before the deadline by which the goal must be achieved or is failed. The rate gradient represents changes in motivation attributable to the rate of progress required to reach the goal.

The GOAL architecture assumes that the motivational value of goal i at time j, denoted M_{ij}, is a weighted additive combination of distance, time, and rate gradients:

$$M_{ij} = w_1 DG_{ij} + w_2 TG_{ij} + w_3 RG_{ij}, \quad (4)$$

where w_1, w_2, and w_3 are weight parameters that determine the relative influence on motivation of the distance, time, and rate gradients, respectively. Without loss of generality we assume the weights sum to one, so only two (w_1 and w_2) need to be estimated (i.e., $w_3 = 1 - w_1 - w_2$).

Overview of Research Aims

Our broad aim is to examine whether any one of the six theoretical perspectives described above is ever sufficient to account for changes in motivation during goal pursuit, or whether multiple perspectives (and if so, which) must be integrated in order to account for the full complexity of motivational effects. To facilitate this broad comparison, we structure our analysis around six focused research questions that address untested predictions or areas of conflict between or within the various perspectives (see Table 1).

The first research question addresses the conflict between the discrepancy and proximity perspectives on the distance gradient when pursuing approach goals. As the distance to the goal reduces and the goal becomes closer to being attained, does motivation increase (as predicted by the proximity perspective) or decrease (as predicted by the discrepancy perspective)?

The second research question addresses the difference in the distance gradients for approach versus avoidance goals. The goal gradient hypothesis assumes that this gradient is steeper for approach goals than for avoidance goals. This implies that motivational increases should be more pronounced when moving closer to an undesired state than to a desired state. However, prospect theory models of motivation imply the opposite prediction. They assume that changes in the distance to the goal have a stronger impact when the individual is below the reference value and attempting to move toward it. This prediction is implied by the fact that the value function is steeper in this domain than it is when the individual is above the reference value and wishes to avoid moving toward the

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Table 1
Summary of Research Questions

Research question	Relevant perspectives
1. Does motivation increase or decrease as the distance to an approach goal decreases?	Proximity vs. Discrepancy
2. Is the effect of distance to goal on motivation stronger when pursuing approach or avoidance goals?	Proximity
3. Is the effect of time to deadline on motivation stronger when pursuing approach or avoidance goals?	Time
4. How does motivation change as the rate of progress required to reach the goal changes?	Expectancy vs. Difficulty vs. Acheivability
5. Are increases in motivation as the distance to goal decreases due to proximity or expectancy?	Proximity vs. Expectancy
6. Are increases in motivation as the time to deadline decreases due to time itself or goal difficulty?	Time vs. Difficulty

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

goal. Hence, motivational increases should be more pronounced as the individual moves toward an approach goal than toward an avoidance goal. We attempt to resolve this conflict by comparing the relative influence of the distance gradient between approach and avoidance goals.

The third research question addresses differences in the time gradient for approach versus avoidance goals. The sign effect in intertemporal choice implies that motivational increases as the deadline looms are stronger when pursuing approach goals than avoidance goals. To our knowledge, this prediction has never been tested in the goal pursuit context. Thus, it is unclear whether the sign effect demonstrated in static, one-shot intertemporal choices generalizes to the more dynamic context of goal pursuit. We test this prediction by comparing the relative influence of the time gradient between approach and avoidance goals.

The fourth research question addresses the conflict between the expectancy, difficulty, and achievability perspectives. The expectancy perspective implies that motivation increases monotonically with the ease of achieving the goal. The difficulty perspective makes the opposite prediction that motivation increases monotonically with the difficulty of achieving the goal. The achievability perspective implies a nonmonotonic relationship between difficulty and motivation, with motivation being the highest when goal achievement is moderately difficult and lower when the goal is easier or more difficult to achieve. We attempt to reconcile this conflict by examining the shape of the rate of progress gradient.

The fifth research question addresses a conflict between the proximity and expectancy perspectives in accounting for the purported increase in motivation as goal achievement comes closer. Research examining the goal gradient hypothesis has traditionally assumed that this increase in motivation arises because the goal become more salient as it becomes closer to being reached (Brown, 1948; Hull, 1932, 1938). This assumption is shared by prospect theory models of motivation, and is consistent with the proximity perspective, suggesting that motivation should increase as the distance to the goal decreases. By contrast, Liberman and Förster (2008) argued that goal gradient effects may be attributable to an increase in the expectancy of reaching the goal. If that argument is correct, motivation should be more strongly influenced by the rate gradient (which captures the effects of expectancy) than the distance gradient (which captures the effects of proximity). We attempt to resolve this conflict by examining the relative influence of the distance versus rate gradients.

The sixth research question addresses a conflict between the time and difficulty perspectives in accounting for the purported motivational increases as deadlines becomes closer to being reached. According to the perspective underlying the time gradient, this increase in motivation arises solely as a result of the looming deadline, and is independent of the amount of progress that needs to be made to reach the goal. According to the difficulty perspective captured by the rate gradient, this increase is triggered by a greater rate of progress being required to reach the goal, causing mobilization of more effort in response to the increasing demands of the task. We attempt to resolve this conflict by examining the relative influence of the time versus rate gradients.

We examine these research questions in two studies. The first concerns a laboratory paradigm that manipulates the distance to goal and time to deadline, and examines their impact on people's prioritization of different goals. We compare a series of alternative models within the GOAL architecture that operationalize different assumptions about which gradients are necessary to account for the dynamics of motivation, and interpret the gradients to answer the six research questions. In the second study, we analyze data from professional basketball, examining how the distance to the goal (i.e., the score difference) and the time remaining in the game affects how much the team invests in trying to win the game. This allows us to examine the research questions in a field setting. Because the value of goal achievement naturally varies across the season, this study also provides an opportunity to examine how the importance of the goal affects these motivational dynamics.

Study 1

In Study 1, we used a revealed preference paradigm, where participants make choices about which of two goals to prioritize. In order to link the GOAL architecture's continuous motivational value to observed binary prioritization decisions we use a logistic model that transforms the difference in motivational values between each choice into a set of predicted choice probabilities. Specifically, the probability of prioritizing Goal A at time j is calculated as follows:

$$pr_{A,j} = \frac{1}{1 + \exp[-\theta(M_{A,j} - M_{B,j})]}, \quad (5)$$

where θ controls the sensitivity to the difference in motivational values, reflecting decisiveness. Larger values of θ result in a sharper transition in response probability as the difference in motivation changes from favoring Goal A to Goal B.[2]

Study 1 involved three experiments. In Experiment 1a, we manipulated the distance to each goal at the start of each goal pursuit episode, while holding the deadline constant between goals.

[2] A more general formulation could assume different values of θ for each goal, instantiating a response bias, but that was not necessary for our paradigm where such systematic biases are unlikely.

In Experiment 1b, we manipulated the time to deadline for each goal, while holding the starting distance constant. In Experiment 1c, we simultaneously manipulated both the starting distance and the deadline. In each experiment, we also manipulated goal type, such that participants sometimes pursued two approach goals and at other times two avoidance goals. As we otherwise used the same paradigm in all three experiments, we report the method for the three experiments together. We then report the results from all three experiments together in the next section. These studies were granted approval by the University of Queensland's behavioral and social sciences ethical review committee (Clearance number: 2015000293).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Method

Participants

The full sample consisted of 330 participants. Participants in Experiments 1a and 1b were undergraduates at the University of Queensland who took part in exchange for course credit. Experiment 1a consisted of 146 participants (57% female, 43% male) with ages ranging from 17 to 51 years ($M = 20.26$, $SD = 5.53$). Experiment 1b consisted of 130 participants (60% female, 40% male) with ages ranging from 17 to 50 years ($M = 19.93$, $SD = 5.01$). Participants in Experiment 1c were members of the University of Queensland local community who received $20 for taking part. Experiment 1c consisted of 54 participants (39% female, 61% male) with ages ranging from 18 to 40 years ($M = 22.41$, $SD = 3.45$). A total of 36 participants were excluded from the analysis because they did not finish the study (16 from Experiment 1a, 16 from Experiment 1b, and 4 from Experiment 1c). This left final samples of 130, 114, and 50 participants for Experiments 1a, 1b, and 1c, respectively.[3]

Experimental Paradigm

We used a paradigm in which participants had to make a series of prioritization decisions over time while attempting to complete different types of goals, and in which the distance to each goal and the time to the deadline for each goal could be manipulated. The paradigm involved a farm simulation task. The task was broken down into a series of trials, which each representing an independent goal pursuit episode. For each trial, participants had to simultaneously manage either two crops (in the approach condition) or two weeds (in the avoidance condition). Each trial was broken down into a series of stages (days), with the participant making a prioritiation decision in each one.

Approach Condition

In approach trials (see left-hand panel of Figure 2), participants had to manage two out of three possible crops (wheat, corn, or rice). Their objective was to ensure that, for as many crops as possible, the height of the crop was above a threshold height by the end the crop's growing season. On each day of the growing season, participants had to choose which crop to treat. Treating a crop increased its potential for growth such that, on any given day, the treated crop would grow more on average than the neglected crop.

Avoidance Condition

In avoidance trials (see right-hand panel of Figure 2), participants had to manage two out of three possible weeds (thistle, nettle, or lantana). Their objective was to ensure that, for as many weeds as possible, the height of the weed was below a threshold height by the end the weed's growing season. On each day participants had to choose which weed to treat, which suppresses its potential for growth on that day, so it would grow less on average than the neglected weed.

Design

In all three experiments, goal type was manipulated within participants, but was held constant across the two goals in each trial. That is, there were no trials in which participants had to simultaneously manage one approach and one avoidance goal.

Experiment 1a

In this experiment, we also manipulated the difference between the height of the crop and the threshold height at the start of each trial (i.e., starting distance). The threshold height was always 240 cm. The combination of starting distances was manipulated within participants. We varied the starting distance for each goal across seven levels (30, 45, 67.5, 90, 112.5, 135, or 150 cm), and including only the nonredundant combinations of distances (e.g., a distance of 30 for Goal A and 45 for Goal B is redundant with a distance of 45 for Goal A and 30 for Goal B). This manipulation resulted in 28 unique starting distance combinations and so the design had 56 unique experimental conditions (Two Goal Type × 28 Starting Distance Combination). Increasing the starting distance made approach goals harder to achieve, but avoidance goals easier to achieve. However, the overall difficulty of the task was constant across goal type conditions because an approach goal with a starting distance of 30 cm was equally easy to achieve as an avoidance goal with a starting distance of 150 cm, an approach goal with a starting distance of 45 cm was equally easy to achieve as an avoidance goal with a starting distance of 135 cm, and so on.

In this experiment, each stage in the trial represented 1 day in the growing season. The growing season for each crop/weed was always 30 days. In the approach condition, the growth of the treated crop was sampled from a normal distribution with a mean of 6 cm and a standard deviation of 3 cm. The growth of the neglected crop was sampled from a distribution with a mean of 0 cm and a standard deviation of 3 cm. In the avoidance condition, the growth of the treated weed was sampled from a normal distribution with a mean of 0 cm and a standard deviation of 3 cm. The growth of the neglected weed was sampled from a distribution with a mean of 6 cm and a standard deviation of 3 cm.

Experiment 1b

Experiment 1b was identical to Experiment 1a except that starting distance was held constant and deadline was manipulated within participants. All crops/weeds started at a height of 90 cm below the threshold (i.e., the starting distance was always 90 cm). Deadline

[3] Due to a technical error, approximately 3% of the observations in Experiments 1a and 1b were not recorded. The missingness of the data was not systematically related to any of the experimental manipulations.

Figure 2

Experimental Paradigm Used in Study 1 (This Example Screen Shot was Taken From Experiment 1b)

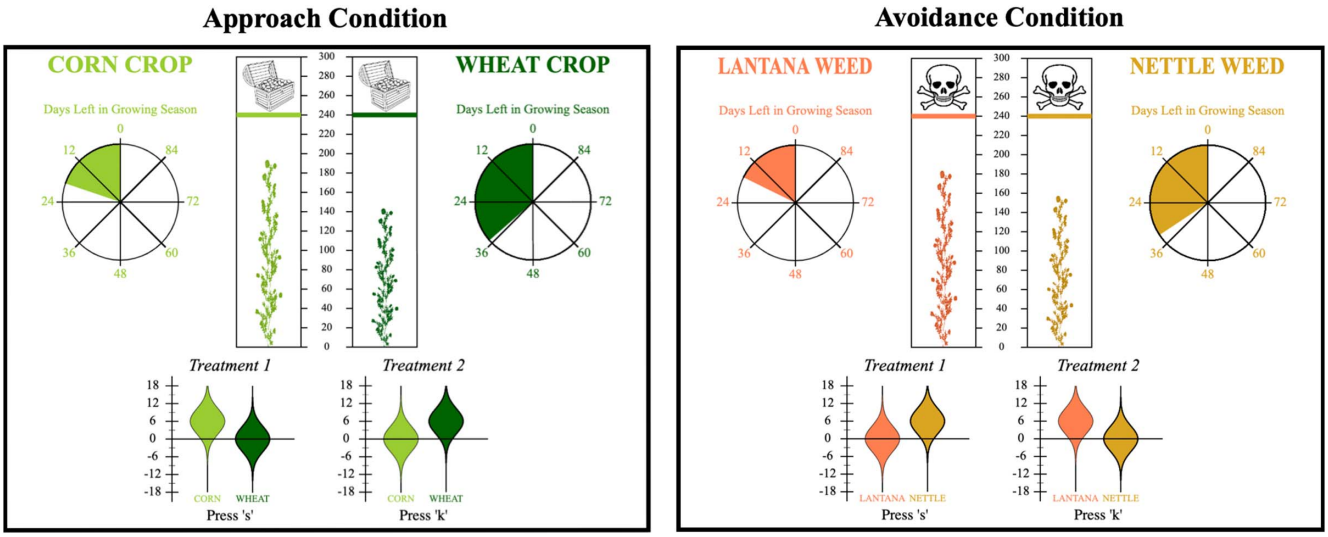

Note. The left-hand panel shows the approach condition, whereas the right-hand panel shows the avoidance condition. Distance to goal is operationalized as the difference between the height of the crop and the target height. Time to deadline is operationalized as the amount of time left in the growing season. See the online article for the color version of this figure.

was manipulated in a similar manner to starting distance in Experiment 1a by varying the length of the growing season for each crop/weed across seven levels (18, 20, 24, 30, 40, 60, or 90 days) and including only the nonredundant combinations of deadlines. This design produced 56 unique experimental conditions (Two Goal Type × 28 Deadline Combination). Increasing the deadline made approach goals easier to achieve because there were more opportunities for growth. For the same reason, increasing the deadline made avoidance goals harder to achieve. However, the overall difficulty of the task was constant because achieving an approach goal with a deadline of 90 was approximately the same in difficulty as achieving an avoidance goal with a deadline of 18, and so on. The deadline manipulation also produced trials that were approximately equal in difficulty to the trials produced by the starting distance manipulation in Experiment 1a.

Experiment 1c

In this experiment, we manipulated both starting distance and deadline within participants, and so more trials were required than in Experiments 1a and 1b. In order to accommodate the extra trials, we reduced the number of decisions per trial so that no more than eight decision stages were included in any trial. Each decision stage in this experiment represented 1 month in the growing season. We varied the length of the growing season across four levels (1, 2, 4 or 8 months) and the starting distance for each goal across four levels (10, 20, 40 or 80 cm). The threshold height was always 200 cm in this experiment. As with Experiments 1a and 1b, we included only nonredundant cells in the design. These manipulations resulted in 120 unique Starting distance × Deadline combinations, and so the design had 240 unique experimental conditions (120 approach trials and 120 avoidance trials).

In the approach condition, the growth of the treated crop was sampled from a normal distribution with a mean of 20 cm and a standard deviation of 20 cm. The growth of the neglected crop was sampled from a distribution with a mean of 0 cm and a standard deviation of 20 cm. In the avoidance condition, the growth of the treated weed was sampled from a normal distribution with a mean of 0 cm and a standard deviation of 20 cm. The growth of the neglected weed was sampled from a distribution with a mean of 20 cm and a standard deviation of 20 cm.

Procedure

Upon starting the experiment, participants were first presented with computerized task instructions and a series of practice trials. Participants then completed each unique experimental condition. In Experiments 1a and 1b, the 56 unique conditions were completed in random order. In Experiment 1c, the approach and avoidance conditions were blocked such that participants either completed all approach trials first or all avoidance trials first (with the order of trials randomized within each block). In each trial, the position of the crop or weed on the screen (i.e., left vs. right) was randomly determined, as well as the type of crop or weed that the participant encountered. Participant treated the left crop/weed by pressing the *a* key and the right crop/weed by pressing the *k* key. After they selected a treatment, the task would immediately update to the next decision stage, so that the participant could see how much the height of each crop or weed had changed and make the next treatment decision. The heights of the crops/weeds, the threshold height, and the number of days/months left in each growing season were displayed on the screen throughout the trial.

For trials where the crops/weeds being treated had the same growing season, the trial ended when the growing season finished. For trials where the crops/weeds did not have the same growing

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

season, the trial ended when the longer growing season finished. In these trials, the crop or weed with the growing season that finished earlier was replaced by a new crop/weed after its growing season ended. These replacements were required to ensure that the participant still had to make prioritization decisions after the first deadline was reached. This ensured that the opportunity cost of prioritizing the goal with the longer deadline was the same across its entire growing season. The starting distance for the replacement crops or weeds was always equal to three times the number of days/months remaining in the other growing season, making the new goal moderately difficult to achieve.

Behavioral Results

We begin by describing the patterns of choices made by participants, then report the analysis using the GOAL architecture. To facilitate interpretation and avoid repetition, we combined the results across the three experiments. Where relevant, we have provided the figures with the results broken down by experiment in the Supplementary Material on the open science framework (https://osf.io/cjynd/). We included in the analysis only decisions that were made when both crops/weeds were below their threshold and before either deadline had been reached, because these decisions are the most theoretically diagnostic. The data set analyzed included a total of 298,463 observations (with an average of 1,015 observations per participant), ensuring that we had sufficient precision of measurement to quantify the influence of the three gradients in each goal type condition for each participant (see Smith & Little, 2018).

We use vector field plots to illustrate how goal prioritization decisions were influenced by the combination of distance to goal and time to deadline for each of the two goals being pursued (see Figure 3). The distance to goal (D) and time to deadline (T) variables were converted to the 0–1 scale assumed by the GOAL architecture by dividing the raw value of that variable by 1 + the maximum value observed in the experiments (for distance, 187 cm; for deadline, 90 days in Experiments 1a and 1b and 8 months in Experiment 1c). The orientation of the vector indicates the average direction which the participants choose to move. The upward component of each vector indicates choices that decrease the distance to the goal on the left whereas the rightward vector component indicates choices that decrease the distance to the goal on the right. Hence, vectors pointing up indicate choices that promote growth of the left crop/weed, vectors pointing right indicate choices that promote growth of the right crop/weed and vectors along the diagonal indicate decisions that, on average, prioritize the growth of both equally.

Approach Condition

The results for the approach condition (the left set of panels) suggest that goal prioritization decisions vary depending on the combination of distance to goal and time to deadline for each goal. In the top-left panel, where the time to both deadlines is relatively short, people tend to prioritize whichever goal is closer to being achieved. This can be seen by the fact that the vectors above the diagonal, which indicate states in which the left-hand goal is closer to being achieved, tend to point upward (i.e., toward the nearest threshold). Vectors below the diagonal, which indicate states in which the right-hand goal is closer to being achieved, point rightward. In the middle and bottom-right panels the deadlines are longer, but are still similar between the two goals. There, the tendency to prioritize the goal that is closer to being achieved is weaker. This can be seen by the fact that there are fewer arrows pointing directly upward or rightward and more arrows pointing diagonally toward the top-right corner of the panel (indicating that goals are equally likely to be prioritized).

In the top-middle, top-right, and middle-right panels the deadline for the left-hand goal is shorter than for the right-hand goal. In this scenario, there is a stronger tendency to prioritize the left-hand goal (i.e., the arrows tend to point more upward than rightward). The opposite is true in scenarios where the deadline for the right-hand goal is shorter (the middle-left, bottom-left, and bottom-middle panels). Here, the results show a tendency to prioritize the right-hand goal (i.e., the arrows tend to point more rightward than upward). Together these results suggest that distance to goal and time to deadline exerts complex effects on motivation when pursuing approach goals. Specifically, when goals have different deadlines, people tend to be more motivated to prioritize the goal with the shorter deadline. This pattern is consistent with the perspective associated with the time gradient, and suggests that this gradient is negative. When goals both have relatively short deadlines, people tend to be more motivated to prioritize the goal that is closer to being achieved, consistent with the negative distance gradient associated with the proximity perspective. When goals both have longer deadlines, the overall preference for prioritizing a particular goal is less clear. The fact the people's preference for prioritizing the goal that is closer to being achieved changes as a function of the deadline suggests that the effects of distance to goal and time to deadline interact in a manner consistent with the rate of progress gradient.

Avoidance Condition

The results for the avoidance condition (the right set of panels) suggest a much simpler pattern: People tend to make choices that allow them to maintain their distance from the undesired threshold that is closer to being reached. As a result, they tend to move closer to the undesired threshold that is farther from being reached. This can be seen by the fact that the vectors tend to point inward toward the main diagonal of each panel. In contrast to the approach context, motivation when pursuing avoidance goals appears much less strongly affected by deadline. This can be seen by the fact that all nine panels show very similar patterns of behavior. These results suggest that motivation is higher for the goal for which the undesired state is closer to being reached, suggestive of a negative distance gradient and consistent with the proximity and discrepancy perspectives.

Computational Modeling Results

First, we compare a series of models that represent different combinations of gradients. This allows us to examine which gradients are necessary to account for the dynamics of motivation. Next we use model parameter estimates to quantify the distance, time, and rate of progress gradients. This allows us to answer the six research questions by examining the nature and relative influence of each gradient.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Figure 3

Vector Field Plots Illustrating the Influence of Distance to Goal, Time to Deadline, and Goal Type on Prioritization Decisions Across the Three Experiments in Study 1

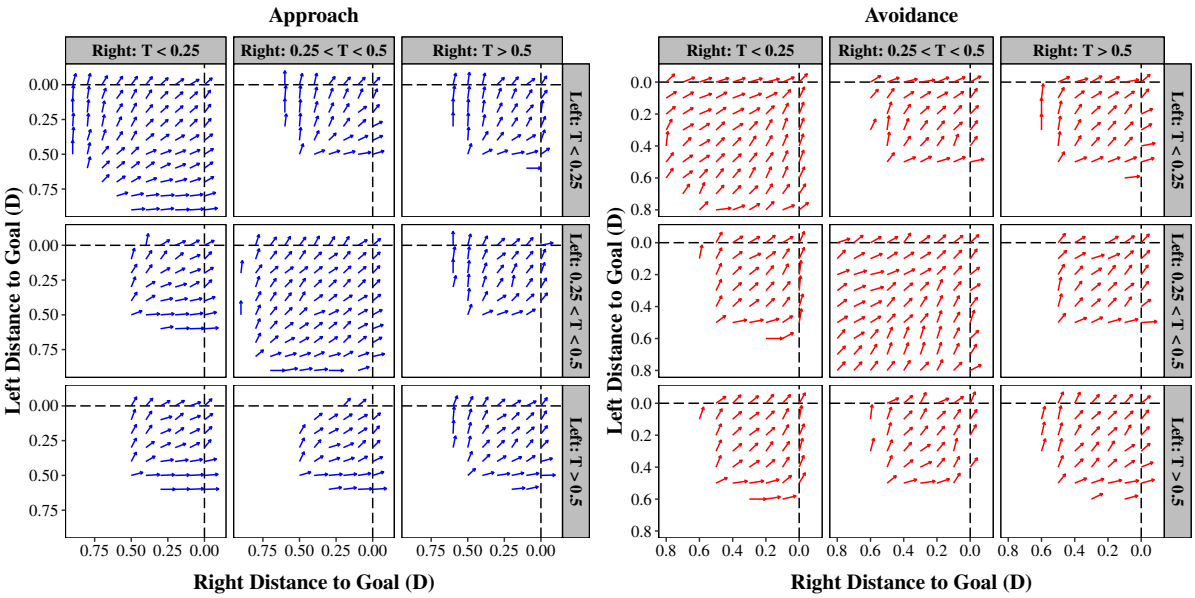

Note. The time to deadline for each goal is broken up into three ranges (<.25, 0.25–0.5 and >0.5) and mapped to different panels, with the columns representing the time to deadline for the right-hand goal and the rows representing the time to deadline for the left-hand goal. The distance to the right- and left-hand goal rounded to the nearest tenth is expressed on the *x* and *y* axes, respectively. The dashed lines at $D = 0$ indicate the threshold that determines whether the goal is achieved or failed. The orientation of the vector indicates the direction which the participants choose to move. Note that only cells containing 10 or more observations are shown. See text for more detail. See the online article for the color version of this figure.

The models were implemented in a hierarchical Bayesian framework using the Stan programming language (Carpenter et al., 2017).[4] The models calculated the predicted motivational value of each goal for every observation according to Equations 1–4, and transformed the difference in motivational values into a set of choice probabilities using Equation 5. The models were fit to the individual decisions separately for the approach and avoidance conditions. However, within each condition, the models assumed that the estimated parameters did not change across levels of the distance to goal and time to deadline manipulations. The models were fit to the data from all three experiments simultaneously so that parameter estimates would be informed by all the available data.[5]

Model Comparison

We compared seven different models (see Table 2), the full model and six special cases constructed by constraining different combinations of parameters to zero. Three of the models assumed that only one gradient was necessary to explain the motivational dynamics at play in this task. The first of these assumed that only the distance gradient was influential (instantiating the proximity and discrepancy perspectives), the second assumed that only the time gradient was influential (instantiating the time perspective), and the third assumed that only the rate gradient was influential (instantiating the expectancy, difficulty, and achievability perspectives). The fourth, fifth, and sixth models assumed that two gradients are necessary to explain the motivational dynamics. Finally, the seventh model assumed that all three gradients were necessary.

In models with a time gradient it was constrained to be decreasing by restricting τ to negative values, in line with the assumption that motivation should decrease with longer deadlines. In models with a distance gradient, the δ parameter could be positive or negative because both signs are theoretically plausible; a positive sign is assumed by the discrepancy perspective whereas a negative sign is assumed by the proximity perspective. Experiment 1a provided no information about the time gradient because the two goals always had the same deadline. Likewise, Experiment 1b provided no information about the distance gradient because the two goals

[4] We ran eight Markov Chain Monte Carlo chains with unique starting values randomly generated by Stan. Each chain had a burn in period of 4,000 samples, after which 10,000 samples were collected, so the final analyses was based on 80,000 samples (e.g., eight chains × 10,000 samples per chain). The chains demonstrated good mixing and there was very little autocorrelation in the final samples. This suggests that the samples provided a good approximation to the models' posterior distributions (Kruschke et al., 2012).

[5] We conducted a parameter recovery analysis to verify that our sample size and design were sufficient to enable reliable inferences regarding the parameters of interest. A parameter recovery analysis assesses the reliability with which the model parameters can be estimated given a study design and sample size (Heathcote et al., 2015). To do this, we generated 100 simulated data sets for each participant using known parameter values and examined the extent to which the estimated parameters from models that were fit to these simulated data sets correspond to the data-generating parameter values. For the parameters we interpret in this paper, the results revealed a high degree of correspondence between data-generating and estimated values, suggesting that our inferences based on those parameters are reliable. We report the full details of the parameter recovery analysis in the Supplemental Materials.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Table 2
Summary of Models Compared in Study 1

Model	Distance		Time		Rate	
	w_1	δ	w_2	τ	w_3	α
Distance	+	+/−				
Time			+	−		
Rate					+	+
Distance + Rate	+	+/−			+	+
Time + Rate			+	−	+	+
Distance + Time	+	+/−	+	−		
Distance + Time + Rate	+	+/−	+	−	+	+

Note. The + symbol indicates parameters constrained to be positive. The − symbol indicates parameters constrained to be negative. The +/− symbol indicates parameters that can be positive or negative. The gradients relevant to each set of parameters are specified above each pair of columns. Parameters that do not have symbols are constrained to be zero, removing the influence of the relevant gradient.

always started the same distance from the goal. As such, in all models, we fixed the time gradient weight (w_2) and shape (τ) to zero for participants from Experiment 1a and the distance gradient weight (w_1) and shape (δ) to zero for participants from Experiment 1b, eliminating the influence of that gradient for those participants. More details about the model and the priors on the model parameters can be found in the Appendix.

We compared the models using the Deviance Information Criterion (DIC; Spiegelhalter et al., 2002, 2014), which accounts for the tradeoff between model fit and parsimony, avoiding overfitting, and hence poorer out-of-sample prediction, through a complexity penalty that counters the ability of more flexible model to provide a better description of the data. As can be seen in Table 3, in the approach condition the best model by far was the one that assumed that all three gradients were necessary. In order to examine the evidence in the approach condition for each gradient at the individual level, we calculated each model's DIC for each participant. In the approach condition, a model that included the distance gradient was selected for 89% of participants, for 81% of participants a model that included the time gradient was selected, and for 88% of participants a model that included the rate gradient was selected.

In the avoidance condition, the model with the lowest DIC assumed that the distance and rate gradients were necessary, but the time gradient was not. However, this model was virtually indistinguishable from a model that assumed that all three gradients were necessary. In the avoidance condition, 76% of participants were best fit by a model that included the distance gradient, 36% of participants were best fit by a model that included the time gradient, and 74% of participants were best fit by a model that included the rate gradient.

Prompted by a reviewer, we also examined results based on the Watanabe–Akaike information criterion (Watanabe, 2010). They closely conform to the DIC results, with the exception that the Distance + Time + Rate model performs slightly better than the Distance + Rate model in the avoidance condition (full details of these analyses are available in the Supplementary Material).

The results of the model comparisons provide clear support for the need for both distance and rate gradients in both approach and avoidance goals, and thus strong evidence for the need for theoretical integration. The evidence for the time gradient is weaker for avoidance goals, but even there the individual difference results suggest that the time gradient is important for a substantial minority of participants. However, on the basis of these results, we cannot make inferences about the direction of certain gradients (e.g., whether the distance gradient is positive or negative). To do so, in the next section interpret the parameter estimates from the full model that includes all three gradients. We used this model even in the avoidance condition in order to obtain a more complete picture of the gradients at the individual level. The weight parameters of this model also provide another way to quantify the importance of each gradient, allowing us to obtain more fine-grained evidence about the relative influence of the time gradient and the other two gradients.

Before doing so, we evaluate the fit of the full model in an absolute sense. Figure 4 shows the relationship between the observed and predicted mean proportion of decisions prioritizing the right-hand goal according to this model across the various conditions in the three experiments. As can be seen, the three-gradient model captures the observed data quite well. The 95% credible interval on the correlation between observed and predicted means for each Times × Distance condition ranges from 0.67 to 0.85

Table 3
Results of Model Comparisons for Experiments 1a, 1b, and 1c

Goal type	Statistic	Models							Statistic	Gradients		
		D	T	R	D + R	T + R	D + T	D + T + R		Distance	Time	Rate
Approach	Overall DIC	26,448	16,965	21,838	16,650	4,290	10,364	0	Overall % using	89	81	88
	Exp 1a % best fit	9		8	82				Exp 1a % using	91		90
	Exp 1b % best fit		4	4		91			Exp 1b % using		95	95
	Exp 1c % best fit	18	8	6	28	4	10	26	Exp 1c % using	82	48	64
Avoidance	Overall DIC	4,366	15,956	3,171	0	2,513	4,099	14	Overall % using	76	36	74
	Exp 1a % best fit	35		21	44				Exp 1a % using	79		65
	Exp 1b % best fit		11	63		25			Exp 1b % using		36	88
	Exp 1c % best fit	16	12	8	42	14	4	4	Exp 1c % using	66	34	68

Note. The columns under the Models heading show the Overall DIC for each model and the percentage of participants' best fit by each model according to the DIC in each experiment. D, T, and R indicate that the model includes the distance, time, and rate of progress gradients, respectively. To facilitate interpretation, the Overall DIC is expressed as a difference between each model and the best performing model in that condition. The columns under the Gradients heading show the percentage of participants for whom the best fitting model included each gradient. DIC = Deviance Information Criterion.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Figure 4

Relationship Between Observed and Predicted Prioritization Decisions in Study 1

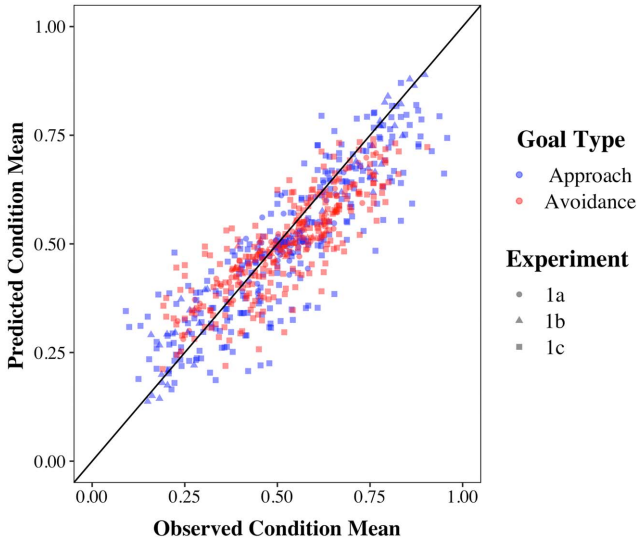

Note. Each point represents the mean proportion of decisions prioritizing the right-hand goal off or a single condition. The condition means were computed by calculating the proportion for that trial for each participant and averaging across participants. This ensures that trials with more decisions did not exert a disproportionate influence on the mean value. See the online article for the color version of this figure.

over the six sets of results (three experiments broken down into approach and avoidance conditions) represented in the graph.

Parameter Analysis

The main focus of the following analysis is on using the parameters of the full model to compute distance, time, and rate gradients (*DG*, *TG*, and *RG*, respectively) in order to answer our six research questions. Before doing so we note that there was a difference between the approach and avoidance conditions in the process of mapping these gradients to choices. Sensitivity to the difference in motivational value between the two goals (i.e., in the average across participants of θ, see Equation 5) was larger in the approach condition, with a 95% credible interval (CI) of [9.25, 9.81], compared with the avoidance condition (95% CI: [4.21, 4.59]), with the 95% CI on average difference clearly not overlapping zero [4.80, 5.47]. This finding suggests that individuals were more sensitive in general to the goal's motivational value when pursuing approach goals than when pursuing avoidance goals.

Figure 5 presents the results of the gradient analysis. The top row of the figure shows the effects of distance to goal and time to deadline on motivational value, which represents the sum of the effects of the distance, time, and rate gradients (with values between 0 and 1, as defined by Equation 4). These plots are equivalent in structure to Figure 1, and therefore can be used to compare the results to the predictions of the six theoretical perspectives. Clearly, the pattern in the approach condition is not consistent with any of the six perspectives shown in Figure 1, suggesting that no one perspective is sufficient to account for these dynamics in its own right. The overall motivational value of an approach goal changes in response to both distance to goal and time to deadline. The motivational value is higher when the goal is relatively close to being achieved and when the deadline is relatively near (i.e., the bottom-left corner of the panel), and decreases as either distance to goal or time to deadline increases. The result is that motivation is the lowest when the goal and deadline are both far away (i.e., the top-right corner).

The overall pattern in the avoidance condition is also not consistent with any of the six theoretical perspectives shown in Figure 1. The results suggest that the motivational value of an avoidance goal also changes as a function of distance and time, but in a different way. Motivation to avoid is the highest when the undesired state is relatively close and the deadline is relatively near (i.e., the bottom-left corner of the panel). In the avoidance context, however, motivation is the lowest when the undesired state is farther away and the deadline is relatively near (i.e., when the goal is very easy to achieve). Taken together, the effects of distance to goal and time to deadline on motivational value reinforce the notion that an integrative framework is needed to account for the motivational dynamics observed in the approach and avoidance conditions.

The second, third, and fourth rows of Figure 5 show the contribution of the distance, time, and rate gradients to the overall motivational value. These plots allow us to ascertain more definitively which theoretical perspectives are necessary to account for the behavior in this task and to answer the research questions regarding the nature and relative strength of the various gradients.

The first research question in Table 1 addresses the nature of the distance gradient in the approach context. Is the distance gradient negative, as predicted by the discrepancy perspective, or positive, as predicted by the proximity perspective? This question is answered in the second row of Figure 5, which shows the distance gradients for each participant. These gradients were calculated based on the w_1 and δ parameters for each participant in Experiments 1a and 1c. In the approach condition, 91% of gradients are negative (shown in red), decreasing with distance to goal, and in the avoidance condition, 93% of distance gradients are negative. These findings provide evidence in favor of the proximity perspective.

The second research question addresses the relative strength of the distance gradient when pursuing approach versus avoidance goals. This is addressed by comparing the value of the w_1 parameter, which quantifies the relative influence of distance gradient compared to the other two gradients, between the approach and avoidance conditions. Figure 6 summarizes the posterior distributions on each model parameter in each goal type condition. The 95% CI of the difference in the w_1 parameter between conditions—computed as the average of w_1(approach)−w_1(avoidance) across participants—is [−0.16, −0.10]. This indicates that the relative influence of the distance gradient is credibly stronger in the avoidance condition than in the approach condition. This result is consistent across experiments (Expt 1a CI: [−0.15, −0.09]; Expt 1c CI: [−0.23, −0.08]).[6]

[6] One might argue that a test of differences in proximity effects across conditions should exclude participants who did not show proximity effects (those who had positive distance gradients). When the 14% of participants who had positive gradients in at least one goal type context were excluded, the same results emerged (Overall CI: [−0.17, −0.10]; Expt 1a CI: [−0.16, −0.10]; Expt 1c CI: [−0.226, −0.074]).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Figure 5

Estimated Overall Motivational Value (Top Row) and Distance (Second Row), Time (Third Row), and Rate (Bottom Row) Gradients in Study 1

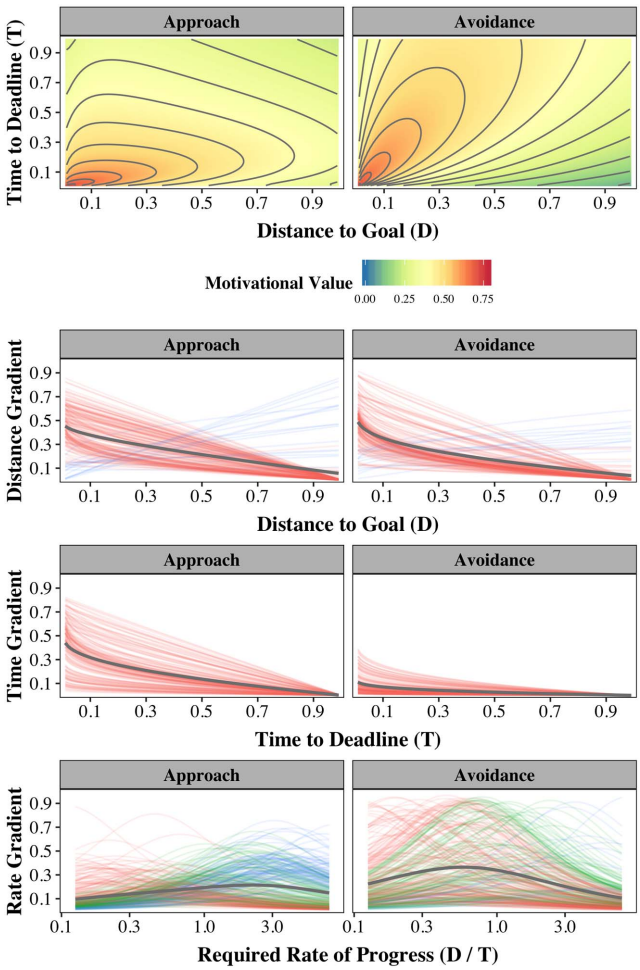

Note. The motivational values shown in the surface plots represent the mean of the posterior predictive distribution on the average motivational value across participants. Distance, time, and rate gradients are shown for each participant, and represent the mean of the posterior predictive distribution on the gradient strength for each individual. Distance and time gradients that are a decreasing function of distance to goal or time to deadline are shown in red. Distance gradients that are an increasing function are shown in blue. Rate gradients shown in blue increase with the rate of progress required to achieve the goal, rate gradients shown in red decrease with the rate required, and those shown in green are nonmonotonic. Distance and rate gradients were classified by identifying the category that was most probable given the posterior distributions on the δ and α parameters. See the online article for the color version of this figure.

The third research question addresses the relative strength of the time gradient when pursuing approach versus avoidance goals. The time gradients are shown in the third row of Figure 5. These gradients were calculated based on the w_2 and τ parameters for each participant in Experiments 2 and 3. This research question is answered by examining the value of the w_2 parameter, which quantifies the relative influence of the time gradient. The 95% CI on the difference in the w_2 parameter between conditions (computed as the average of w_2(approach)−w_2(avoidance) across participants) is [0.32, 0.57]. This indicates that the relative influence of the time gradient was credibly, and substantially, stronger in the approach condition than in the avoidance condition. Again, this result is consistent across experiments (Expt 1b CI: [0.41, 0.71]; Expt 1c CI: [0.10, 0.28]).

The fourth research question addresses the prevalence of the three different types of rate of progress gradients. Is the relationship between the rate of progress required and motivation consistent with the expectancy, difficulty, or achievability perspectives? The rate gradients are shown in the fourth row of Figure 5. These gradients were calculated based on the w_3 and α parameters for each participant across all three experiments. The *x*-axis for these panels represents the rate of progress required to reach the goal, which corresponds to the ratio of the distance to goal and time to deadline. We classified participants' gradients as being increasing, decreasing, or nonmonotonic functions of the required rate of progress based on the shape of their rate gradient across the rates that were required to achieve the goals in the task. Rate gradients that were increasing (such that increases in the required rate at that state would result in an increase in the gradient) in at least 75% of task states on which all participants' decisions were based were classified as "increasing." Rate gradients that were decreasing in at least 75% of states were classified as "decreasing." Rate gradients were otherwise classified as "nonmonotonic."[7]

As can be seen in Figure 5, there was substantial variation in the nature of the rate gradient between individuals in the approach condition. In this condition, 31% of people had nonmonotonic gradients (shown in green), suggesting that, consistent with the achievability perspective, they were most motivated by goals that required a moderate rate of progress in order to achieve. 33% of participants had gradients that increased with the required rate of progress (shown in blue), suggesting that, consistent with the difficulty perspective, they were most motivated by the hardest goals. 36% of participants had rate gradients that decreased with the required rate (shown in red), suggesting that, consistent with the expectancy perspective, they were most motivated by the goals that were easiest to achieve.

In the avoidance condition, the balance of rate gradients was somewhat different to the approach condition. 35% of participants had nonmonotonic rate gradients, consistent with the achievability perspective. 59% had gradients that decreased, which is consistent with the difficulty perspective, as in the avoidance condition a decreasing gradient indicates that motivation is the highest for the hardest goals (i.e., goals for which the undesired state is difficult to avoid). Only 6% of participants in the avoidance condition had

[7] On the domain of 0 to infinity, the rate gradient only ever becomes monotonic as α approaches 0 or 1. However, a rate gradient that is nonmonotonic in this domain may be *effectively* monotonic in the experimental paradigm if it reaches its peak at rates of progress that are extreme compared to those encountered in the experiments. In classifying participants' rate gradients there is, therefore, a tradeoff between (a) avoiding inappropriately classifying participants as having nonmonotonic gradients when their rate gradient peaks at an extreme value that is rarely encountered in the experiment, in which case, their gradient is effectively monotonic; and (b) avoiding inappropriately classifying participants as having increasing or decreasing gradients when there is clear nonmonotonicity. The 75% rule for classifying gradients as monotonic is, ultimately, an arbitrary cutoff. However, we believe it strikes a good balance between (a) and (b). Further details on the distribution of rate gradients across participants can be found in the Supplemental Material.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Figure 6

Posterior Distributions on the Seven Model Parameters in Each Goal Type Condition in Study 1

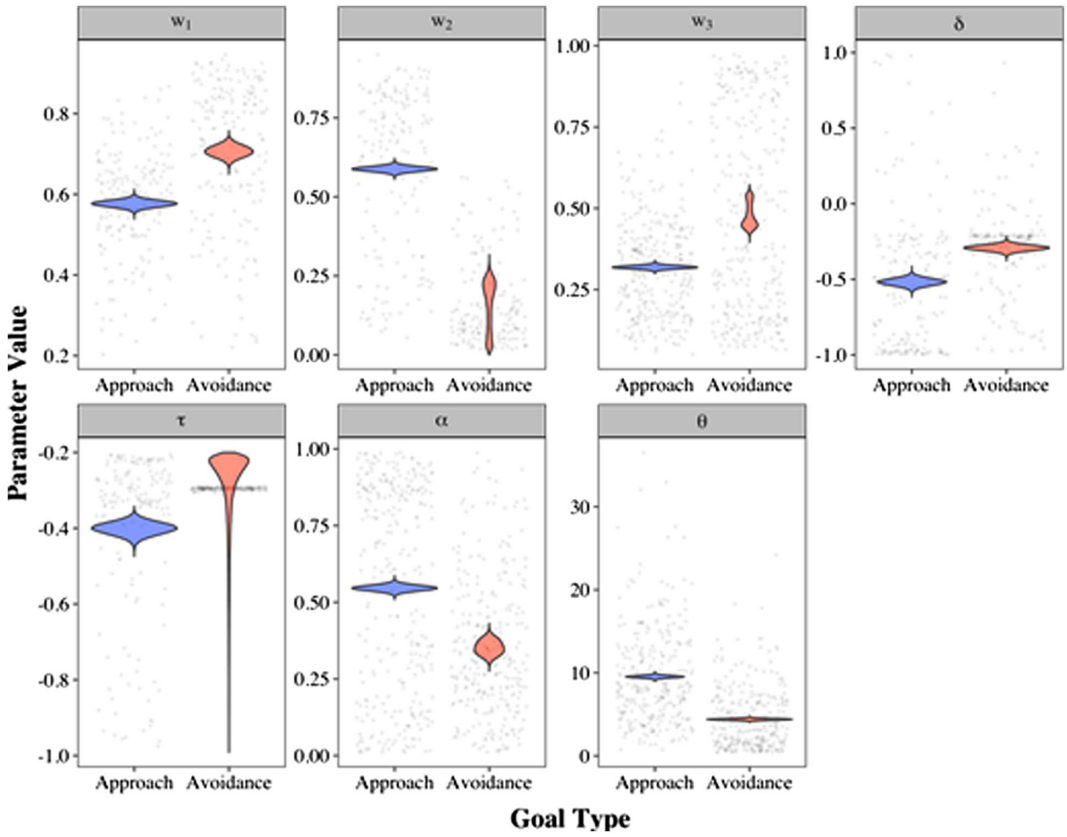

Note. The violin plots represent the posterior distribution on the mean value for each parameter when aggregated across participants. The jittered points represent the mean of each participants' posterior distribution. See Supplementary Material for a breakdown of the posteriors across the three experiments. See the online article for the color version of this figure.

gradients that increased with the required rate of progress. These findings suggest that the achievability and difficulty perspectives account for the behavior of almost all participants, with the expectancy perspective being required for at best a small minority of participants in the avoidance context.

The fifth research question asks whether any observed increase in motivation as the distance to the goal decreases is attributable to proximity effects or to the goal becoming easier to achieve. The preceding analyses have confirmed the presence of proximity effects (even when controlling for possible expectancy effects via the rate gradient). However, the presence of decreasing and nonmonotonic rate gradients also indicates that part of the increase in motivation in response to a looming goal may be attributable to the goal becoming more achievable. We can assess the relative contribution of these effects by evaluating the difference between w_1 and w_3 parameters across participants in Experiments 1a and 1c. The 95% CI on this difference (computed as the average of $w_1 - w_3$ across participants) is [0.22, 0.31] in the approach condition and [0.38, 0.48] in the avoidance condition. The results of this analysis were consistent across experiments (Approach Expt 1a CI: [0.074, 0.215]; Approach Expt 1c: [0.30, 0.41]; Avoidance Expt 1a CI: [0.38, 0.50]; Avoidance Expt 1c: [0.33, 0.51]).[8] This result indicates that, although both proximity and expectancy effects are present, the change in motivation observed as the goal draws nearer is more strongly influenced by the goal's proximity than the expectancy of achieving the goal.

Finally, the sixth research question asks whether any increase in motivation as the deadline becomes closer is due to the effects of the deadline itself or an increase in the rate of progress required to achieve the goal. Answering this question requires a comparison of the time and rate gradients. The model comparisons suggested that both time and rate effects were present when pursuing approach goals, but that time effects may have been weaker or absent in the avoidance context. We can quantify the relative contribution of these effects by evaluating the difference between w_2 and w_3 parameters across participants in Experiments 1b and 1c. The 95% CI on this difference (computed the average of $w_2 - w_3$ across participants) was [0.30, 0.38] in the approach condition and [−0.75, −0.29] in the avoidance condition. The results were consistent

[8] We excluded from this analysis all participants who had positive distance gradients and/or rate gradients that increased monotonically as a function of goal difficulty, as such gradients always produced decreases in motivation as distance to the goal reduces. Thus, it was not appropriate to evaluate the fifth research question using these participants. After these exclusions, 40% of participants in each condition were retained for this analysis.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

across experiments (Approach Expt 1b (CI: [0.46, 0.54]), Approach Expt 1c (CI: [0.010, 0.167], Avoidance Expt 1b CI: [−0.95, −0.38], Avoidance Expt 1c CI: [−0.256, −0.050]).[9] This result indicates that, when pursuing approach goals, the change in motivation observed as the deadline looms was more strongly influenced by the deadline itself compared to the goal's difficulty. However, when pursuing avoidance goals, this change was more strongly influenced by the goal's difficulty.

Discussion

In Study 1, we used the GOAL architecture to compare six theoretical perspectives on goal-directed motivation using data from an experimental task in which participants made a series of prioritization decisions in an attempt to achieve different types of goals. Our results suggested that the distance, time, and rate of progress gradients were all influential. We found that five of the six perspectives—all except the discrepancy perspective—need to be integrated in order to account for the full complexity of the motivational dynamics at play in this task. The proximity perspective accounted for the observed increase in motivation as participants got closer to reaching a goal. The time gradient accounted for the motivating effects of the looming deadline when pursuing approach goals, but less so when pursuing avoidance goals. The achievability and difficulty perspectives accounted for individual differences in how people responded to the rate of progress required to reach the goal, with the expectancy perspective also being needed in the approach context. These findings support the argument that an integrative theoretical framework is needed in order to develop a more comprehensive understanding of goal-directed motivation and decision-making.

Study 2

In Study 2, we use data from the National Basketball Association (NBA) to examine the robustness of the inferences generated in Study 1 in a field context. The NBA consists of 30 teams that each play 82 games over the course of a season. Each game consists of four, 12-min quarters, with the winning team being the one that has scored the most points at the end of the game.[10] An NBA game involves both approach and avoidance goals. When a team has fewer points than the opposing team they need to eliminate the gap between their score and their opponent's score (an approach goal). When a team has more points than the opposing team they need to prevent the gap between their score and their opponent's score from narrowing (an avoidance goal). The distance to the goal can be expressed as the absolute value of the difference between the teams' scores (see the analysis of professional ice hockey data by Beck et al., 2019, for a similar operationalization of distance to goal). The time to the deadline can be expressed as the number of minutes remaining in the game.

Following Beck et al. (2019), we operationalize the motivation to expend effort in terms of the amount of workload allocated to a team's strongest players. In an NBA game, each team has five players on the court at a given time, with each team having up to eight additional reserve players. The five strongest players on the team generally begin the game on the court. These five players are referred to as the "starters." Teams are allowed an unlimited number of substitutions in a game, and coaches regulate players on-court time strategically in order to maximize the chances of winning the game while minimizing player fatigue and the risk of injury. In this study, we operationalize the coach's motivation to expend effort as the proportion of the five starters that are on the court at a given point in time. All else being equal, playing the starters generally increases the rate at which a team will score points and decreases the rate at which the opposing team can score points. At the same time, playing the starters increases the risk of these players becoming fatigued or experiencing an injury.

In addition to providing an opportunity to evaluate the research questions introduced earlier in a field setting, the NBA context allows us to examine the role of goal importance. To what extent does the importance of achieving the goal influence the distance, time, and rate of progress gradients? Theoretically, there are a number of ways in which goal importance may affect these dynamics. First, importance may influence the overall level of sensitivity to the goal's motivational value, effectively amplifying changes along any of the three gradients in the same way (Ballard et al., 2018; Vancouver et al., 2010;). Second, importance may affect the relative influence of the gradients on a goal's motivational value, causing certain gradients to become more influential and others to become less so (Estle et al., 2006; Schmidt & DeShon, 2007). Finally, importance may exert an effect that is independent of the three gradients (Locke & Latham, 1990), for example, by increasing the resources allocated by a constant amount that is invariant to the position along the gradients. The effects of goal importance can be examined by comparing the motivational dynamics in games played during the 82-game "regular season" to those from preseason and playoff games. Preseason games are lower stakes because they do not count toward a team's season record. Games played during the playoffs—the end-of-season tournament that determines that year's champion—are higher stakes because a win can advance a team closer to the championship and a loss can bring a team's season to an end.

Data

The analysis was conducted using data freely available from the NBA. These data contain play-by-play information for nearly every NBA game since 1997, with variables such as the game score, time remaining, and any shot attempts, fouls, or substitutions recorded for each possession. The data were downloaded in May, 2020 using the nbastatR package (Besler, 2020), which provides an R interface to the NBA.com data repository. In total, we analyzed data from 28,136 games played between 1997 and 2019.

The analysis was conducted on data from the fourth quarter of each game. We chose to analyze only the fourth quarters because substitutions in the first three quarters are often determined by a preplanned rotation and are less influenced by the score.

[9] We excluded from this analysis all participants whose rate gradients that decreased monotonically as a function of goal difficulty as such gradients always produce a decrease in motivation as the deadline looms. Thus, it was not appropriate to evaluate the sixth research question using these participants. After these exclusions, 44% of participants were retained for the analysis of the approach condition and 93% were retained for the analysis of the avoidance condition.

[10] If teams are tied at the end of the fourth quarter, an overtime period is played. We did not analyze overtime periods in this study.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

The fourth quarter is when coaches become more strategic in how they regulate their starters' playing time, increasing or decreasing starters' on-court time depending on how the game unfolds. Substitutions made in the fourth quarter are, therefore, more diagnostic of the underlying motivational dynamics than substitutions made in earlier quarters.

The outcome variable of interest in this study is the number of seconds in a given minute of play which were occupied by starters. We refer to this variable as *starters' on-court time*. Because a team has five players on the court at any given time, there are 300 s of playing time in each minute of play (5 players × 60 s). If all five starters are on the court for the entire minute of play, the starters' on-court time would be 300 s for that minute. If no starters were on for the entire minute, this variable would have a value of zero seconds. Because substitutions can happen at any time, the starters' on-court time variable can take on any integer value between zero and 300 s. The distance to goal and time to deadline variables were operationalized as the absolute value of the score difference and the number of minutes remaining, respectively, at the start of that minute of play. Each game provided 24 observations (12 min of fourth quarter play × 2 teams), which resulted in a total of 675,264 data points.

We used the beta-binomial distribution to transform the motivational value predicted by the GOAL architecture into a prediction regarding the starters' on-court time. We chose the beta-binomial in favor of a simple binomial distribution because it allows us to better account for noise in the decision-making process. For example, coaches differ from each other, and some coaches might be more consistent than others in how they respond to particular situation. Similar to the logistic model used in Study 1, the beta-binomial directly models this noise and isolates it from the systematic variance attributable to distance, time, and the rate of progress required. This enables clearer conclusions to be made about the motivational dynamics. The starters' on-court time in minute j (y_i) is assumed to be beta-binomially distributed, with a maximum of 300 possible seconds of play. The model is formalized below in Equations 6–8 (which, for simplicity, represent a model of a single game).

$$y_j \sim \text{betabinomial}(300, A_j, B_j), \tag{6}$$

$$A_j = M_j \theta \kappa, \tag{7}$$

$$B_j = \frac{M_j \theta}{\kappa}. \tag{8}$$

M_j represents the motivational value predicted by the GOAL architecture (according to Equation 4) at the start of minute j, and κ is a bias parameter that captures the coaches' overall tendency to allocate greater or fewer resources toward goal achievement. This parameter is constrained to be positive, with values greater than one indicating a bias toward playing the starters and values less than one indicating a bias toward not playing them. The θ parameter controls how sensitive coaches are to the motivational value of the goal. This parameter is also constrained to be positive, with higher values of indicating greater sensitivity to changes in motivational value, and so it has a similar interpretation to the θ parameter in Study 1.[11] For brevity, we report results for only full model that allowed all three gradients, and focused our analysis on the parameters of this model, which revealed that all three were influential.

Results

Figure 7 shows the average number of starters on the court in preseason (left column), regular season (middle column), and playoff games (right column) as a function of the number of minutes remaining in the game and the score difference. The top row of panels shows the observed data. As can be seen, in preseason games, which are relatively low in importance, the starters' on-court time is generally low and mostly unrelated to the score difference or the time remaining in the game. In regular season games, which are more important, the starters' on-court time is regulated more strategically. The starters receive more playing time when the score difference is low (i.e., close to zero) and when there is less time remaining in the game. The yellow diagonal contour separating the orange and green/blue regions (which is particularly clear when there are fewer than 7 min remaining) shows that coaches integrate information about the distance to the goal and time remaining when making decisions about how much of the starters' energy to expend. Specifically, as the time remaining in the game decreases, a smaller score difference is required in order to justify keeping the starters in the game. In playoff games, which are the most important, this same pattern holds, except that a larger score difference is required before coaches are willing to remove their starters from the game.

The bottom row shows the predictions of the model after having been fit to the data (using the same procedure as in Study 1). Unique parameters were estimated for each team, with the parameters associated with each team being drawn from a common population distribution.[12] Parameters associated with the approach context were estimated by fitting the model to observations in which the team had fewer points than its opponent (and therefore needed to narrow the gap in scores). Parameters associated with the avoidance context were estimated by fitting the model to observations in which the team had more points than its opponent (and therefore had to prevent the gap from narrowing). The preseason, regular season, and playoff games were also analyzed using separate models. As can be seen, the model accurately captures the qualitative effects of score difference, time remaining, and game type on the starters' on-court time. However, the model does appear to underestimate the magnitude of the effects of these variables in regular season and playoff games, producing less extreme effects than are shown in the data.

Parameter Analysis

The gradients from Study 2 are shown in Figure 8. This figure has the same structure as Figure 5, with the exception that it is broken down by game type. We begin by discussing the gradients associated with the regular season games, as these games represent the typical circumstances under which teams operate. In the next section, we consider differences between preseason, regular season, and playoff games in order to assess how the importance of success influences the gradients.

[11] A parameter recovery analysis confirmed that the parameters reported below are reliably estimated in this data set (see Supplemental Material for more details).

[12] For the purposes of our analysis, games associated with the same organization in different NBA seasons were considered to be played by different "teams". For example, the Chicago Bulls in the 1997–1998 season were modeled as a separate team from the Chicago Bulls in the 1998–1999 season.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Figure 7

Observed and Predicted Number of Starters On the Court as a Function of Score Difference, Time Remaining, and Game Type in Study 2

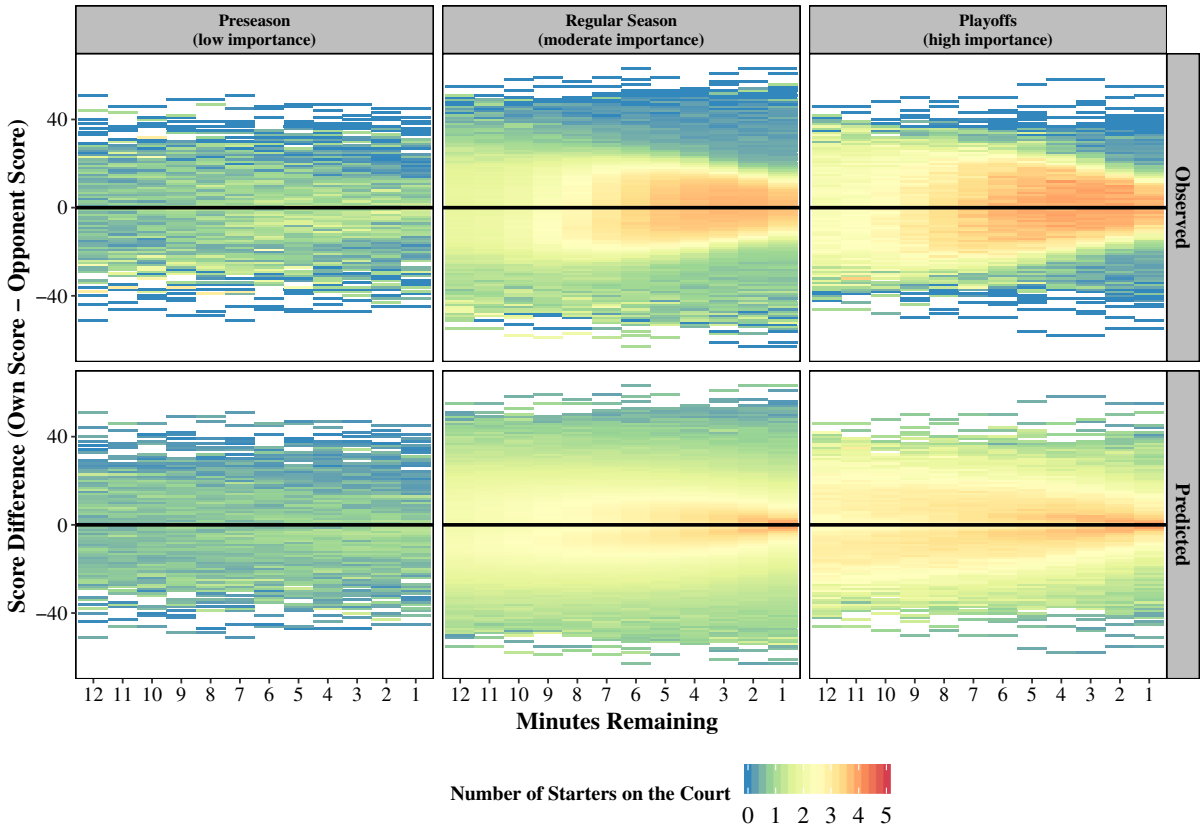

Note. The effects of score difference (mapped to *y*-axis), time remaining (mapped to *x*-axis), and game type (mapped to panel rows) on the number of starters on the court in a given minute of play (mapped to color) in NBA games between 1997 and 2019. The top row represents the observed values whereas the bottom row represents the values predicted by the GOAL architecture. See the online article for the color version of this figure.

As can be seen in the lower three rows of Figure 8, all three gradients play a role in influencing the regulation of starters' playing time in regular season NBA games. The top panel shows that the gradients combine to effect motivational value in a similar manner to in Study 1 in that the point of maximum motivation occurs when distance to goal and time to deadline are both low (i.e., the bottom-left-hand corner of each panel). To address the first research question in Table 1, we once again examined the direction of the distance gradients (second row of panels). In Study 2, all 653 teams had negative distance gradients in both approach and avoidance situations. This finding aligns with Study 1 and provides strong support for the proximity perspective.

To address the second research question, we assessed the relative strength of the distance gradient under approach versus avoidance in the same manner as in Study 1. Figure 9 summarizes the posterior distributions on the model parameters in each goal type condition. In Study 2, the 95% CI on the average difference in the strength of the distance gradient (w_1) between approach and avoidance contexts was [−0.054, −0.035]. These results are consistent with Study 1 and the goal gradient hypothesis, suggesting that the relative influence of the distance gradient is stronger during avoidance compared with approach.

To address the third research question, we assessed the relative strength of the time gradient (shown in the third row of Figure 8) under approach versus avoidance in the same manner as in Study 1. The 95% CI on the average difference in the strength of the time gradient (w_2) between approach and avoidance contexts was [−0.026, −0.013], These results are in the opposite direction to findings of Study 1, suggesting that the relative influence of the time gradient is slightly stronger during avoidance compared to approach.

The fourth research question addresses the nature of the rate of progress gradients (see the fourth row of Figure 8). We classified teams as having rate gradients that were increasing, decreasing, or nonmonotonic functions of the required rate of progress in the same manner as we did for participants in Study 1. There was less variation in the nature of the rate gradients in Study 2, with nonmonotonic gradients dominating. In approach situations, 95% of teams had a nonmonotonic rate gradient (shown in green), 4% of teams had a decreasing rate gradient (shown in red), and less than 1% had an increasing rate gradient (shown in blue). In avoidance situations, 97% of teams had a nonmonotonic rate gradient, less than 1% had a decreasing rate gradient, and 3% had an increasing rate gradient. These findings strongly support the achievability

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Figure 8

Estimated Overall Motivational Value (Top Row) and Distance (Second Row), Time (Third Row), and Rate (Bottom Row) Gradients in Preseason, Regular Season and Playoff NBA Games Analyzed in Study 2

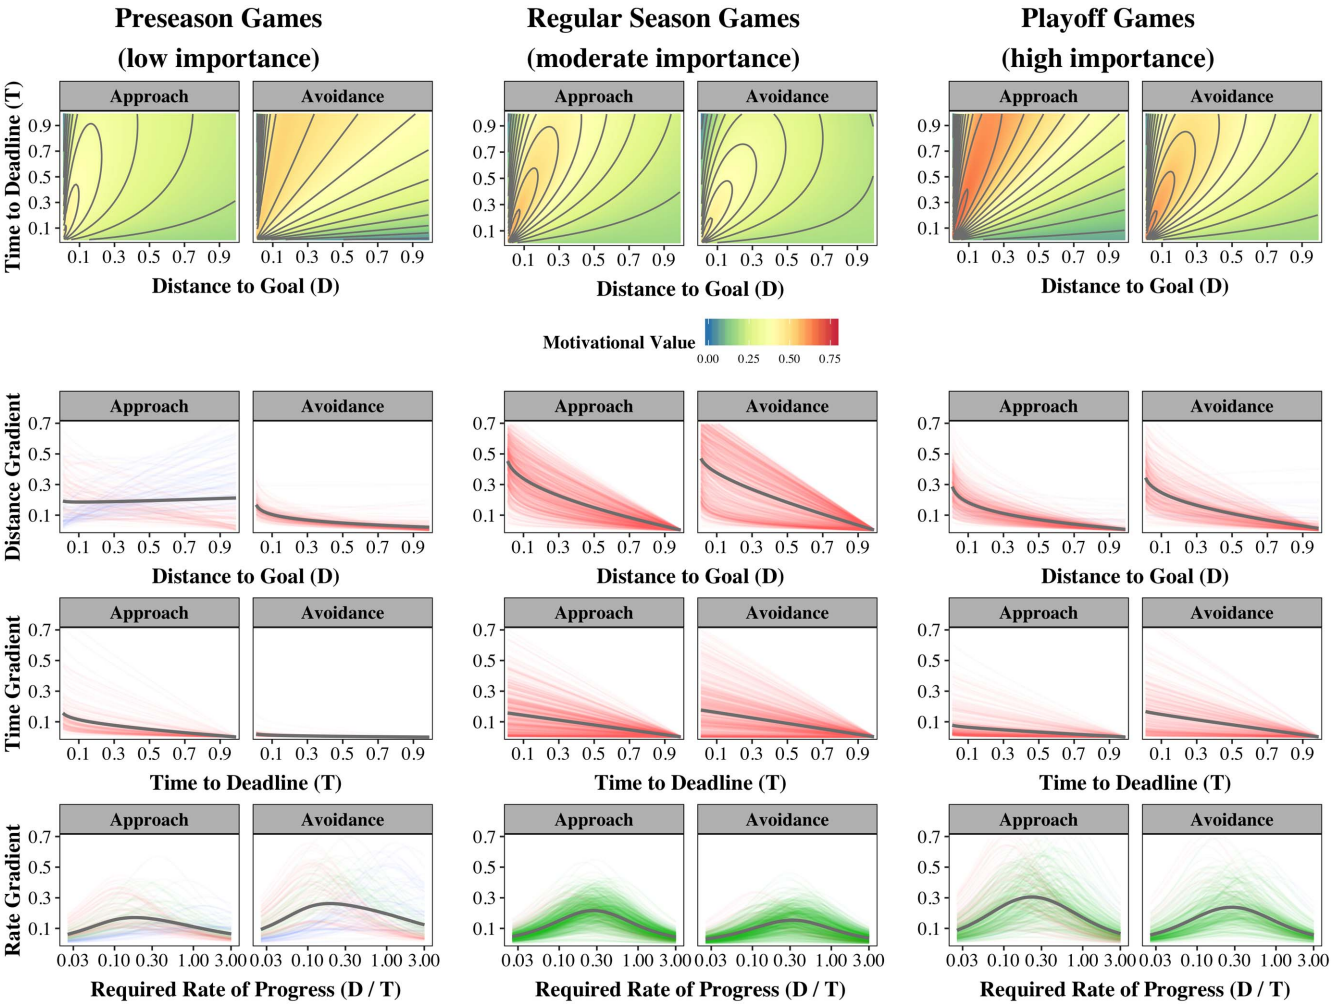

Note. The motivational value shown in the surface plots represents the mean of the posterior predictive distribution on the average motivational value across teams. Distance, time, and rate gradients are shown for each team, and represent the mean of the posterior predictive distribution on the gradient strength for each team. All distance or time gradients are a decreasing function of distance to goal or time to deadline and are shown in red. Rate gradients that increase with the rate of progress required to achieve the goal are shown in blue, rate gradients that decrease with the rate of progress required are shown in red, and that are nonmonotonic are shown in green. The distance and rate gradients were classified by identifying the category that was most probable given the posterior distributions on the δ and α parameters. See the online article for the color version of this figure.

perspective for both approach and avoidance goals, with greater motivation to play starters when the rate at which a team (in the approach context) or their opponent (in the avoidance context) needed to reduce the score difference was moderate than when it was higher or lower.

The fifth research question examines whether any increase in motivation as the distance to the goal decreases is attributable to the proximity or expectancy effects. As with Study 1, we assess the relative contribution of these effects by evaluating the difference between w_1 and w_3 parameters across teams.[13] The 95% CI on this difference (computed as the average of $w_1 - w_3$ across teams) was [0.37, 0.39] in approach situations and [0.47, 0.49] in avoidance situations. Consistent with Study 1, this result indicates that the changes in motivation observed as the score differences gets closer to zero is more strongly influenced by the goal's proximity than the expectancy of achieving it.

The sixth research question examines whether any increase in motivation as the deadline becomes closer is due to the effects of the deadline itself or an increase in the rate of progress required to reach the goal. As in Study 1, we answered this question by comparing the time gradients with the rate gradients. We did this by evaluating the

[13] As with Study 1, we excluded from this analysis the teams that had rate gradients that increased monotonically as a function of goal difficulty. There were no teams with positive distance gradients in this study so no further exclusions were necessary. After these exclusions, we retained >99% of teams for the analysis of the approach context and 96% of teams for the analysis of the avoidance context.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Figure 9

Posterior Distributions on the Nine Model Parameters as a Function of Goal Type and Game Type in Study 2

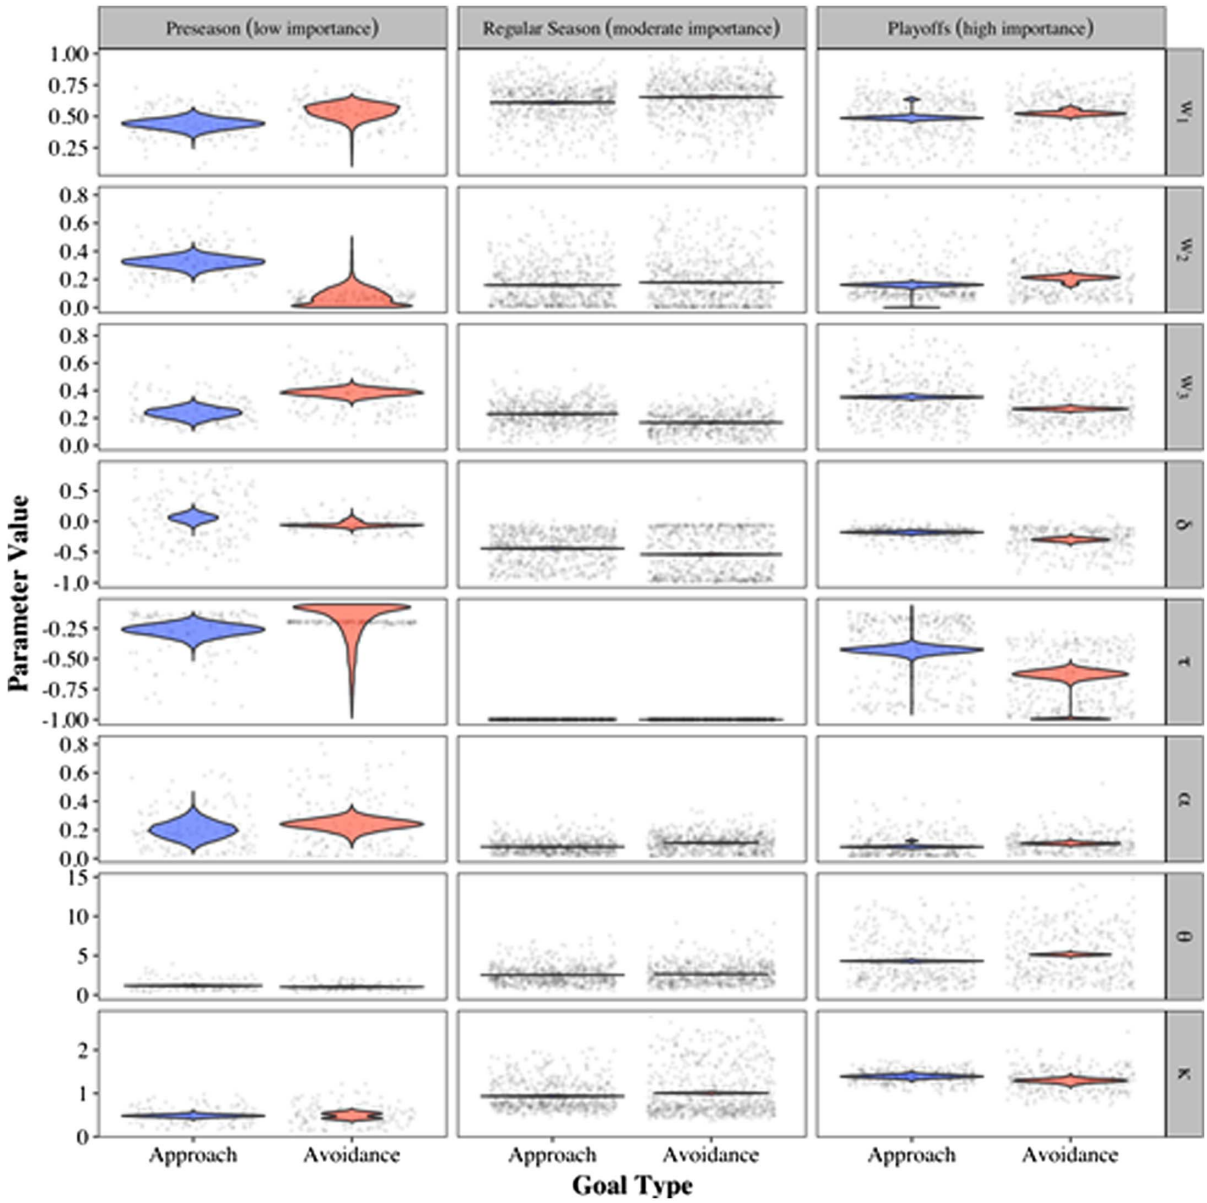

Note. The violin plots represent the posterior distribution on the mean value for each parameter when aggregated across teams. The jittered points represent the mean of each team's posterior distribution. See the online article for the color version of this figure.

difference between w_2 and w_3 parameters across teams.[14] The 95% CI on this difference (computed the average of $w_2 - w_3$ across teams) was [−0.071, −0.061] in approach situations and [0.007, 0.018] in avoidance situations. This result runs in the opposite direction to that observed in Study 1. It suggests that when teams are behind and working to catch up, the change in motivation as the end of the game looms is more strongly influenced by the difficulty of doing so rather than the deadline itself. By contrast, when teams are ahead and working to prevent the other team from catching up, the change in motivation is more strongly influenced by the looming deadline itself than the goal's difficulty.

Finally, we report the results regarding coaches' bias and overall sensitivity to the motivational value of the goal. The bias is measured by the κ parameter and captures the overall tendency to allocate greater or fewer resources to goal achievement independent of a goal's motivational value. The results suggested a small

[14] As in Study 1, we excluded from this analysis who had rate gradients that decreased monotonically as a function of goal difficulty as such gradients always produce an decrease in motivation as the deadline looms. After these exclusions, we retained 96% of teams for the analysis of the approach context and 99% of teams for the analysis of the avoidance context.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

bias toward allocating fewer resources (i.e., allocating the starters less playing time) in the approach context (95% CI: [0.92, 0.96]), with no indication of bias in the avoidance context (95% CI: [0.99, 1.03]). The difference in bias between the approach and avoidance context was small, but credible (95% CI on difference: [0.04, 0.10]). The θ parameter captures the overall sensitivity to the motivational value of the goal. In contrast to Study 1, overall sensitivity was slightly lower in the approach context (95% CI: [2.51, 2.55]) than in the avoidance context (95% CI: [2.63, 2.68]; 95% CI on difference: [0.094, 0.157]).

The Influence of Goal Importance

To examine the relationship between goal importance and the overall sensitivity to the gradients, we compared the θ parameter across less important preseason games, moderately important regular season games, and highly important playoff games. The results suggested the sensitivity to the gradients generally increased with goal importance. Overall sensitivity was higher in regular season games than in preseason games (95% difference CI: [1.42, 1.55], both in approach, [1.25, 1.43], and avoidance, [1.54, 1.72], situations) and higher in playoff games than regular season games (95% difference CI: [1.98, 2.26], again both in approach, [1.63, 1.92], and avoidance [2.24, 2.72], situations).

To examine the relationship between goal importance and the relative influence of the three gradients, we compared w_1, w_2, and w_3 parameters across game types. The analysis revealed a difference in the relative influence of the distance and rate gradients between regular season and playoff games. Overall, the distance gradient was more influential in regular season compared to playoff games (95% CI on overall difference in w_1: [0.050, 0.145]). When this difference was broken down by goal type, it was credible in avoidance situations (95% CI: [0.084, 0.158]), but not in approach situations (95% CI: [−0.027, 0.144]). The rate gradient was more influential in playoff games compared to regular season games (95% CI on overall difference in w_3: [0.10, 0.12]; in approach [0.11, 0.14], and in avoidance [0.082, 0.118]). There was no credible difference in the influence of the time gradient between playoff and regular season games (95% CI on overall difference in w_2: [−0.066, 0.033]; in approach [−0.158, 0.021], and in avoidance [−0.022, 0.060]).

The low overall sensitivity in preseason games means it is difficult to make inferences about the relative influence of the gradients in this context, as the gradients collectively had little influence on resource expenditure in these games. So the following results must be interpretted with caution. According to our analysis, the distance gradient was more influential in regular season games (95% CI on the overall difference in w_1: [0.079, 0.224]; in approach [0.091, 0.251], and avoidance [0.021, 0.248]). There was no evidence for an overall difference in the time gradient (95% CI on the difference in w_2: [−0.027, 0.122]) nor when considering only avoidance situations (95% CI on difference in w_2: [−0.177, 0.039]). However, in approach situations, the time gradient was more influential in the preseason (95% CI on the difference in w_2: [0.089, 0.234]). Overall, the rate gradient was more influential in preseason games (95% CI on the difference in w_3: [0.069, 0.156]). This difference was credible in avoidance situations (95% CI on the difference in w_3: [0.17, 0.27]) but not in approach situations (95% CI on the difference in w_3: [−0.065, 0.074]).

We assessed whether goal importance influenced the extent to which coaches allocated greater or fewer resources in general by examining how the κ parameter varied across the three different types of games. It had a higher value in regular season games than in preseason games (95% CI: [0.42, 0.55]; in approach [0.39, 0.51], and in avoidance [0.40, 0.63]) and a higher value in playoff games than regular season games (95% CI: [0.31, 0.43]; in approach [0.38, 0.53], and in avoidance [0.19, 0.37]). These results suggested that goal importance increased the overall resources allocated to goal achievement independently of the distance to the goal, time to the deadline, and the rate of progress required to reach the goal.

Discussion

Study 2 provided an opportunity to assess the robustness and generalizability of the conclusions obtained in Study 1 by applying the GOAL architecture to a naturalistic data set from the field setting of professional basketball. As with Study 1, we found that all three gradients appeared to be influential in the approach context, and that the distance and rate gradients were influential in avoidance. In Study 2, the time gradient also appeared to be influential during avoidance, which was less the case in Study 1.

Overall, the results of Study 2 were fairly consistent with Study 1. Study 2 supported the prediction of the goal gradient hypothesis, revealing a negative relationship between distance to goal and motivation, and an overall tendency for this effect to be stronger in avoidance than in approach. It also corroborated the Study 1 finding that the increase in motivation as the goal becomes closer appears to be more strongly influenced by proximity effects than by expectancy effects. Finally, Study 2 reinforced the Study 1 finding that the achievability perspective is necessary to account for the effects of the rate of progress gradient.

Unlike Study 1, Study 2 revealed only a very small percentage of rate gradients that monotonically increasing or decreased as a function of the rate of progress required to reach the goal. As such, this study does not provide much evidence for necessity of the difficulty or expectancy perspectives. The increased prevalence of the nonmonotonic rate gradients in Study 2 may be due to coaches having a high level of domain expertise and therefore being more attuned to their chances of winning the game. The only other major point of difference in the two sets of results concerned the relative effects of the time and rate gradients. In contrast to Study 1, the time gradient was more influential in the avoidance context compared to the approach context in Study 2. Moreover, the rate gradient was found to be more influential than the time gradient in the approach context, whereas the time gradient was found to be more influential than the rate gradient in the avoidance context. This was the opposite pattern of results than was obtained in Study 1.

The differences between approach and avoidance were also generally weaker in Study 2 than in Study 1. This may be because the approach–avoidance distinction is more fluid in an NBA game than it is the experimental paradigm used in Study 1. Teams may oscillate between approach and avoidance many times within a single game. When this happens, the distinction between approach and avoidance may become less salient. There are also other factors at play that may have obscured the impact of the approach–avoidance distinction in Study 2. One such factor is the number of fouls each player has committed. NBA players are disqualified from the game after committing six fouls, so coaches will often

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

delay returning a player to the court when they are close to reaching the six-foul limit, potentially counteracting motivation to have their best players on the court. It is also the case that the coaches' motivation to rest their top players is a governed by a dynamic process and likely fluctuates depending on how many games the team has played recently and how much the top players have played in the current game. These fatigue dynamics will also affect the coach's decisions about which players to put into the game.

Study 2 also provided an opportunity to assess the impact of goal importance on goal-directed motivation. We found that the overall sensitivity to the three gradients was higher when the importance of success was greater, which is in keeping with prior work that assumes a goal's importance amplifies the signals that indicate the attractiveness of prioritizing it (Ballard et al., 2018; Vancouver et al., 2010). We also found that importance increased the resources allocated to the goal independently of the three gradients, which is consistent with work showing that importance can directly impact motivation (Locke & Latham, 1990). Goal importance also produced some systematic differences in the relative influence of the three gradients. For instance, the rate gradient was more influential when goal achievement was highly important (i.e., in playoffs games) compared to when it was only moderately important (i.e., in regular season games). By contrast, the distance gradient was more influential when goal achievement was moderately important compared to when it was highly important, though this was only the case in the avoidance context.

General Discussion

When you start work at the beginning of the day, and you have a long list of goals that you need to achieve, which one do you work on first? Do you choose the one that is closest to completion, or furthest from completion? Do you choose the one that is due first? Do you choose the one that is the easiest to finish, hardest to finish, or one that is neither too easy nor too hard? How does your motivation change once you start working on your chosen goal? There are many theories that purport to answer these questions, but they provide different answers. The GOAL architecture integrates these diverse perspectives and provides a framework for formalizing, testing, and comparing them in terms of a quantitative model with distance, time, and rate of progress gradients.

Our findings demonstrate the need for theoretical integration when attempting to understand the dynamics of motivation during goal pursuit. In particular, we found that the distance, time, and rate of progress gradients were all necessary to account for the general patterns of decision-making observed in these studies. Yet there was some evidence of individual differences in how these gradients influenced motivation. For example, although the achievability perspective was arguably the most well-supported of the three relevant to the rate gradient, many participants in Study 1 had rate gradients that were most consistent with the difficulty perspective, while some had rate gradients that were most consistent with the expectancy perspective (answering Research Question 4, see Table 1). These individual differences may reflect variation in the tendency for effort to be perceived as costly or valued (Inzlicht et al., 2018). Some individuals may see a difficult goal as an exciting and motivating challenge. Others may see it as aversive. Our findings also demonstrate that a complete account cannot be achieved without considering the type of goal being pursued, as the influence of the three gradients differed depending on whether people are pursuing approach or avoidance goals (particularly in Study 1). Taken together, these findings suggest that past perspectives on the motivational dynamics of goal pursuit are inadequate, and that a broader and more inclusive account, such as that provided by the GOAL framework, is required to provide a unified account of the dynamics of motivation.

In addition to integrating different theoretical perspectives, the GOAL architecture allowed us to resolve number areas of conflict between those theoretical perspectives. In both studies, we found evidence that motivation tends to increase as the distance to the goal decreases. This result is consistent with the proximity perspective but inconsistent with the discrepancy perspective, answering Research Question 1. Both studies also showed that the increase in motivation as distance decreases is more pronounced for avoidance goals. This result is consistent with the goal gradient hypothesis, but inconsistent with the implications of prospect theory models of goal pursuit, answering Research Question 2. The GOAL architecture also allowed us to disentangle the influence of proximity (i.e., the goal becoming *closer*) and expectancy (the goal becoming *easier*) on the relationship between distance to goal and motivation. Although both proximity and expectancy contributed to this relationship, the findings from both studies suggest that the influence of proximity was stronger than the influence of expectancy, answering Research Question 5.

We found conflicting results regarding the relative influence of the time gradient in approach and avoidance goal pursuit addressed by Research Question 3. In Study 1, the tendency for motivation to increase as the deadline drew nearer was stronger when pursuing approach goals than avoidance goals, which is consistent with the sign effect often observed in studies of intertemporal choice (Gonçalves & Silva, 2015; Murphy et al., 2001; Read et al., 2017; Thaler, 1981). However, we found the opposite pattern in Study 2, although the reversal was fairly small. We also attempted to disentangle the influence of the looming deadline itself (i.e., the deadline becoming *closer*) and the increase in the rate of progress required to reach the goal (i.e., the deadline becoming *harder*) on the increase in motivation as the time remaining decreases. These results, which address Research Question 6, were mixed. In Study 1, the effect of the deadline itself was stronger than the rate of progress when pursuing approach goals, but the opposite was true when pursuing avoidance goals. In Study 2, these relationships were weaker, and reversed (i.e., rate of progress was more important for approach goals and the deadline itself for avoidance goals).

Toward an Explanatory Model

Although the GOAL architecture can further our understanding of motivation and decision-making, the model is descriptive. It makes no assumptions about the cognitive processes that give rise to the form of its three gradients. So what type of theoretical framework is needed to explain the underlying mechanisms? We believe a key component that should underpin an explanatory framework is the assumption that motivation is derived from the perception that one's actions change the prospect of goal attainment. That is, a person is motivated to prioritize a goal to the extent that doing so increases their chances of attaining the goal and/or failing to do so decreases their chances. The person's current chances of goal achievement act as a reference point that changes over time depending on the state of

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

progress, and that is used to evaluate potential future courses of action (e.g., Jarecki & Rieskamp, 2020; Marbacher et al., 2021). This mechanism can explain why people prefer goals that are challenging but not impossible to achieve (i.e., the inverted-*u*-shaped rate gradient assumed by the achievability perspective). Goals that are too easy are poor motivators because the chance of goal attainment is already so high that further efforts do little to raise these chances. Goals that are too difficult fail to inspire because the chances of success are so low that effort expenditure will not make the goal any more achievable.

The assumption that motivation depends on expected changes in goal attainability also helps to explain the proximity and deadline effects captured by the distance and time gradients. When the goal is further from being achieved, the prospects of goal achievement are more uncertain, which weakens the perceived impact of one's actions on those prospects, and in turn, motivation. As the goal becomes closer to being achieved, the prospects of goal achievement become more certain. This makes the impact of one's actions on those prospects more salient, which increases motivation (the proximity effect). Likewise, the impact of one's actions on the prospects of goal achievement are most salient when there is less time available, as inaction during this time is more certain to result in goal failure than inaction when one has more time available. This makes motivation increase as the deadline looms (the deadline effect). The effects of distance and time are likely to be exacerbated if we adopt the widely held assumption that people seek certainty, and are more sensitive to the difference between a certain outcome and one that is merely probable (e.g., a 100% chance of goal achievement and a 95% chance) than the difference between two uncertain outcomes with probabilities that differ by the same amount (e.g., a 95% chance vs. a 90% chance). This assumption means that people are particularly motivated when given the opportunity to "lock in" a goal by ensuring its completion, or when threatened with the prospect of certain failure.

Theories of multiattribute decision-making that account for the time to make decisions (Bhatia, 2013; Noguchi & Stewart, 2018; Roe et al., 2001; Trueblood et al., 2014; Usher & Mcclelland, 2004) may offer insights into how these gradient arise from the way people process goal-related information. These theories of intradecision dynamics generally assume that preference for a particular course of action accumulates over time as people shift attention between the different attributes of the choices. It is possible that people process the distance to the goal and the time to the deadline attributes in a similar manner. Future research that seeks to test these assumptions might also benefit from using process-tracing methods, which can help to further clarify how information is being processed (Schulte-Mecklenbeck et al., 2017).

A model that captures how people process goal-related information will also help us understand the effects of time pressure and capacity limitations. The principle of bounded rationality assumes that people have limited information processing resources and so must often rely on relatively simple decision rules or heuristics, particularly when time or capacity is limited (Gigerenzer & Goldstein, 1996; Simon, 1956). In the context of goal pursuit, one strategy might be to simply prioritize the goal with the shortest deadline. An alternative would be to prioritize the goal that is closest to being reached. We might also expect people to rely on simpler rules or heuristics when the number of goals being pursued is greater, or when pursuing different combinations of approach versus avoidance goals because the presence of additional goals or combinations of different types of goals would increase the complexity of the computations required to integrate and compare information regarding distance and time across the goals.

What role is there for individual differences in this integrative perspective? We know that there is a general tendency for people to be more sensitive to undesirable outcomes than to desirable outcomes (Kahneman & Tversky, 1979), which may help explain why the proximity effects were stronger under avoidance than under approach in the current studies. However, there is also variation in this tendency. Those more sensitive to desirable outcomes (e.g., promotion focused individuals) are likely to be more motivated by approach goals, and the salience of goal attainment may increase more strongly as distance to the goal decreases (Förster et al., 1998; Zou et al., 2020). Among those more sensitive to undesired outcomes (e.g., prevention focused individuals) the same may be true about the salience of failing avoidance goals. People are also likely to differ in their optimism regarding the likelihood of goal achievement (Zhang et al., 2007). This variation may be due to genuine differences in skill or biases in the perception of their own ability (e.g., overconfidence). Likewise, people may differ in their tolerance of uncertainty, with some preferring to secure goal achievement as soon as possible and others being more comfortable with leaving goals unresolved for longer periods of time. It is also possible that people may have different internal representations of distance and time in the context of goal pursuit. For example, people may vary in the extent to which they exhibit diminishing sensitivity to each of these quantities (Dehaene, 2003; Scholten & Read, 2010).

A Linking Architecture

The GOAL architecture also has the potential to help bridge the gap between theories of motivation and models of decision-making, opening up new areas for future investigation. For example, models of choice and response time (e.g., Brown & Heathcote, 2008; Ratcliff & McKoon, 2008) consider the dynamics that play out within a single decision episode. These models assume that people make decisions by accumulating evidence in favor of possible courses of action until the evidence for one alternative breaches a threshold. However, most research applying these models has only examined independent, one-shot decisions, so fails to capture the motivational dynamics that play out across a series of interdependent decisions, such as those involved in goal pursuit. The GOAL architecture provides a front-end to models of choice and response time that could help investigate these dynamics and answer questions about threshold and the rate of evidence accumulation change across decisions as people work toward a goal.

Another area in which the GOAL architecture could contribute is in understanding the potential for effort to be viewed as both costly and valuable (Inzlicht et al., 2018). Work examining the cost of effort has shown that people tend to choose options that require less work (Kool et al., 2010; Westbrook & Braver, 2015; Westbrook et al., 2013), a principal that has come to be known as the "law of less work." Yet studies in this literature typically do not consider the effects that goals and deadlines have on the willingness to engage in a cognitively demanding task. In the GOAL architecture, the willingness to exert effort may change as a function of the amount of time one has available and the length of time that effort must

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

sustained in order to achieve the goal. These insights can be used to generate new predictions regarding how people regulate cognitive effort.[15]

On the other side of the effort paradox is research that has demonstrated the value of effort. The GOAL architecture provides a way of linking a host of empirical effects regarding the satisfaction that can be derived from effort within a common framework. For example, the martyrdom effect (Olivola & Shafir, 2013), Ikea effect (Mochon et al., 2012; Norton et al., 2012), the preference for earned income versus windfall gains (Arkes et al., 1994; Kameda et al., 2002), and the notion of contra-freeloading (Inglis et al., 1997) all reflect a tendency for outcomes to be valued more when they are achieved through an effortful process. These effects are encapsulated by the difficulty perspective. By contrast, the tendency for individuals to self-handicap by disengaging from a task so that they can attribute poor performance to lack of effort (Hendrix & Hirt, 2009; Urdan & Midgley, 2001) is captured by the expectancy perspective. Individual differences in the rate of progress gradient may be partially accounted for by need for cognition, which reflects the extent to which a person engages in and enjoys effortful activity (Cacioppo & Petty, 1982; Sandra & Otto, 2018).[16]

The GOAL architecture can also be used to assess the effectiveness of people's decisions, for example, by comparing the decisions of human participants to those of a normative model. In many experimental paradigms, including the one used for Study 1, methods such as dynamic programming or reinforcement learning can be used to identify the decisions that maximize the rate of goal achievement. Fitting the GOAL architecture to these decisions provides insight into the nature of the distance, time, and rate gradients that would be expected if people's decisions align with a normative model. By comparing these so-called normative gradients with the gradients observed from human participants, we can identify whether people over- or under-weight certain gradients in such a way that compromises their prospects of goal achievement. The Supplemental Materials contains an example of such an analysis using the data from Study 1. This analysis shows that many participants systematically over-weighted the distance gradient and under-weighted the rate of progress gradient.

Limitations and Extensions

It should be noted that the scope of the GOAL architecture does not extend to all types of goals. It is applicable to goals that have concrete deadline, and for which progress toward the goal can be quantified, but not more abstract goals like "being a good person." However, we believe this framework is relevant to many of the goals that people encounter, particularly so-called "SMART" goals (Specific, Measurable, Achievable, Realistic, and Time-bound), which goal-setting researchers have long regarded as highly effective for improving performance and encouraging behavior change (e.g., Doran, 1981; Fielding, 1999; Locke & Latham, 2002). Nevertheless, there are possible extensions of this framework that apply to goals that do not fall into the SMART category. For example, for goals where distance and time have more qualitative representations ("I'm close to achieving this goal" vs. "I'm far from achieving this goal" or "I need to be complete this soon" vs. "I have plenty of time"), distance and time could be represented as random variables with distributions that account for the fact that the individual may not have a clear sense of their current status with respect to these variables.

We examined scenarios where no more than two goals needed to be represented in the model, and when two goals were being simultaneously pursued they were always the same type. However, the GOAL architecture can be applied to problems involving any number of approach and avoidance goals. This could be achieved using a multinomial logistic decision rule, or using a racing accumulator model (e.g., Ravenzwaaij et al., 2019) when modeling intradecision dynamics. Such extensions can be used to examine, for example, how people respond to increases in complexity in terms of the number or types of goals being pursued simultaneously.

Although we believe that the consistency of the results between our experimental and field contexts provides evidence for the generality of our inferences, both studies had goals that were not set by the participants. We opted for this setup because individual differences in the goals that people set can obscure the effects of the variables we have attempted to isolate. However, by building on the understanding of how distance to goal and time to deadline combine to influence the dynamics of motivation developed here, it becomes possible to begin to explore these dynamics in more complex settings, such as when individuals (or teams) set their own goals. Indeed, future work may benefit from integrating the GOAL architecture with formal models of goal setting (e.g., Gee et al., 2018) in order to develop a more general understanding of the goal regulation process.

Conclusion

The aim of this research was to provide a formal, unifying framework that can be used to integrate and compare theories of

[15] The GOAL architecture also provides a framework for linking behavioral measures of effort and motivation to physiological data that captures underlying neural mechanisms (e.g., Palestro et al., 2018; Palmeri et al., 2017; Turner et al., 2016). EEG or fMRI measurements could be correlated with the strength of each gradient in order to elucidate their neural underpinnings, or the model fit jointly to behavioral and neural data (see Forstmann & Wagenmakers, 2015).

[16] Financial decision-making is a related area in which the GOAL architecture also has the potential to enhance our understanding. For example, Gathergood et al. (2019) describe several heuristics that consumers might use to determine how to allocate repayments when managing multiple credit card debts, all of which are all linked to the distance gradient. The tendency to prioritize the card with the highest balance, or to allocate repayments in proportion to the balance on each card (balance-matching), is consistent with a positive distance gradient in a context where the person has the approach goal of reducing each balance to zero. The tendency to prioritize the card with the lowest balance (the debt snowball method) is consistent with a negative distance gradient. The tendency to allocate repayment based on card capacity can be accounted for by treating it as an avoidance goal (i.e., avoid a capacity of zero). In this case, prioritizing the card with the higher capacity is consistent with a positive distance gradient, whereas prioritizing a card with the lower capacity is consistent with a negative distance gradient. The GOAL architecture can also capture the unique effects that assets and debts have on financial decision-making in different contexts (Sussman & Shafir, 2012), for example, by treating asset growth and debt reduction as distinct goals and assuming that the weights associated with the distance gradients for these goals change across settings. However, our findings suggest that these effects may also be influenced by time. For example, if a person has savings or debt reduction goals with concrete deadlines attached, they may deviate from these tendencies and instead use strategies such as prioritizing the goal with the shorter deadline, or the one that requires larger repayments in order to complete within the allotted time. The GOAL architecture has, therefore, potential as a more general framework to think about the factors that are likely to influence these types of decisions.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

goal-directed motivation and decision making. The GOAL architecture provides this framework, extending our understanding of the goal pursuit process by combining previously disparate accounts into a more general architecture that describes the distance and temporal dynamics of goal pursuit. Using this architecture, we demonstrated that motivation during goal pursuit changes in a complex manner that cannot be easily accounted for by any one existing theoretical perspective, but that could be captured by the more general framework. Our findings highlight the importance of quantitative theoretical integration when attempting to understand the factors driving motivation and decision-making in the context of goal pursuit.

References

Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, *82*(4), 463–496. https://doi.org/10.1037/h0076860

Anderson, A. C. (1933). Runway time and the goal gradient. *Journal of Experimental Psychology*, *16*(3), 423–428. https://doi.org/10.1037/h0074765

Arkes, H. R., Joyner, C. A., Pezzo, M. V., Nash, J. G., Siegel-Jacobs, K., & Stone, E. (1994). The psychology of windfall gains. *Organizational Behavior and Human Decision Processes*, *59*(3), 331–347. https://doi.org/10.1006/obhd.1994.1063

Atkinson, J. W., & Birch, D. (1970). *The dynamics of action*. Wiley.

Austin, J. T., & Vancouver, J. B. (1996). Goal constructs in psychology: Structure, process, and content. *Psychological Bulletin*, *120*(3), 338–375. https://doi.org/10.1037/0033-2909.120.3.338

Ballard, T., Farrell, S., & Neal, A. (2018). Quantifying the psychological value of goal achievement. *Psychonomic Bulletin and Review*, *25*(3), 1184–1192. https://doi.org/10.3758/s13423-017-1329-1

Ballard, T., Neal, A., Farrell, S., Lloyd, E., Lim, J., & Heathcote, A. (2021, July). *A general architecture for modeling the dynamics of goal-directed motivation and decision making*. OSF. https://doi.org/10.17605/OSF.IO/CJYND

Ballard, T., Yeo, G., Loft, S., Vancouver, J. B., & Neal, A. (2016). An integrative formal model of motivation and decision making: The MGPM. *Journal of Applied Psychology*, *101*(9), 1240–1265. https://doi.org/10.1037/apl0000121

Ballard, T., Yeo, G., Neal, A., & Farrell, S. (2016). Departures from optimality when pursuing multiple approach or avoidance goals. *Journal of Applied Psychology*, *101*(7), 1056–1066. https://doi.org/10.1037/apl0000082

Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, *84*(2), 191–215. https://doi.org/10.1037/0033-295X.84.2.191

Beck, J. W., Schmidt, A. M., & Natali, M. W. (2019). Efficient proximal resource allocation strategies predict distal team performance: Evidence from the National Hockey League. *Journal of Applied Psychology*, *104*(11), 1387–1403. https://doi.org/10.1037/apl0000407

Berg, I. A., & Weisman, R. L. (1942). The goal gradient in a maze of variable length. *The Journal of Psychology*, *14*(2), 307–315. https://doi.org/10.1080/00223980.1942.9917129

Besler, A. (2020). *nbastatR*. http://asbcllc.com/nbastatR/

Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review*, *120*(3), 522–543. https://doi.org/10.1037/a0032457

Biner, P. M. (1987). Effects of difficulty and goal value on goal valence. *Journal of Research in Personality*, *21*(3), 395–404. https://doi.org/10.1016/0092-6566(87)90020-1

Bonezzi, A., Brendl, C. M., & De Angelis, M. (2011). Stuck in the middle: The psychophysics of goal pursuit. *Psychological science*, *22*(5), 607–612. https://doi.org/10.1177/0956797611404899

Brehm, J. W., & Self, E. A. (1989). The intensity of motivation. *Annual Review of Psychology*, *40*, 109–131. https://doi.org/10.1146/annurev.ps.40.020189.000545

Brehm, J. W., Wright, R. A., Solomon, S., Greenberg, J., & Wright, R. (1983). Perceived difficulty, energization, and the magnitude of goal valence. *Journal of Experimental Social Psychology*, *19*(1), 21–48. https://doi.org/10.1016/0022-1031(83)90003-3

Brown, J. S. (1948). Gradients of approach and avoidance responses and their relation to level of motivation. *Journal of Comparative and Physiological Psychology*, *41*(6), 450–465. https://doi.org/10.1037/h0055463

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178. https://doi.org/10.1016/j.cogpsych.2007.12.002

Bryan, J. F., & Locke, E. A. (1967). Parkinson's law as a goal-setting phenomenon. *Organizational Behavior and Human Performance*, *2*(3), 258–275. https://doi.org/10.1016/0030-5073(67)90021-9

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology*, *42*(1), 116–131. https://doi.org/10.1037/0022-3514.42.1.116

Campion, M. A., & Lord, R. G. (1982). A control systems conceptualization of the goal-setting and changing process. *Organizational Behavior and Human Performance*, *30*(2), 265–287. https://doi.org/10.1016/0030-5073(82)90221-5

Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal Of Statistical Software*, *76*(1), 1–32. https://doi.org/10.18637/jss.v076.i01

Carver, C. S., & Scheier, M. F. (1990). Origins and functions of positive and negative affect: A control-process view. *Psychological Review*, *97*(1), 19–35. https://doi.org/10.1037/0033-295X.97.1.19

Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. Cambridge University Press.

Cryder, C. E., Loewenstein, G., & Seltman, H. (2013). Goal gradient in helping behavior. *Journal of Experimental Social Psychology*, *49*(6), 1078–1083. https://doi.org/10.1016/j.jesp.2013.07.003

Csikszentmihalyi, M. (2014). *Flow and the foundations of positive psychology*. Springer.

Dehaene, S. (2003). The neural basis of the weber–fechner law: A logarithmic mental number line. *Trends in Cognitive Sciences*, *7*(4), 145–147. https://doi.org/10.1016/s1364-6613(03)00055-x

DeShon, R. P., & Gillespie, J. Z. (2005). A motivated action theory account of goal orientation. *Journal of Applied Psychology*, *90*(6), 1096–1127. https://doi.org/10.1037/0021-9010.90.6.1096

Doran, G. (1981). There's a S.M.A.R.T way to write management's goals and objectives. *Management Review*, *70*(11), 35–36.

Doyle, J. R. (2013). Survey of time preference, delay discounting models. *Judgment and Decision Making*, *8*(2), 116–135. https://doi.org/10.2139/ssrn.1685861

Estle, S. J., Green, L., Myerson, J., & Holt, D. D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. *Memory & Cognition*, *34*(4), 914–928. https://doi.org/10.3758/BF03193437

Fielding, M. (1999). Target setting, policy pathology and student perspectives: Learning to labour in new times. *Cambridge Journal of Education*, *29*(2), 277–287. https://doi.org/10.1080/0305764990290210

Förster, J., Higgins, E. T., & Idson, L. C. (1998). Approach and avoidance strength during goal attainment: Regulatory focus and the "goal looms larger" effect. *Journal of Personality and Social Psychology*, *75*(5), 1115–1131. https://doi.org/10.1037/0022-3514.75.5.1115

B. U. Forstmann, & E. J. Wagenmakers (Eds.). (2015). Model-based cognitive neuroscience: A conceptual introduction. In *An Introduction to Model-Based Cognitive Neuroscience* (pp. 139–156). Springer. https://doi.org/10.1007/978-1-4939-2236-9_7

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Gathergood, J., Mahoney, N., Stewart, N., & Weber, J. (2019). How do individuals repay their debt? The balance-matching heuristic. *American Economic Review*, *109*(3), 844–875. https://doi.org/10.1257/aer.20180288

Gee, P., Neal, A., & Vancouver, B. (2018). A formal model of goal revision in approach and avoidance contexts. *Organizational Behavior and Human Decision Processes*, *146*, 51–61. https://doi.org/10.1016/j.obhdp.2018.03.002

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*(4), 650–669. https://doi.org/10.1037/0033-295x.103.4.650

Gonçalves, F. L., & Silva, M. T. A. (2015). Comparing individual delay discounting of gains and losses. *Psychology & Neuroscience*, *8*(2), 267–279. https://doi.org/10.1037/h0101057

Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, *130*(5), 769–792. https://doi.org/10.1037/0033-2909.130.5.769

Heath, C., Larrick, R. P., & Wu, G. (1999). Goals as reference points. *Cognitive psychology*, *38*(1), 79–109. https://doi.org/10.1006/cogp.1998.0708

Heathcote, A., Brown, S. D., & Wagenmakers, E. J. (2015). An Introduction to good practices in cognitive modeling. In B. U. Forstmann & E. J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience* (pp. 25–48). Springer.

Hendrix, K. S., & Hirt, E. R. (2009, January). Stressed out over possible failure: The role of regulatory fit on claimed self-handicapping. *Journal of Experimental Social Psychology*, *45*(1), 51–59. https://doi.org/10.1016/j.jesp.2008.08.016

Hendy, K. C., Liao, J., & Milgram, P. (1997). Combining time and intensity effects in assessing operator information-processing load. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *39*(1), 30–47. https://doi.org/10.1518/001872097778940597

Higgins, E. T. (1997). Beyond pleasure and pain. *American Psychologist*, *52*(12), 1280–1300. https://doi.org/10.1037/0003-066X.52.12.1280

Hockey, G. R. J. (1997). Compensatory control in the regulation of human performance under stress and high workload: A cognitive-energetical framework. *Biological Psychology*, *45*(1–3), 73–93. https://doi.org/10.1016/S0301-0511(96)05223-4

Hull, C. L. (1932). The goal-gradient hypothesis and maze learning. *Psychological Review*, *39*(1), 25–43. https://doi.org/10.1037/h0072640

Hull, C. L. (1938). The goal-gradient hypothesis applied to some "field-force' problems in the behavior of young children. *Psychological Review*, *45*(4), 271–299. https://doi.org/10.1037/h0053885

Humphreys, M. S., & Revelle, W. (1984). Personality, motivation, and performance: A theory of the relationship between individual differences and information processing. *Psychological Review*, *91*(2), 153–184. https://doi.org/10.1037/0033-295X.91.2.153

Inglis, I. R., Forkman, B., & Lazarus, J. (1997). Free food or earned food? A review and fuzzy model of contrafreeloading. *Animal Behaviour*, *53*(6), 1171–1191. https://doi.org/10.1006/anbe.1996.0320

Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The effort paradox: Effort is both costly and valued. *Trends in cognitive sciences*, *22*(4), 337–349. https://doi.org/10.1016/j.tics.2018.01.007

Jagacinski, R. J., & Flach, J. M. (2003). *Control theory for humans: Quantitative approaches to modeling performance*. Lawrence Erlbaum.

Jarecki, J. B., & Rieskamp, J. (2020). *Prospect theory and optimal risky choices with goals*. Proceedings of the 42nd Annual Conference of the Cognitive Science Society (virtual meeting).

Kahneman, D., & Tversky, A. (1979). Propsect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. https://doi.org/10.2307/1914185

Kameda, T., Takezawa, M., Tindale, R. S., & Smith, C. M. (2002). Social sharing and risk reduction: Exploring a computational algorithm for the psychology of windfall gains. *Evolution and Human Behavior*, *23*(1), 11–33. https://doi.org/10.1016/S1090-5138(01)00086-1

Karabenick, S. A., & Youssef, Z. I. (1968). Performance as a function of achievement motive level and perceived difficulty. *Journal of Personality and Social Psychology*, *10*(4), 414–419. https://doi.org/10.1037/h0026735

Kernan, M. C., & Lord, R. G. (1990). Effects of valence, expectancies, and goal-performance discrepancies in single and multiple goal environments. *Journal of Applied Psychology*, *75*(2), 194–203. https://doi.org/10.1037/0021-9010.75.2.194

Kivetz, R., Urminsky, O., & Zheng, Y. (2006). The goal-gradient hypothesis resurrected: purchase acceleration, illusionary goal progress, and customer Retention. *Journal of Marketing Research*, *43*(1), 39–58. https://doi.org/10.1509/jmkr.43.1.39

Klein, H. J. (1989). An integrated control theory model of work motivation. *Academy of Management Review*, *14*(2), 150–172. https://doi.org/10.2307/258414

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General*, *139*(4), 665–682. https://doi.org/10.1037/a0020198

Koop, G. J., & Johnson, J. G. (2012). The use of multiple reference points in risky decision making. *Journal of Behavioral Decision Making*, *25*(1), 49–62. https://doi.org/10.1002/bdm.713

Kruschke, J. K., Aguinis, H., & Joo, H. (2012). The time has come: Bayesian methods for data analysis in the organizational sciences. *Organizational Research Methods*, *15*(4), 722–752. https://doi.org/10.1177/1094428112457829

Kukla, A. (1972). Foundations of an attributional theory of performance. *Psychological Review*, *79*(6), 454–470. https://doi.org/10.1037/h0033494

Latham, G. P., & Locke, E. A. (1975). Increasing productivity and decreasing time limits: A field replication of Parkinson's law. *Journal of Applied Psychology*, *60*(4), 524–526. https://doi.org/10.1037/h0076916

Lempert, K. M., & Phelps, E. A. (2016). The malleability of intertemporal choice. *Trends in Cognitive Sciences*, *20*(1), 64–74. https://doi.org/10.1016/j.tics.2015.09.005

Lewin, K. (1935). *A dynamic theory of personality*. McGraw-Hill.

Liberman, N., & Förster, J. (2008). Expectancy, value and psychological distance: A new look at goal gradients. *Social Cognition*, *26*(5), 515–533. https://doi.org/10.1521/soco.2008.26.5.515

Locke, E. A., & Latham, G. P. (1990). *A theory of goal setting and task performance*. Prentice-Hall.

Locke, E. A., & Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American Psychologist*, *57*(9), 705–717. https://doi.org/10.1037/0003-066X.57.9.705

Loewenstein, G., & Issacharoff, S. (1994). Source dependence in the valuation of objects. *Journal of Behavioral Decision Making*, *7*(3), 157–168. https://doi.org/10.1002/bdm.3960070302

Loewenstein, G., Rick, S., & Cohen, J. D. (2008). Neuroeconomics. *Annual Review of Psychology*, *59*, 647–672. https://doi.org/10.1146/annurev.psych.59.103006.093710

Lord, R. G., & Levy, P. E. (1994). Moving from cognition to action: A control theory perspective. *Applied Psychology: An International Review*, *43*(3), 335–398. https://doi.org/10.1111/j.1464-0597.1994.tb00828.x

Louro, M. J., Pieters, R., & Zeelenberg, M. (2007). Dynamics of multiple-goal pursuit. *Journal of Personality and Social Psychology*, *93*(2), 174–193. https://doi.org/10.1037/0022-3514.93.2.174

Luce, R. D., & Raiffa, H. (1957). *Games and decisions*. Wiley.

Marbacher, L., Jarecki, J. B., & Rieskamp, J. (2021). *How goals erase framing effects in risky decision making*. Proceedings of the 43rd Annual Conference of the Cognitive Science Society (virtual meeting).

Mochon, D., Norton, M. I., & Ariely, D. (2012). Bolstering and restoring feelings of competence via the IKEA effect. *International Journal of Research in Marketing*, *29*(4), 363–369. https://doi.org/10.1016/j.ijresmar.2012.05.001

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Murphy, J. G., Vuchinich, R. E., & Simpson, C. A. (2001). Delayed reward and cost discounting. *The Psychological Record*, *51*(4), 571–588.

Neal, A., Ballard, T., & Vancouver, J. B. (2017). Dynamic self-regulation and multiple-goal pursuit. *Annual Review of Organizational Psychology and Organizational Behavior*, *4*, 410–423. https://doi.org/10.1146/annurev-orgpsych-032516-113156

Noguchi, T., & Stewart, N. (2018). Multialternative decision by sampling: A model of decision making constrained by process data. *Psychological Review*, *125*(4), 512–544. https://doi.org/10.1037/rev0000102

Norton, M. I., Mochon, D., & Ariely, D. (2012). The IKEA effect: When labor leads to love. *Journal of Consumer Psychology*, *22*(3), 453–460. https://doi.org/10.1016/j.jcps.2011.08.002

Olivola, C. Y. (2011). When noble means hinder noble ends: the benefits and costs of a preference for martyrdom in altruism. In D. M. Oppenheimer & C. Y. Olivola (Eds.), *The science of giving: Experimental approaches to the study of charity* (pp. 49–62). Taylor & Francis.

Olivola, C. Y., & Shafir, E. (2013). The Martyrdom Effect: When pain and effort increase prosocial contributions. *Journal of Behavioral Decision Making*, *26*(1), 91–105. https://doi.org/10.1002/bdm.767

Palada, H., Neal, A., Tay, R., & Heathcote, A. (2018). Understanding the causes of adapting, and failing to adapt, to time pressure in a complex multistimulus environment. *Journal of Experimental Psychology: Applied*, *24*(3), 380–399. https://doi.org/10.1037/xap0000176

Palestro, J. J., Bahg, G., Sederberg, B., Lu, Z. -L., Steyvers, M., & Turner, B. M. (2018). A tutorial on joint models of neural and behavioral measures of cognition. *Journal of Mathematical Psychology*, *84*, 20–48. https://doi.org/10.1016/j.jmp.2018.03.003

Palmeri, T. J., Love, B. C., & Turner, B. M. (2017). Model-based cognitive neuroscience. *Journal of Mathematical Psychology*, *76*(Part B), 59–64. https://doi.org/10.1016/j.jmp.2016.10.010

Peters, L. H., O'Connor, E. J., Pooyan, A., & Quick, J. C. (1984). The relationship between time pressure and performance: A field test of Parkinson's Law. *Journal of Occupational Behaviour*, *5*(4), 293–299. https://doi.org/10.1002/job.4030050406

Powers, W. T. (1973). *Behavior: The control of perception*. Aldine.

Powers, W. T. (1978). Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review*, *85*(5), 417–435. https://doi.org/10.1037/0033-295X.85.5.417

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Ravenzwaaij, D. V., Brown, S. D., Marley, A. A. J., & Heathcote, A. (2019). Accumulating advantages: A new conceptualization of rapid multiple choice. *Psychological Review*, *127*(2), 186–215. https://doi.org/10.1037/rev0000166

Read, D. (2004). Intertemporal choice. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 424–443). Blackwell Publishing. https://doi.org/10.1002/9780470752937.ch21

Read, D., Olivola, C. Y., & Hardisty, D. J. (2017). The value of nothing: Asymmetric attention to opportunity costs drives intertemporal decision making. *Management Science*, *63*(12), 4277–4297. https://doi.org/10.1287/mnsc.2016.2547

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, *108*(2), 370–392. https://doi.org/10.1037/0033-295X.108.2.370

Rung, J. M., & Madden, G. J. (2018). Experimental reductions of delay discounting and impulsive choice: A systematic review and meta-analysis. *Journal of Experimental Psychology: General*, *147*(9), 1349–13831. https://doi.org/10.1037/xge0000462

Sandra, D. A., & Otto, A. R. (2018). Cognitive capacity limitations and need for cognition differentially predict reward-induced cognitive effort expenditure. *Cognition*, *172*, 101–106. https://doi.org/10.1016/j.cognition.2017.12.004

Schmidt, A. M., & DeShon, R. P. (2007). What to do? The effects of discrepancies, incentives, and time on dynamic goal prioritization. *Journal of Applied Psychology*, *92*(4), 928–941. https://doi.org/10.1037/0021-9010.92.4.928

Schmidt, A. M., & Dolis, C. M. (2009). Something's got to give: The effects of dual-goal difficulty, goal progress, and expectancies on resource allocation. *Journal of Applied Psychology*, *94*(3), 678–691. https://doi.org/10.1037/a0014945

Schmidt, A. M., Dolis, C. M., & Tolli, A. P. (2009). A matter of time: Individual differences, contextual dynamics, and goal progress effects on multiple-goal self-regulation. *Journal of Applied Psychology*, *94*(3), 692–709. https://doi.org/10.1037/a0015012

Scholten, M., & Read, D. (2010). The psychology of intertemporal tradeoffs. *Psychological Review*, *117*(3), 925–944. https://doi.org/10.1037/a0019619

Schulte-Mecklenbeck, M., Johnson, J. G., Böckenholt, U., Goldstein, D. G., Russo, J. E., Sullivan, N. J., & Willemsen, M. C. (2017). Process-tracing methods in decision making: On growing up in the 70s. *Current Directions in Psychological Science*, *26*(5), 442–450. https://doi.org/10.1177/0963721417708229

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138. https://doi.org/10.1037/h0042769

Smith, P. L., & Little, D. R. (2018). Small is beautiful: In defense of the small-N design. *Psychonomic Bulletin and Review*, *25*(6), 2083–2101. https://doi.org/10.3758/s13423-018-1451-8

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, *64*(Part 4), 583–639. https://doi.org/10.1111/1467-9868.00353

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van der Linde, A. (2014). The deviance information criterion: 12 years on. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, *76*(3), 485–493. https://doi.org/10.1111/rssb.12062

Steel, P., & König, C. J. (2006). Integrating theories of motivation. *Academy of Management Review*, *31*(4), 889–913. https://doi.org/10.2307/20159257

Sussman, A. B., & Shafir, E. (2012). On assets and debt in the psychology of perceived wealth. *Psychological Science*, *23*(1), 101–108. https://doi.org/10.1177/0956797611421484

Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, *8*(3), 201–207. https://doi.org/10.1016/0165-1765(81)90067-7

Trew, J. L. (2011). Exploring the roles of approach and avoidance in depression: An integrative model. *Clinical Psychology Review*, *31*(7), 1156–1168. https://doi.org/10.1016/j.cpr.2011.07.007

Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological review*, *121*(2), 179–205. https://doi.org/10.1037/a0036137

Turner, B. M., Rodriguez, C. A., Norcia, T. M., Mcclure, S. M., & Steyvers, M. (2016). Why more is better: Simultaneous modeling of EEG, fMRI, and behavioral data. *NeuroImage*, *128*, 96–115. https://doi.org/10.1016/j.neuroimage.2015.12.030

Unsworth, K., Yeo, G., & Beck, J. (2009). Multiple goals: A review and derivation of general principles. *Journal of Organizational Behavior*, *30*, 839–862. https://doi.org/10.1002/job

Urdan, T., & Midgley, C. (2001). Academic self-handicapping: What we know, what more there is to learn. *Educational Psychology Review*, *13*(2), 115–138. https://doi.org/10.1023/A:1009061303214

Usher, M., & Mcclelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*(3), 757–769. https://doi.org/10.1037/0033-295X.111.3.757

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Vancouver, J. B. (2005). The depth of history and explanation as benefit and bane for psychological control theories. *Journal of Applied Psychology*, *90*(1), 38–52. https://doi.org/10.1037/0021-9010.90.1.38

Vancouver, J. B. (2008). Integrating self-regulation theories of work motivation into a dynamic process theory. *Human Resource Management Review*, *18*(1), 1–18. https://doi.org/10.1016/j.hrmr.2008.02.001

Vancouver, J. B., Weinhardt, J. M., & Schmidt, A. M. (2010). A formal, computational theory of multiple-goal pursuit: Integrating goal-choice and goal-striving processes. *Journal of Applied Psychology*, *95*(6), 985–1008. https://doi.org/10.1037/a0020628

von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton University Press.

Vroom, V. H. (1964). *Work and motivation*. Wiley.

Wang, X. T., & Johnson, J. G. (2012). A tri-reference point theory of decision making under risk. *Journal of Experimental Psychology: General*, *141*(4), 743–756. https://doi.org/10.1037/a0027415

Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, *11*(116), 3571–3594. https://doi.org/10.1002/job.1963

Weiss, R. F. (1960). Deprivation and reward magnitude effects on speed throughout the goal gradient. *Journal of Experimental Psychology*, *60*(6), 384–390. https://doi.org/10.1037/h0043425

Westbrook, A., & Braver, T. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive, Affective, & Behavioral Neuroscience*, *15*(2), 395–415. https://doi.org/10.3758/s13415-015-0334-y

Westbrook, A., Kester, D., & Braver, T. S. (2013). What is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. *PLOS ONE*, *8*(7), Article e68210. https://doi.org/10.1371/journal.pone.0068210

Wright, R. A. (2008). Refining the prediction of effort: Brehm's distinction between potential motivation and motivation intensity. *Social and Personality Psychology Compass*, *2*(2), 682–701. https://doi.org/10.1111/j.1751-9004.2008.00093.x

Wright, R. A., Contrada, R. J., & Patane, M. J. (1986). Task difficulty, cardiovascular response, and the magnitude of goal valence. *Journal of Personality and Social Psychology*, *51*(4), 837–843. https://doi.org/10.1037/0022-3514.51.4.837

Zhang, Y., Fishbach, A., & Kruglanski, A. W. (2007). The dilution model: How additional goals undermine the perceived instrumentality of a shared path. *Journal of Personality and Social Psychology*, *92*(3), 389–401. https://doi.org/10.1037/0022-3514.92.3.389

Zou, X., Scholer, A. A., & Higgins, E. T. (2020). Risk preference: How decision maker's goal, current value state, and choice set work together. *Psychological Review*, *127*(1), 74–94. https://doi.org/10.1037/rev0000162

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Appendix

Details of Model Estimation

In order to efficiently estimate the model parameters, we fit a reparameterized, but formally equivalent, version of the model in which the direction of the distance and time gradients was controlled by the w_1 and w_2 parameters, respectively, as opposed to the δ and τ parameters. In the reparameterized model, w_1 was allowed to take on any real value and w_2 was allowed to take on any negative value (constraining the time gradient to be negative). A w_3 parameter was estimated and allowed to take on any positive value, whereas θ was fixed to 1. This reparameterization makes the model much more efficient to fit because it does not require an implementation of the conditional in Equations 1 and 2.

The following transformations were used to convert the parameters of the model that was fit to the data to the parameters interpreted in the article. We use the prime symbol (′) to denote parameters from the reparameterized version of the model. The θ parameter was calculated by summing the absolute values of the weight parameters from the reparameterized model as shown in Equation A1:

$$\theta = |w_1'| + |w_2'| + w_3'. \tag{A1}$$

Note that w_3' was constrained to be positive, so $w_3' = |w_3'|$. The normalized weight parameters that we interpret in the article could then be calculating by dividing $|w_1'|$, $|w_2'|$, and w_3' by θ as shown in Equations A2–A4:

$$w_1 = \frac{|w_1'|}{\theta}, \tag{A2}$$

$$w_2 = \frac{|w_2'|}{\theta}, \tag{A3}$$

$$w_3 = \frac{w_3'}{\theta}. \tag{A4}$$

The δ parameter was then calculated according to Equation A5:

$$\delta = \begin{cases} \delta', & \text{if } w_1' >= 0 \\ -\delta', & \text{if } w_1' < 0. \end{cases} \tag{A5}$$

Since w_2' was constrained to be negative in the reparameterized model, $\tau = -\tau'$. The α parameter was not transformed in the reparameterized model, so $\alpha = \alpha'$.

The models that were fit to the observed decisions had hierarchical structures. The w_1', w_2', and w_3' parameters were sampled from normal (or truncated normal) distributions. Each of these parameters had a population mean that was drawn from a *Normal*(0, 5) distribution. The mean of the w_2' distribution was constrained to be negative whereas the mean of the w_3' distribution was constrained to be positive. The w_1', w_2', and w_3' parameters each had a population standard deviation that was drawn from *Normal*(0, 5) distribution that was truncated at zero.

The δ', τ', and α' parameters were estimated by sampling values that could fall anywhere on the real line and transforming the sampled value using a probit function. Each value was sampled from a normal distribution with a population mean drawn from a *Normal*(0, 0.75) distribution and a population standard deviation also drawn from a *Normal*(0, 0.75) distribution (truncated at zero). The sampled value was then transformed using a probit function according to Equation A6:

$$y = \Phi(x) \times (U - L) + L, \tag{A6}$$

where Φ is the standard normal cumulative distribution function, x is the sampled value, y is the value of the transformed parameter used to generate model predictions, U is the upper bound on the parameter, L is the lower bound. This parameterization produced priors that

were approximately uniform across the allowable ranges of these three parameters. The δ and τ parameters were bounded between 0.2 and 1 in Study 1 and 0.05 and 1 in Study 2. The α parameter was bounded between 0.01 and 0.99 in both studies.

The κ parameter in Study 2 was estimated by sampling its logged value from a normal distribution with a population drawn from a *Normal*(0, 0.5) distribution bounded between $-ln(10)$ and $ln(10)$ and a population standard deviation also drawn from a *Normal*(0, 0.5) distribution (truncated at zero).

Received December 9, 2019
Revision received July 22, 2021
Accepted July 25, 2021 ■

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.