

Change one can believe in: Adding learning to computational models of self-regulation



Jeffrey B. Vancouver^a, Justin M. Weinhardt^{b,*}, Ronaldo Vigo^a

^aDepartment of Psychology, Ohio University, United States

^bHaskayne School of Business, University of Calgary, Canada

ARTICLE INFO

Article history:

Received 29 January 2013

Accepted 5 December 2013

Available online 21 January 2014

Accepted by Douglas Brown

Keywords:

Computational model

Motivation

Dynamics

Learning

ABSTRACT

Theories of self-regulation describe motivation as a dynamic process of goal choice and goal striving. To facilitate those processes, individuals learn about themselves and their environment, which is an internal dynamic process. However, the precise nature of the relationship between these learning and motivational processes is not well specified. This article integrates formal models of learning, goal choice, and goal striving using a single information processing structure found in self-regulatory models of motivation. Results from two published studies (DeShon & Rench, 2009; Schmidt & DeShon, 2007) validate the model. In both cases, the integrated model accounts for findings that previous theories of self-regulation could not explain. Discussion focuses on additional tests to validate the model and on the value of incorporating formal models from the cognitive, learning, and motivational literatures to account for behavior in complex settings and over time.

© 2013 Elsevier Inc. All rights reserved.

Introduction

Behavior is inherently complex and dynamic (Atkinson & Birch, 1970). One reason for this complexity is that organisms act to seek or maintain numerous goals based on changing internal and external conditions (Austin & Vancouver, 1996). Humans can also think about the possible consequences – good and bad – of their actions and the likelihood those consequences will materialize using their understanding of the environment and themselves (Krantz & Kunreuther, 2007). These capacities, their operation (i.e., their dynamics), and their development are especially important when considering motivation and behavior in work settings (Kanfer, 2012; Mitchell & James, 2001). Perhaps not surprisingly, developing a firm understanding of the mechanisms involved in these processes is a challenge. Moreover, several scholars are concerned that path models describing the relationships among variables and the verbal explanations accompanying these models are insufficient for explicating and communicating theories of human behavior (Busemeyer & Diederich, 2010; DeShon, 2012; Farrell & Lewandowsky, 2010; Sun, 2008; Vancouver, 2012). To address this concern, these scholars suggest using formal theories because they can provide a more precise, transparent, and internally consistent approach to theorizing. In particular, computational models of

subsystem processes are a particularly appropriate type of formal theorizing because they can be simulated and provide predictions of behavior over time (Adner, Polos, Ryall, & Sorenson, 2006).

Toward that end, Vancouver, Weinhardt, and Schmidt (2010) recently presented a computational model of multiple-goal pursuit that describes processes for thinking and acting over time. In this model, they integrated theories of goal choice and goal striving, which tend to be considered in separate, middle-range theories (Klein, Austin, & Cooper, 2008). Moreover, the model was dynamic, predicting how individuals respond to changes in the environment brought on by their own actions as well as outside forces. The multiple-goal pursuit model (MGPM; Vancouver et al., 2010) represents an instantiation of a larger formal theory of self-regulation presented by Vancouver (2008). That theory, which borrows heavily from control theory perspectives of human behavior (e.g., Carver & Scheier, 1998; DeShon & Gillespie, 2005; Ford, 1992; Lord & Levy, 1994; Powers, 1978), describes mechanisms alleged to represent action, thinking, feeling, and learning processes. One of the more remarkable elements of the theory is that the basic mechanism for all these processes is the same. Specifically, at the core of control theory is a weighted difference function driving a negative feedback loop that arcs through the environment. This weighted difference function represents a self-regulatory agent, which is a very simple subsystem of the human system (Vancouver, 2005). Via this conceptualization, the complexity of human behavior arises from a combination of environmental dynamics and, more importantly, the organization of multiple agents within individuals.

* Corresponding author. Address: 2500 University Drive NW, Calgary, AB T2N 1N4, Canada.

E-mail addresses: vancouve@ohio.edu (J.B. Vancouver), justin.weinhardt@haskayne.ucalgary.ca (J.M. Weinhardt), vigo@ohio.edu (R. Vigo).

At this time, much scholarly work is needed to conceptualize and verify the specific nature and organizations of agents that might account for the variety of phenomena representing human behavior. By developing a working simulation of goal choice and goal striving, the MGPM represents an example of that kind of scholarly work. However, that model was limited to the dynamics of the environment and information (i.e., signals) flowing through the individual. It did not address dynamic processes that might occur within the individual (e.g., processes that change what or how individuals process signals). Such mechanisms would lead to the relatively permanent change in individuals commonly referred to as learning (Weiss, 1990).

Although applied psychologists often consider learning the province of training and development, learning theories have also long been of central interest to organizational behavior and motivation researchers (Pinder, 2008). Indeed, learning is an important process in human functioning that supports goal striving, goal choice, and other self-regulatory processes (Sitzmann & Ely, 2011). For instance, expectancy and social cognition theories (e.g., Bandura, 1986; Porter & Lawler, 1968; Vroom, 1964) assume much of the variance in motivation is based on learned beliefs about environmental contingencies, properties of the self (e.g., beliefs about capacity), and experience with outcomes. However, the bulk of the research on learning within organizational science does not address the sub-system processes that lead to learning, focusing instead on a higher-level analysis. For example, research has shown that error management training increases learning (Keith & Frese, 2005). Yet, as Keith and Frese note, the specific processes on which error management training capitalizes are less well known. Here we offer an account of the subsystem processes that lead to learning and intra-individual change.

Of course, formal models of learning mechanisms are common in cognitive psychology (Young & Wasserman, 2005) and some have appeared in organization science (Gibson, Fichman, & Plaut, 1997; March, 1996). We incorporate some of that knowledge here. Integrating these mechanisms with modern conceptualizations of goal choice and goal striving provide a more comprehensive theory of motivation and self-regulation. Specifically, we want to explain how an individual develops understandings of themselves and the environment that are useful for self-regulation. For example, individuals might create beliefs of the effectiveness of actions (i.e., expectancies; self-efficacy) or future conditions, like when professors learn to anticipate the kinds of questions they might receive on a set of material presented to a class or the reviews they might get on a journal article and use those anticipated events to adjust their presentations or articles. A better understanding and integration of these dynamic processes is likely to facilitate the development of interventions targeted at improving self-regulation (Boekaerts, Maes, & Karoly, 2005; Vancouver & Day, 2005). Moreover, if the learning subsystem is consistent with the goal-choice and goal-striving subsystems, then the representation of self-regulation will be conceptually parsimonious as well comprehensive (Vancouver, 2008).

Toward these ends, the goals of the current project are fourfold. First, we seek to add understanding of individual change to the MGPM, which at this time only includes behavior and environment change. Second, we use that understanding to account for how individuals might handle uncertainty in the environment as well as learn helpful information about oneself that might be useful for planning and making decisions. Third, we seek to maintain parsimony by using the same core concept (i.e., the self-regulatory agent) used in the MGPM. Finally, we accomplish the above formally by using a computational representation of self-regulation. In the following sections, we review the core concept in our approach and how it is used in the MGPM. We then review learning concepts and some formal modeling concepts found in the

cognitive literature. Next, we explain how learning might support multiple goal pursuit, incorporating the learning agents into the MGPM. Finally, we assess the validity of the resulting model by assessing its ability to account for the phenomena it purports to explain.

A computational theory of self-regulation

In recent years, there has emerged an increasing desire to develop a comprehensive, integrative theory of human work motivation and behavior (Locke & Latham, 2004; Pinder, 2008; Steel & König, 2006). Toward that end, self-regulation theories have shown promise (Diefendorff & Chandler, 2011; Lord, Diefendorff, Schmidt, & Hall, 2010). Self-regulation theories highlight the agentic, goal-directed, quality of behavior (Bandura, 1997; Vancouver & Day, 2005). A central construct in self-regulation theories are goals, which are internally represented desired states (Austin & Vancouver, 1996), and a central process in these theories is discrepancy reduction (Vancouver, 2005).¹ In particular, the individual maintains or achieves goals by perceiving the states of variables and acting on discrepancies between the perceived states and the desired states (i.e., goal). The actions affect the states of interest, moving them toward the goals and thereby reducing the discrepancies. Because the actions partially determine the states of the variables and perceived states partially determine the actions, the process reflects a feedback loop. Moreover, because discrepancies signal actions that reduce the discrepancies, the sign of the loop is negative. When operating properly, a negative feedback loop keeps “regular” or “controls” the variable’s perceived state much like a cruise control system in a car keeps regular or controls the speed of the car as perceived by the car’s speedometer (Vancouver & Day, 2005). It is this process that accounts for the self-regulation or control theory labels often used interchangeably in the literature (e.g., Carver & Scheier, 1981, 1998).²

Fig. 1 represents the basic negative feedback, or discrepancy-reducing, loop. The figure makes explicit the key functions needed for self-regulation. Specifically, the input function translates signals external to the system (i.e., inputs) into a single signal used by the system. Typically, this involves translating stimuli (s) indicating the state of some variable (v) into a perception (p) of that state. The perception is then available to other functions in the system. Mathematically, we represent the input function as a product of two vectors: a vector of inputs and a vector of weights for the inputs. Multiplying the two vectors gives a scalar (i.e., single) value much like a regression equation gives a single prediction based on a set of weighted inputs. Typically, in self-regulation models, the inputs are actually the outputs from other functions or nodes (Powers, 1973). We should also note that connectionist (i.e., neural network) models describe a similar arrangement, except that the input function is called a node or processing unit (Gibson et al., 1997). Given that the outputs from input functions are labeled perceptions (p) in self-regulation models, Vancouver (2008) represented the input function as follows:

$$p = \mathbf{w}_s \mathbf{p}_s, \quad (1)$$

¹ Another central process in self-regulation theories is discrepancy production via the adoption of goals that are beyond current conditions. Some suggest that discrepancy production is a byproduct of discrepancy reduction (e.g., Scherbaum & Vancouver, 2010); whereas, others suggest that discrepancy reduction is dependent on the discrepancy production, making discrepancy production more interesting (Bandura, 1997). In this paper, we merely assume discrepancies exist. We do not address the possible processes that produce them.

² The loop can also represent the motivation for one-time achievements, like obtaining a Ph.D., where the regularity element is less obvious and control seems to qualify the actions as opposed to the perceptions. This achievement type of context appears to have created some semantic confusion (Vancouver, 2000; Vancouver & Scherbaum, 2008).

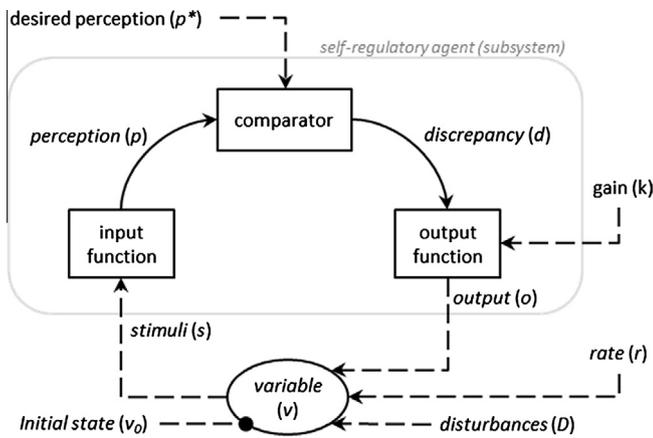


Fig. 1. Self-regulatory agent as part of cybernetic negative feedback loop.

where the bold letters refer to vectors of weights (\mathbf{w}) and perceptions (\mathbf{p}), respectively, the subscripts refer to the stimuli the perceptions are based on, and the italicized p refers to a scalar value.³

The idea behind the function is that multiple, simple sensory signals need to be combined to create complex perceptions (Powers, 1973). Moreover, those sensory signals or their combination may imperfectly represent the environment (Grush, 2004). What is a primary concern for neural network researchers are the structures and processes that determine the weights and thus the nature of the results from these functions. For example, one might be interested in understanding what information a professor uses, and how the professor combines it, to infer a class does not following the lecture. Fortunately, for many theoretical problems the details of how stimuli regarding the state of the variable (v) are translated into perceptions (p) of that state are essentially isomorphic and thus of relatively little importance for the particular theoretical issue or phenomenon. For example, a student and his or her advisor both perceive a schedule with four courses as a schedule with four courses. In these cases, one merely needs to describe a function that translates the variable state directly into the perception (e.g., $p = v$) or represents a translation with some bias (e.g., $p = b + v$ or $p = bv$, where b is a bias term). We will primarily be using these simpler functions here.

Perceptions are relatively meaningless without further processing. To give them meaning, they are assessed against a standard, referent, threshold, or goal level (Carver & Scheier, 1998; Powers, 1973; Vancouver, 2008). This role is represented in the comparator function in the negative feedback loop (see Fig. 1). In this function, perceptions derived from the input function are compared with a desired perception (p^*). The function returns the difference between these two values, called error or discrepancy (d). For example, an advisor might see that a student's schedule only includes four classes instead of the five classes that constitute a full load. This difference of one class represents a discrepancy of the advisor's five-class goal for advisees needing a full load.

The discrepancy value is the heart of self-regulation models of motivation. It represents tension (Lewin, 1951), which the negative feedback loop reduces via its structure (i.e., the organization and operations of the functions). The exact nature of the comparator function can vary depending on the nature of the goal (e.g., maximum, minimum, or optimal; Edwards, 1992) or requirements of the theorizing. For example, neural network models often use a sigmoid to represent this threshold function (Thomas & McClelland, 2008). They also use optimal goals, where error on either side of the standard motivates change. A simpler representation is a conditional difference function, where only positive differences ($d > 0$) are forwarded; negative differences ($d < 0$) are truncated to a null signal (0). The conditional represents the idea that only perceptions below the desired state motivate action (i.e., a maximal goal). A conditional that represents a minimal goal is one where only perceptions above the desired state (e.g., noise levels above a threshold) motivate action. To keep it general, we use $f(p^* - p)$ to represent some function, f , that uses the inputs (i.e., desired and actual perception) within the parenthetical in the function; however, we include the negative operator to make explicit that the function takes a difference between these two inputs. The specific functions used in the comparators in the subsequent models described here are provided in the appendices.⁴

Once a discrepancy is detected (or not), its degree of importance needs to be addressed. This is accomplished by the output function (see Fig. 1). The output function uses a weight (i.e., multiplier) called gain (k) to augment or dampen the discrepancy (d) signal from the comparator to create its output (o). This is the weighted difference function at the core of control theory (Vancouver, 2005). In psychological terms, the gain is often thought of as an importance weight (e.g., Hyland, 1988; Schmidt & DeShon, 2007), but it also determines whether or when discrepancies are passed on as outputs (e.g., Vancouver, Putka, & Scherbaum, 2005). Depending on the nature of the action (i.e., discrete or continuous), the output function might also categorize the weighted discrepancy signal (e.g., 0 or 1). For example, the advisor may or may not think that carrying a full load is critical for students, which determines whether the advisor mentions a small discrepancy arising from only four of the desired five-course load in an advisee's schedule.

The output (o) from the output function represents a set of actions by the individual on variable (v). An important aspect of the feedback loop's focal variable is that it is dynamic, which means that it retains its value over time (Powers, 1978). For example, adding a course to a schedule increments that schedule by one course (e.g., from four to five). Because dynamic variables involve change from one state to another, they need to have an initial value (v_0). We used a black dot to connect the initial value to the variable (see Fig. 1). In addition, to indicate a variable is dynamic, we use an oval to encase it.

One way to represent dynamic variables mathematically is in terms of change (Δ), where the domain of the equation includes the factors that change the variable. In particular, when depicting negative feedback loops, one can dichotomize the factors affecting a variable between the effects from the loop and the effects from

³ Alternatively, coming from a neural network tradition, Gibson et al. (1997) described the same function in terms of the output from a node using the following notation: $output_i = \sum_j w_{ij} output_j$, where the i subscript refers to the focal node and the j subscript refers to the nodes whose outputs provide input to the focal node. Note that this function represents summing across potential output signals from multiple j units, where there is only one i (although this i unit could become just another j unit for another node). Moreover, the use of the summation symbol and the subscripts represents an alternative way to describe the multiplication of two vectors, making the two functions equivalent. Other connectionist (i.e., neural network) modelers use a instead of $output$ because the signals are referred to as activation levels (e.g., Thomas & McClelland, 2008). This later notation shows up later in this paper.

⁴ All the models described here can be downloaded from the first author's website (<http://www.ohiopsychology.com/Research-Lab-Index.html?lab=22>). They were rendered in Vensim®, a system dynamics modeling platform. One can simulate them in VensimPLE®, which is a free version of Vensim for educational or personal use that can be downloaded at (www.vensim.com). The VensimPLE® software can also be used for examining the specific equations used and for replicating the sensitivity analysis and experiments described here. The Vensim web site provides help and links for help on using the software and modeling more generally. Dr. Vancouver also has a web site devoted to supporting the understanding and use of computational modeling (<https://sites.google.com/site/motivationmodeling/home>). This site includes a tutorial based on a Vancouver and Weinhardt (2012) paper in *Organizational Research Methods*.

outside the loop, called disturbances (D). Disturbances might advance a variable to the desired state (e.g., the student adds courses on their own) or move it away from the desired state (e.g., a course is dropped by the college because of insufficient enrollments). Such effects can be complex and at varying rates. At this point, we represent them generally with the symbol D . On the other hand, we more explicitly represent the rate (r) of the effects of output (o) from the feedback loop as a weight that qualifies the output (e.g., ro); thus,

$$\Delta v = ro + D \quad (2)$$

For example, the rate (r) might represent the speed at which a class picks up the material presented by the instructor given some level of instruction, which is the instructor's output (o). Positive disturbances might occur as a result of help from a tutor for the students, but negative disturbances could include students forgetting previously learned material.

Note that output (o) is somewhat a function of the state of the variable (i.e., output is a function of discrepancy, discrepancy is a function of perception, and perception is a function of the state of the variable). Thus, the variable's state (v) is somewhat a function of itself, which is the definition of a feedback loop. Also, provided disturbances are not overwhelming, and both gain and rate are non-zero and positive, the actions from the loop will eventually bring (or maintain) the state of the variable to (or at) the desired state (p^*) as perceived (p); at which point the comparator function will be zero and thus cease action on the variable. That is, the above equations represent a discrepancy-reducing system.

The agent in the negative-feedback loop

When psychological theories that use the concepts depicted in Fig. 1 were introduced (Miller, Galanter, & Pribram, 1960; Powers, Clark, & McFarland, 1960), the theoreticians recognized that a single loop or subsystem could not account for much. Thus, these original theorists described hierarchies of subsystems. These subsystems often only include the three boxed functions shown in Fig. 1 (i.e., the input, comparator, and output functions), not necessarily the entire loop. To facilitate description of the theory, Vancouver (2008) suggested calling the combination of the three functions in the loop a self-regulatory agent (see the grayed rectangle in Fig. 1) because the structure represents an entity that pursues, via an output signal, its purpose of getting the variable state, as perceived, to the desired state.

Typically, theorists describe the multiple agents in terms of an action or perceptual hierarchy (e.g., Carver & Scheier, 1998; Powers, 1973). Specifically, the perceptual hierarchy label highlights the notion that the perceptions created from lower-level input functions feed into higher-level input functions. For example, the perception of an accepted paper somewhat determines the perception of a reputation. This conceptualization is also consistent with the multilayered notion of neural networks (Thomas & McClelland, 2008). The action hierarchy description focuses on the output side, where outputs from higher-level agents determine the desired perceptions (i.e., goals) or gains for other agents. For example, the completion and acceptance of high quality papers are the means to improving reputation. At the lowest level of the action hierarchy, the output signals determine muscle tensions, which transduce neural signals into action on the environment (Austin & Vancouver, 1996; Lord & Levy, 1994; Miller, 1978). For example, the words one wants to use to communicate about a study are translated into words on a document via parallel and sequential muscle tensions in the professor's fingers hovering over a keyboard.

Another important issue is that the variables monitored and acted upon need not be in the person's environment; they can also

be variables within the person (e.g., glucose levels, affect levels). That is, a variable is considered external to the agent (i.e., subsystem), but not necessary to the larger system in which the agent is embedded. For example, the professor may feel the need to eat while typing up her next article.

Given the comprehensive nature of theories based on the negative feedback loop (Lord & Levy, 1994; Vancouver, 2005), sub-disciplines within psychology and related fields have often only focused on certain types of agents. For example, much of the original work on control subsystem agents focused on physiological control systems (Cannon, 1932). Since then, the largest amount of work has been on motor control (Jagacinski & Flach, 2003). Other work has focused on the perceptual side of the theory (e.g., Grush, 2004).

Perhaps of most interest to applied psychologists, negative feedback processing systems have been implicated in many higher-order cognitive and social processes (e.g., Carver & Scheier, 1998; Cropanzano, James, & Citera, 1993; Klein, 1989; Lord & Levy, 1994). However, most of this later work has been informal. That is, theoreticians have used natural language, not mathematics, to describe the theory. Moreover, after a recent review of self-regulation theories, DeShon and Rench (2009) claimed that no self-regulation theory was up to the task of explaining important aspects of goal-pursuit behavior. In particular, self-regulation theories tend to be vague regarding processes for handling "multiple, potentially conflicting, goals that are pursued in dynamic environments that require flexibility and robustness" (p. 218). In addition, there is controversy surrounding descriptions of (a) the perceptions individuals attempt to control, (b) the connections or organization among the agents, and (c) the development process (i.e., how agents evolve and develop connections) (e.g., Bandura & Locke, 2003; Vancouver, 2005). DeShon and Rench (2009) also noted that due to the complexity and dynamics involved, the field would benefit from computational models of the theorized processes (see also, Vancouver, 2012).

Modeling processes of multiple goal pursuit

Fortunately, since the DeShon and Rench (2009) review, several advances in self-regulation theory have emerged to address previously unspecified processes using a computational modeling approach (e.g., Scherbaum & Vancouver, 2010; Vancouver et al., 2010). In particular, the MGPM addressed some of these processes by going beyond the typical action hierarchy (Carver & Scheier, 1998) to a more networked structure that can account for thinking and decision-making processes like those involved in goal choice (Powers, 1973; Vancouver, 2008). Specifically, the MGPM described expectancy agents, whose output represented expectancies for reaching each task goal as a function of (a) the changing discrepancy between the current state and the goal, (b) a resource limit, like a deadline for achieving the goal, and (c) an expected lag (i.e., time to reduce a unit of discrepancy) for achieving the goal. The MGPM also included a time agent that compared the current time to the deadline for the tasks to create an output used as a threshold value by the expectancy agent and a choice agent that used outputs from task agents and expectancy agents to create and compare expected utilities to determine goal choice.

The MGPM is a general model of multiple-goal striving and choice; however, model testing involves specific protocols that require specific representations in an instantiation of the model (e.g., the disturbances used in a protocol). Relatively few rigorous protocols exist that examine multiple goal striving. An exception is a study by Schmidt and DeShon (2007), which served as a testing platform for the MGPM. In that protocol, the researchers asked participants to construct class schedules for students in two colleges over a 30 m period. The overall objective was to provide schedules

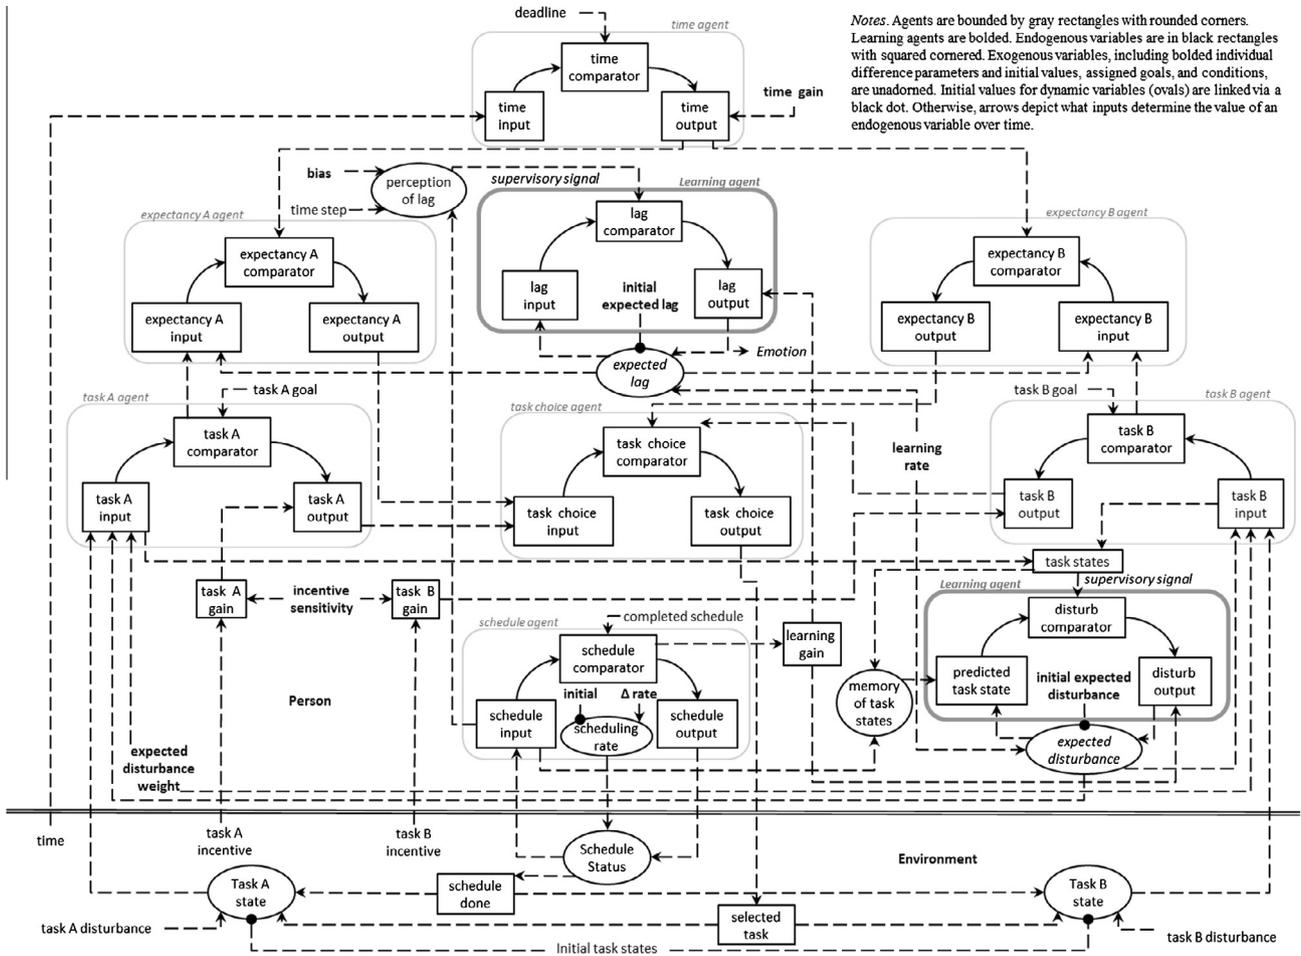


Fig. 2. The multiple-goal pursuit model with learning agents.

to students seeking them (i.e., get number of students needing schedules to zero in each school). Thus, each college represented a task with a goal of zero. Initially, both colleges had five students seeking schedules, but over time, more students queued up for schedules in each school via an unpredictable timetable. These additional students represented a disturbance to the variables participants were seeking to control. The primary objective of the study was to determine on which goal (i.e., which college) participants worked over time.

Given the specifics of the protocol and the self-regulatory agents needed for the multiple goal pursuit model described above, Vancouver et al. (2010) built and simulated a computational model of multiple goal pursuit. Fig. 2 presents the structure of that model; the mathematics are presented in Appendix A. The model includes the self-regulatory agents they felt were involved in determining goal pursuit across the course of the 30-min study session. This includes the agents mentioned above (task and expectancy agents for each task [i.e., school], the time agent, and the choice agent) as well as one more task agent, called the schedule agent, which was the subordinate task needed to reach the schools' goals. The figure also shows a double line between the variables and processes within a person (i.e., top) and those occurring in the environment (i.e., bottom). The environmental variables and processes as well as the exogenous⁵ desired states (i.e., the task goals, deadline, and completed schedule state) represent the protocol. The

⁵ Exogenous variables are ones without any arrows pointing at them. Endogenous variables have arrows pointing at them.

other exogenous variables are bolded in the figure; they represent possible individual difference variables and can be free parameters in the model. Some of these, as well as two learning agents (bolded gray rectangles), are additions to the MGPM that we described in a subsequent section.

The original MGPM (Vancouver et al., 2010) accounted for three key phenomena found by Schmidt and DeShon (2007). First, the simulated model, like most of the participants, spent at least the first half of the time working on the goal with the greatest discrepancy. Because the actions reduced the discrepancies, this meant that the model and the participants switched frequently between the two goals. The second major finding from the study was that in conditions where incentives for the two goals differed, the participants preferred to work toward the goal with the higher incentive. Likewise, when differing incentive conditions were represented in the model via Task A incentive and Task B incentive values (see Fig. 2, near bottom, left side) the model also preferred to work on the higher incentive task, provided the incentive sensitive parameter was greater than zero. The final major finding was that, for some of the participants, the preference for working on the task with the greater discrepancy reversed as the deadline approached. That is, during the latter part of the study session some individuals tended to remain working on the low discrepancy task until reaching the goal or beyond (creating a buffer of banked schedules) before switching to the other goal. The model also revealed this reversal of the role of discrepancies toward the end of the simulated session time given certain values in the model parameters. Using the

computational model, Vancouver et al. (2010) demonstrated that the expectancy agents, who used as inputs information about the approaching deadline, the goal discrepancies, and the expected lag for creating schedules, could explain the reversal.

The MGPM not only accounted for important effects found in a study of multiple-goal pursuit (Schmidt & DeShon, 2007), but also integrated concepts from decision-making theories (e.g., subjective expected utility), economic theories (e.g., hyperbolic discounting), and motivation (e.g., control theory). It also explicitly added dynamic elements to previously conceptualized static constructs (i.e., expectancy and valence). However, there were patterns of behavior observed in Schmidt and DeShon and other studies of dynamic goal striving (e.g., DeShon & Rench, 2009; Louro, Pieters, & Zeelenberg, 2007) that the model did not address. In particular, the researchers observed that individuals exhibited goal-directed behavior toward goals already achieved. Moreover, the MGPM depicted expected lag as a direct function of actual lag (i.e., the actual time it takes to reduce a unit of discrepancy). Yet, beliefs like expected lag likely emerge from interactions with the environment. To address these issues, we added the learning agents shown in Fig. 2. We discuss these agents and learning more generally next.

Learning and learning agents

Originally the province of behaviorists (e.g., reinforcement theory), learning theories have become highly cognitive and computational (Weiss, 1990). For example, neural network models have become a mainstay for both cognitive psychology and for computational modelers (Young & Wasserman, 2005). Among the various formal models of learning, most can be classified as supervised or unsupervised, depending on the context. Unsupervised learning models describe how learning might occur when there is no information regarding what might be correct. It is inefficient compared to supervised learning, where there is corrective information that can be used by the learner. Importantly, supervised learning models are described in terms that are remarkably similar to the negative feedback loop described above. For example, Young and Wasserman (2005) describe supervised learning models in the following way:

“In supervised learning, the learner is modeled as producing a response, comparing that response to the one that should have been produced (this information is provided by a ‘teacher’), computing the error (the difference between the actual response and the correct response) and then altering its internal representations in order to reduce future error. The response does not have to be overt, but may be an internal mental state (e.g., the anticipation of a future stimulus configuration). Likewise, the teacher may be an instructor, the environment (e.g., the outcome of an action) or an internally generated signal from another brain system” (p. 168).

The above description is foundational for several theories of learning, including the leading animal learning model (i.e., Rescorla & Wagner, 1972) and connectionist (e.g., neural network) models (e.g., Gluck & Bower, 1988; Thomas & McClelland, 2008). Of interest is that these “supervised” models do not require an external teacher, as noted in the last line of the quote. In particular, our model will assume the supervisor signal comes from the environment, but not necessarily from an instructor or employee’s supervisor (though it can). Likewise, it is interesting that the “response” is not necessarily some overt action, but it might be some anticipated perception (i.e., future stimulus configuration). These kinds of supervisor and response signals allow easy incorporation of

the learning model within control theory models of self-regulated behavior.⁶

Consider, for example, the issue of expectancies. Most generally, expectancies are beliefs in contingencies (Lewin, 1951). Applied psychologists have used several labels to refer to expectancies that are likely useful for understanding human behavior, motivation, and decision-making, including self-efficacy (Bandura, 1977), capacity and context beliefs (Ford, 1992), outcome expectancies (Carver & Scheier, 1998), and instrumentalities (Vroom, 1964). Moreover, Austin and Vancouver (1996) noted that these kinds of beliefs are central when attempting to understand goal processes (e.g., goal choice, striving, and revision). Yet a central question might be how can one understand the processes by which these beliefs arise? Bandura (1997) and Porter and Lawler (1968) both suggest that the beliefs change as a function of experienced past performance – a finding substantiated in a recent meta-analysis (Sitzmann & Yeo, 2013) – but are unclear regarding how the change occurs. The supervised learning model would imply that one way they arise is via adjustments to a nascent expectancy belief based on a comparison of a prediction using the nascent belief with a reality, or at least a perceived reality, coming from the environment. Indeed, we use the basic supervised learning mathematical model, known as the delta-learning rule (Widrow & Hoff, 1960), to account for the formation of how individuals learn both capacity and context beliefs.

Delta-learning rule

Several forms of the delta-learning rule exist (Anderson, 1995), but we reproduce a connectionist version here from Thomas and McClelland (2008). Specifically,

$$\Delta w_{ij} = \eta(t_i - a_i)a_j, \quad (3)$$

where w_{ij} represents a weight indicating the association between two units i and j , η is the rate of learning, t_i is the target value of unit i (i.e., the supervisor signal), a_i is the current value (i.e., activation level) of unit i , and a_j is the value of unit j . The rule represents a process that moves, by some degree (η), the value of unit i toward the target value (t_i), given the activation level of unit j , by changing the weight of unit j ’s contribution to unit i ’s activation. In most connectionist models of learning, the weights are dynamic variables, constituting the memory or knowledge of the system (Lewandowsky & Farrell, 2011; Thomas & McClelland, 2008).

What might also be clear from Eq. (3) is that it represents a weighted difference function like that found in control theories (Vancouver, 2005). The similarity between the delta learning rule and the control theory formulation can be readily seen by substituting the weighted discrepancy $f(p^* - p)k$ from the comparator and output functions for output (o) in Eq. (2). The resulting equation for change to the variable is:

$$\Delta v = rf(p^* - p)k + D \quad (4)$$

Note the similarity between Eqs. (3) and (4). That is, the learning rate (η) is a rate (r) that represents the effect of the agent per time step, the supervisor signal (t_i) serves the role of the goal or reference standard (p^*), the current unit i value (a_i) is the perception (p) arising from the unit’s input function, and the unit j value

⁶ Learning has been incorporated in control theories of human behavior since their inception (Powers et al., 1960), but the conceptualization was different. Specifically, Powers (1973) conceptualized learning as random change to functions in the network of agents based on the total error in the network. The approach taken here is local, focusing on signed change to a specific agent as a function of the sign of the error in the associated learning agent. More recently, Gibson et al. (1997) described a computational model of learning using a more local, neural network conceptualization with some reference to control theory concepts, but the reference was oblique and the learning component more complex than the one we describe here.

(a_i) is the gain (k) for the agent. That is, the equations are equivalent except for the lack of a disturbance term and the general function (f) symbol. The function symbol is unnecessary because the simple difference function is sufficient (i.e., it is an optimizing system agent) and because nothing else should be disturbing the weight, D is unnecessary.⁷ Thus, the delta-learning rule represents a specific type of self-regulatory agent; one we call a learning agent. To complete the analogy between the delta-learning rule and a self-regulatory agent, one only needs to understand that the weight (w_{ij}) is the controlled variable (v) for learning agents.

Using the delta rule and its generalization, known as the *back propagation algorithm*, researchers in cognitive science have modeled a wide variety of psychological phenomena ranging from classical conditioning (Schmajuk & DiCarlo, 1991), memory (Vogel, 2005), human categorization (Kruschke, 1992), and human personality (Read et al., 2010), to name a few. Indeed, these approaches have increased our understanding of adaptive behavior in humans. However, they mostly provide formal explanations for low-level cognitive phenomena (e.g., word recognition), often using large numbers of very basic components (i.e., artificial neurons). The application of such modeling to higher-order cognitive is less common (see March, 1996, and Gibson et al., 1997, for exceptions), but potentially relevant to a wide range of higher-level phenomena like self-regulation. We aim to alleviate this deficit with our proposed approach.

Applications of the learning agent

To get a better handle on the nature of learning agents, and how they interact with other self-regulatory agents, illustrations are useful. These illustrations also allow us to assess the validity of the model. Here, we describe applications of learning agents to the development of capacity and context beliefs (Ford, 1992) within the MGPM. Specifically, we model the mechanism for learning the efficiency (i.e., change over time) of one's actions on a variable of interest (i.e., a capacity belief), though it could be more generally used to form a belief about the time it takes something to occur (e.g., the completion of a paper for submission to a journal). We also describe how one might develop and use information about environmental impacts on a variable of interest (i.e., context beliefs). Specifically, we propose that individuals can learn about disturbances and use the learned value to adjust their perception of the variable an agent might be seeking to regulate. These two kinds of beliefs are useful for decision making and planning given the many concurrent goals individuals are likely regulating and the dynamics of the environment and actions on the variables (i.e., states) involved. For example, academics may not be able to predict the exact nature of the reviews an article might receive, but they can learn to anticipate that some work will be required to revise and resubmit. Likewise, they can learn that the first week of classes of any semester will be chaotic for students and faculty, suggesting that then might not be a good time to work on the revision or plan a site visit for a new research project.

Learning the efficiency of one's actions

The first application involves a process that would represent the development of the belief about how long it takes to accomplish aspects of the task, which Vancouver et al. (2010) termed *expected lag*. This belief can subsequently determine self-efficacy (Bandura, 1997) or expectancy beliefs (Vroom, 1964). In particular, lag, which is one over rate (i.e., $\frac{1}{r}$), is likely a common internal metric for

representing associations (e.g., the *efficiency* of ones' actions on the environment). In Fig. 2, expected lag is located slightly above and to the left of the center of the figure. It is a dynamic variable (i.e., retains its value over time), so we enclosed the label in an oval. In the MGPM, expected lag and each task agent's discrepancy feed into the respective task agent's expectancy agent to create a perception of the time needed to reach each goal (see the arrows pointing at the expectancy input functions). That is, expected lag is a weight (w) modifying the task discrepancy to create the perception the expectancy agent uses. However, unlike Eq. (1), where weights qualify a lower-level agent's *perceptions* of stimuli (p_s), the weight qualifies a *discrepancy* from another agent.⁸ In this way, expected lag scales the discrepancy into time units so that the perception within the expectancy agent is a representation of the time one believes is needed to reduce the discrepancy (i.e., reach the goal). For example, if there are currently five students seeking schedules, and the person expects that each schedule will take 2 min to complete, than the expected time to complete the outstanding schedules will be 10 min.

In the original instantiation of the MGPM, expected lag was a direct function of actual lag (i.e., the length of time it took an individual to complete a schedule and thus decrement the queue of students by one), plus a bias parameter to reflect individual differences in calibration between the actual lag and the belief the individual had for lag. Yet, the process by which such a belief arises within the individual may be important not only for understanding how humans adapt to their environment, but also the possible sources of bias in such beliefs.

To understand how an expected lag belief might be developed, consider our mapping of the delta-learning rule (i.e., Eq. (3)) to the negative feedback loop (i.e., Eq. (4)). To facilitate this mapping, Table 1 provides the names of the functions and variables used by the learning agent, the symbols from Eqs. (3) and (4), a description of the meaning of the functions and variables, and an example. To begin, expected lag is the variable (v) of a negative feedback loop (see Fig. 1) and is the dynamic variable for a learning agent (see Fig. 2). Thus, expected lag changes based on *lag output* times the *learning rate*. The lag output is a function of (a) *learning gain* and (b) the discrepancy between the perception arising from the *lag input* and the supervisor signal calculated in the *lag comparator*. Learning gain is zero until there is a signal indicating the presence of a supervisor signal, which in this case is when the discrepancy in the scheduling agent's comparator drops to zero. The supervisor signal is a potentially biased perception of lag. In this particular case, where we are modeling the Schmidt and DeShon (2007) paradigm, the perception of lag arises from tracking the time to complete a schedule. It involves short-term memory to track the passage of time. Bias could then arise due to decay in short-term memory, a common element in theories and models of short-term memory (e.g., Wang & Arbib, 1993). Alternatively, bias might arise from an imperfect representation of the passage of time. Here we represented *bias* as a weight modifying the time interval that cumulates in the *perception of lag* variable (see Fig. 2). Bias values greater than one mean that people perceive that time passes slowly, and as a result, they over-estimate how long it takes to complete a schedule. Bias values less than one mean that people perceive that time passes quickly, and as a result, they under-estimate how long it takes to complete a schedule. Finally, the *lag input* function takes expected lag and multiplies it by the discrepancy reduced during a learning event. A learning event is a meaningful demarcation of action on the discrepancy. In this particular case, the event is bookended by the beginning of work on a schedule and the completion of

⁷ In some theories for some kinds of weights, decay is possible, and could be represented as a disturbance.

⁸ The use of discrepancy as an input to the expectancy agent is why the MGPM represents a networked structure as opposed to a hierarchical structure.

Table 1
Functions and variables for expected lag learning agent.

Function/var.	Eq. (3)	Eq. (4)	Meaning	Example
Expected lag	w_t	v_t	Belief of time to reduce one unit of discrepancy	Expected time to produce a schedule
Learning gain	a_j	k	Teachable moment (i.e., end of learning event)	Recognize schedule completed
Lag output	$(t_i - a_i)a_j$	$(p^* - p)k$	Correction for belief at time of teachable moment	Took 2 m to complete schedule but predicted it would take 1 m
Lag comparator	$(t_i - a_i)$	$(p^* - p)$	Difference between supervisor signal and predicted lag	Meaningless until schedule completed because no supervisor signal
Perception of lag	t_i	p^*	Time since beginning of learning event, as best remembered	2 m to complete schedule
Bias	na	na	Bias in perception of passage of time	Perception of time for learning event to occur
Lag input	a_i	p	Prediction of time learning event (subtask) ends	Prediction of 1 m to produce a schedule
Initial expected lag	w_{t_0}	v_{t_0}	Initial belief of time thought needed to reduce one unit of discrepancy	Belief it would take 1 m to complete schedule

the schedule, which reduces the task discrepancy one unit. Thus, lag input is expected lag times one (i.e., a unity function).

Given the above described agent, expected lag changes by a fraction of the difference between what the agent predicted the lag would be, beginning with an initial value (i.e., *initial expected lag*; see Fig. 2), and the perceived actual lag every time a lag is perceivable (i.e., at the end of the learning event, like the completion of a schedule). The fraction is the learning rate. For example, a researcher might expect it to take two months to write a paper. However, it actually takes four months to write the paper. Therefore, the next time the researcher writes a paper she or he believes it will take three months (i.e., a learning rate of .5 is applied to the difference between the two and four months and added to create a new expectation). Fig. 2 also includes another effect of the output of this particular learning agent, emotion, which we return to in the discussion section. Importantly, the addition of this learning agent to the MGPM only adds two individual difference parameters (i.e., learning rate and the initial value of the expected lag). After discussing the addition of the second learning agent, we evaluate the effect of these additions on the MGPM.

Learning to predict outside effects

The second application involves learning about the occurrence of disturbances that affect a variable one is attempting to control (i.e., regulate) and using that information to adjust perceptions of where the variable stands. Specifically, in control theory models, disturbances (D) represent the collection of other sources of change on variables besides the action of the agent (see Eq. (2)). Disturbances may be regular and predictable or irregular and unpredictable (Powers, 1973). When they are unpredictable, simple control system architectures (i.e., Fig. 1) provide a simple and efficient mechanism for regulation. Indeed, a central element of Powers' control theory model of human behavior is that individuals do not need complex understanding of the environment to survive. To demonstrate this, he used simple tracking tasks with unpredictable disturbances and found that individuals maintained desired states with aplomb (Powers, 1978).

However, more sophisticated control mechanisms include internal representations of elements of the environment to increase the efficiency of the control process. For example, reviewers for a specific journal often have a sense of how the journal's audience will likely respond to an article they are reviewing. When engineers develop sophisticated control systems, they can create the representations of the environment that the control systems they design might use (Ioannou & Sun, 1996); however, when the control system of interest is human, the question is how such a system might self-organize its own representation of the environment. This representation could include how long it takes the control systems to affect the environment, which is the role of

the learning agent that develops the expected lag. In this application, a belief about the possible effects of outside factors (i.e., disturbances like the interruptions to a professor's time in the first week of classes) is developed.

In the lower right quadrant of Fig. 2 is an agent used to learn the degree to which disturbances affect the variable of interest (e.g., the task state). In the Schmidt and DeShon (2007) protocol, disturbances were new students queuing up for schedules. The *expected disturbance* reflects the effect of these new students on the predicted states of the task. As with expected lag, expected disturbance is a function of the output from the learning agent, the learning rate, and the expected disturbance's current value, initialized via *initial expected disturbance*. To simplify the model, we used the learning rate from the previous learning agent.⁹ The learning agent's output is a function of gain, which is a signal that the learning event (e.g., some aspect of the task) was completed, and the difference between the observed task states and the predicted task states. The observed task state, which is the supervisor signal, is a combination of the perceptions of the task states. The predicted task state is based on the memory of the task states since the beginning of the event (i.e., obtained from schedule agents' input function) and the expected disturbances. For example, School A might have had three students and School B had two students waiting for schedules when one began to work on a schedule. This would create a memory of task states equal to five. If expected disturbances are two, then once the schedule begun is finished, the predicted task state would be six because the person expects one of the original five students now has the schedule just completed, but two more queued up while working on the schedule. If the actual total number of students waiting for schedules in the two schools was seven when the schedule was completed, the individual would adjust their expected disturbance up. Thus, via this learning agent, we suggest that *expected disturbance* changes to reflect the change in task state per event not attributable to one's action.

The next issue is to represent how the existing control mechanisms use this expected disturbance to help regulate task states. The primary concern inherent in control theory logic is maintaining a variable at a desired level (Powers, 1973). Disturbances move the variable from this desired referent. Thus, if one can anticipate those disturbances, or simply their impending effects on the variable, one can potentially buffer the variable from those disturbances (Ashby, 1958). This is particularly important when an individual is engaged in multitasking because as the individual engages in one task, disturbances that increase the discrepancies between the current state and desired state on the non-focal tasks can begin accruing on other tasks. In our example of academics, they must allocate their time between writing their papers and

⁹ In traditional models of supervised associative learning, rates may be a function of the stimulus (Wills, 2005). Thus, one could potentially extend the model to include a distinct learning rate parameter per agent.

tasks involved with the first week of the semester. If they anticipate little opportunity to write during the first week of the semester, they may work harder on the paper the week before to meet their goal for getting the paper done.

Consistent with Vancouver's (2008) argument that the input function determines the content of a self-regulatory agent, we also include in this function a term for anticipated disturbance. That is, anticipated disturbances could be a weighted input within the input function vector (see Vancouver, 2008). The input function weight for the anticipated disturbance is a parameter we call *expected disturbance weight* (located in about the center of the lower left quadrant of Fig. 2). Because this addition affects perceptions of the state of the variables, it affects goal discrepancies and, most importantly, it effects when the individual thinks the goal has been met (i.e., not until the task state is beyond the goal by the amount of expected disturbance). Thus, this new addition only adds two parameters: the weight by which expected disturbances are multiplied in the task input functions and the initial expected disturbance.

Validating the model

Assessing a model's validity, like theory testing, involves several methods spanning protocols and researchers (Vancouver & Weinhardt, 2012). An initial assessment is whether the model can produce the phenomena it is purported to explain. In the present case, we combined a well-established computational model of supervised learning with a computational model of multiple-goal pursuit to provide a more comprehensive computational theory of self-regulation. The combination model should presumably account for how an individual could develop a sense of their efficiency at a subtask and the effects of disturbances on task states.

To assess the validity of a model built with this combination of elements, we simulated the model depicted in Fig. 2 (see also Appendix A or downloadable model⁴) and compared the behavior of the model with the behavior observed in the study by Schmidt and DeShon (2007). We then conduct sensitivity analysis on the key parameters added to the MGPM. We also experimented with the model to test its capacity beyond specific protocols (Davis, Eisenhardt, & Bingham, 2007). These tests allow us to evaluate the model more generally. Finally, to cross-validate the model we created a new instantiation of the MGPM to represent a study described by DeShon and Rench (2009). In both paradigms examined, participants appeared to develop beliefs regarding the lag of actions on subtasks involved in realizing higher-level goals, as well as beliefs about disturbances to the variables the participants are attempting to control. Moreover, both studies describe phenomena that traditional self-regulation theories cannot easily explain.

Accounting for phenomena

To see if a model can account for the phenomena that it purports to explain, simulations are very useful (Davis et al., 2007). To simulate the model, one needs values for the exogenous variables. As mentioned, several of these values come from the

protocol instantiated (e.g., goal levels). Others are parameters that come from data. For example, Schmidt and DeShon (2007) observed an average initial actual lag just shy of a minute (i.e., 59.38 s) that improves to just shy of 40 s by the end of the session. These values reflect the fact that participants got more efficient in producing schedules during their session. We used them to determine the initial and change in rate values for the simulations. Some other parameters (e.g., incentive sensitivity; time gain) were not of interest for the current simulations and thus set to their simplest values (e.g., 1). Table 2 lists the model parameters that required values for our simulations. As noted above, initial expected lag represents the time individuals might think it takes them to move the variable of interest one unit. We used an initial expected lag of half a minute because individuals tend to be overconfident (Dunning, Heath, & Suls, 2004) until they have more experience (Mitchell, Hopper, Daniels, George-Falvy, & James, 1994). To simplify interpretations, we set bias to one, which meant there was no bias in the perception of the passage of time. For the learning rate, we used 0.04 because connectionist modelers typical find that values between 0.03 and 0.05 provide optimal fit in experimental protocols examined (Fagot, Kruschke, Depy, & Vauclair, 1998; Hutter & Legg, 2008; Kruschke, 1992). For the initial expected disturbance, we assumed participants begin with a conservative estimate of no disturbance, but use whatever expected disturbance value they developed. Therefore, we set the value for initial expected disturbance to zero and the value for expected disturbance weight to one. All these values are the default values listed in Table 2 and are the values given in Appendix A and the downloadable model.

Using the Vensim[®] Professional Version 5.10b, we ran a simulation of the model shown in Fig. 3.⁴ Fig. 3 shows the actual (dotted line) and expected (solid black line) lag over time of a simulated individual. The trajectory shows the expected lag increased from its initial value of half a minute until it reached a reasonably accurate assessment of lag. The upward bias in the expected lag is because actual lag kept decreasing and there is, forgive us, a lag in the learning. Nonetheless, this simulation demonstrates that the model is capable of describing a method by which individuals can learn how long it takes them to complete a subtask as well as continually adjust this assessment.

The second learning agent added to the MGPM model represents how an individual might develop a sense of the disturbances affecting the task states. In the Schmidt and DeShon (2007) protocol, 36 students began or joined the queues over the 30 m of the session. These 36 students were disturbances to the task state variables. Given the rate of scheduling, based on the average initial and changing rate across all the participants found in the study, the simulated individual completed 35 schedules over 30 min. Given that a completed schedule marked the learning event, the expected frequency of disturbance per completed schedule would be 1.03, which the dashed gray line in Fig. 3 represents. Fig. 3 also shows what happens to the expected disturbance over time when it initially began at zero. In particular, because of the sporadic nature of the disturbances, expected disturbances fluctuates somewhat, but eventually converged to a value near one. This finding shows the model's capacity to learn expected disturbances.

Table 2
Model parameters.

Parameter	Values tested; default	Meaning
Initial expected lag	0–120 s; 30	See Table 1
Bias	0–6; 1	See Table 1
Learning rate	0–1; 0.04	Degree of change in belief given difference between a prediction and an observation
Initial expected disturbance	0–3; 0	Initial belief regarding outside influences on task states
Expected disturbance weight	0–3; 1	Weight applied to expected disturbance in task input function

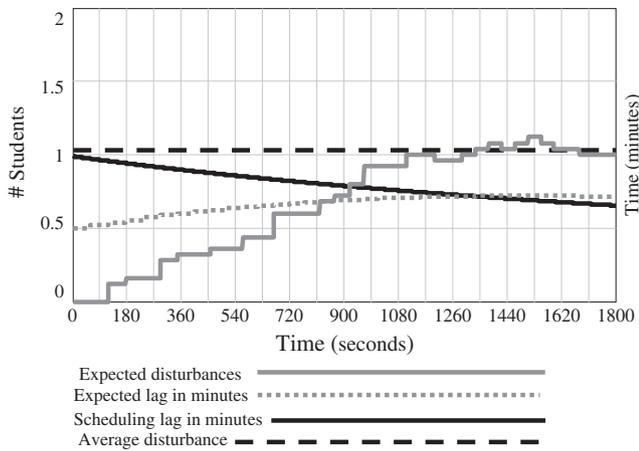


Fig. 3. Expected lag and expected disturbances over time.

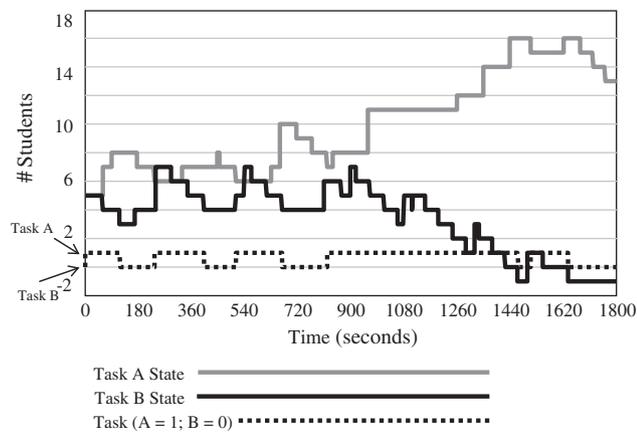


Fig. 4. Model results over time (Task A incentive twice Task B incentive). When the dotted line is up that indicates the simulated individual is working on Task A. When the dotted line is down that indicates the simulated individual is working on Task B.

To provide additional confirmation of the validity of the model to account for the phenomena the model claims to explain, we also examined the goal-directed behavior of the model. In the simulation, the incentive for Task A was twice that of Task B. Fig. 4 shows the states of the two variables and the task worked on over time, which matched the behavior exhibited by many of Schmidt and DeShon's (2007) participants in this condition. That is, both simulated and real individuals worked on both tasks initially, though with a preference for Task A unless the discrepancy for Task B became large enough to counter the incentive's effect. However, about halfway through the work period, the preference for Task A remained until a schedule was banked (i.e., the state of the number of students in School A was -1) at around 1450 s into the simulation (see Fig. 4). At this point, the simulated individual switched back to Task B to complete a schedule. Interestingly, while that schedule was being worked on, two more students queued up to get schedules in School A. This returned the state of Task A to one. Thus, the simulation switched back to working on schedules for School A until the Task A state was returned to -1 . Finally, the simulation finished the session working on Task B.

In sum, the findings from these preliminary simulations of the MGPM confirm that the model, and particularly the learning components, produced behavior that is consistent with observations of phenomena. However, the behavior of models can be a function of the values the modelers choose for the parameters. Assessing the importance of those choices is the next step in model evaluation

(Busemeyer & Diederich, 2010; Davis et al., 2007). This is called sensitivity analysis, which we discuss next.

Sensitivity analysis

For the sensitivity analysis, we wanted to evaluate the effect of different parameter values within reasonable ranges to gage their relevance for determining model behavior and their possible role as individual difference constructs. Table 2 lists the range of values assessed. For example, we assessed the effect of the bias parameter across a range from 0, representing a lower bound (i.e., perceiving time as not passing at all; learning events occur instantaneously), to 3, representing a perception of time passing many times slower than it actually passes and well beyond the level of bias one might normally expect. As predicted, bias affected expected lag. Low numbers result in short expected lags and large numbers in long expected lags. The expected lags affected the timing of the reversal effect. Specifically, an individual represented with a negative bias (i.e., low numbers) never exhibits the reversal, whereas an individual represented with a positive bias (i.e., high numbers) exhibits the reversal early in the simulation. Recall, an individual with no bias (i.e., bias = 1) exhibits the reversal about halfway through the simulation. Thus, variations in the bias parameter in our model produced results similar to the bias parameter in the MGPM model, which Vancouver et al. (2010) noted could account for some of the individual differences observed in the behavior of participants in the Schmidt and DeShon (2007) study.

Unlike the bias parameter, variations in initial expected lag and initial expected disturbance had no effect on model behavior using the default settings for the other parameters (see Table 2). In all cases, the model acquired approximately accurate beliefs by about two thirds of the way through the simulation and these beliefs did not affect behavior until beyond the halfway point. Likewise, learning rate, which is typically bounded between 0 and 1 (Anderson, 1995; March, 1996), made little difference except when it was zero, which resulted in no learning and thus behavior was a function of the initial values for expectancy and disturbance. Indeed, we ran 1000 simulations where initial expected lag was pulled randomly from a uniform distribution ranging from 0 to 120 and where the learning rate was pulled randomly from a uniform distribution ranging from 0.01 to 1. Regardless of the learning rate or the initial expected lag, expected lag converged on the actual lag, typically very quickly. However, the exact point of the reversal changed as a function of expected lag. Because the task with the smaller discrepancy at any specific time varied, varying initial expected lag when the learning rate was low resulted in changing which task achieved its goal when incentives for the two tasks were equal. That is, the model could predict that one task might eventually get favorable treatment over another task, but it could not predict which without information about the differences in the value of the tasks (e.g., differential incentives).

Also of note, a learning rate of one produced perfect and immediate learning of the lag, but it produced an unrealistically oscillating expected disturbance belief due to the instability in disturbances over time. That is, when dealing with a probabilistic or noisy set of stimuli, high learning rates create oscillating weights. Learning modelers have long known this (e.g., Ratcliff, 1990), which is why learning rates are usually relatively small or models have learning rates decrease over learning trials (Anderson, 1995).¹⁰ Finally, changes in the expected disturbance weight affected the number of schedules considered sufficient before switching back to the task that the individual had neglected when

¹⁰ However, there are also cases where one-shot learning is observed, suggesting a learning rate of one (Wang & Arbib, 1993), which might be more reasonable for the expected lag learning agent.

incentives differed or time was running out. However, overall, the expected disturbance weight, learning rate, and initializing values did not dramatically change the patterns of behavior produced by the model (i.e., discrepancies still positively predicted choice at first and then negatively predicted choice later). This result indicates that the model is robust (Davis et al., 2007).

Experimentation

In addition to sensitivity analysis, Davis et al. (2007) described various ways to experiment with a simulation to further theory development and testing. A theoretical question of some importance to understanding self-regulation regards the role of feedback. For example, DeShon and Rench (2009) noted that definitions of self-regulation did not include the concept of feedback despite its heavy presence in control theories of self-regulation (Vancouver & Day, 2005). Moreover, they argued that feedback might not be necessary for people to self-regulate because they could use learned associations and perceptions of relevant aspects of the environment to make choices and engage in behavior. Stated another way, if one had an internal understanding (i.e., a mental model) of the environment, one could use that as a substitute for feedback from the environment. Interestingly, this is the main element of modern control theory (Ioannou & Sun, 1996), where engineers add models of the environment to control systems to improve the functioning of the controllers in noisy or low feedback environments.

The mechanisms described here provide a possible explanation for how humans might develop such internal understandings and use them to enhance self-regulation. To explicitly test this idea, we created a modified version of the model that not only included the learning elements described here, but used perception from the predicted task state function (see Fig. 2) as a substitute for feedback from the environment for Task B. Running this model had no discernible effect on the behavior of the system. Fig. 5 shows why. Specifically, it shows the actual and estimated state of Task B during the simulation. Note the estimated state is based on (a) the initial state of Task B, (b) the expected disturbance learned from observing the difference between predicted and actual Task A states, and (c) the expected effect of an action on Task B. The estimated state also required a memory capacity as suggested by DeShon and Rench (2009).

The above finding indicates that a self-regulation theory can explain goal-seeking behavior without information about the current state of the regulated variable; yet, it does not obviate feedback as an important concept for understanding self-regulation for two reasons. First, in this case, there is an observable variable closely paralleling the unobservable variable available

for developing a viable model of the effect of actions and disturbances – a development process based on feedback. Second, and more importantly, the internal model creates input for the agent (i.e., the estimated state of the variable), which at least partially determines the output of the agent (i.e., action), which is then used by the internal model, closing a feedback loop. The only issue is whether feedback is filtered through the environment or the model of the environment. This is a profound distinction, though not one that can be understood in terms of the presence or absence of a feedback process (cf., DeShon & Rench, 2009).

Cross-validation

Thus far, we have only described instantiating the multiple-goal pursuit with learning model using the Schmidt and DeShon (2007) paradigm. To further assess the validity of the extended model, and to further illustrate its generalizability, we cross-validated it with data from another paradigm. In particular, DeShon and Rench (2009) collected repeated-measures behavior from nine individuals attempting to manage two goals symbolizing hunger and thirst. Their task included choosing between three options in each of 200 trials: seek food, seek drink, or do nothing. During each 10-s trial, it took individuals time to reach the food or drink icons, depending on the distance of the icon from the center of the screen, which is the position of the individual at the beginning of the trial. Reaching the icon meant the food or drink was obtained and consumed, which helped quench hunger or thirst, respectively. Meanwhile, as time passed, hunger and thirst got steadily worse (i.e., were disturbed), though consuming the food or drink could more than make up for two rounds of this disturbance. Thus, if participants alternated between food and drink for each trial, they would readily quench their hunger and thirst (i.e., reach their assigned goals for these variables).

Importantly, DeShon and Rench (2009) found two effects they believed no self-regulation theory of the time could explain. In particular, and inconsistent with typical descriptions of control theories (e.g., Carver & Scheier, 1998), all but one individual let relative opportunity associated with the closeness of items sometimes trump relative deprivation when choosing which goal to pursue. That is, they strove for items that were closer even though the respective state was in better shape (i.e., had a lower discrepancy) than the state associated with the other item. For example, they might have gone for the food because it was closer than the drink even though they were thirstier than they were hungry. The second effect DeShon and Rench found was that individuals would select and strive for an item even though the state of the variable that item improved was at or above the goal level. For example, all but one individual continued to obtain food and drink even though their hunger and thirst states were above their goals (i.e., their appetites were satiated). The first issue can easily be handled by a variation on the MGPM and the second by a model that learns to use a belief of the disturbance to adjust perceptions of one's state relative to the goal.

We begin with the variation to the MGPM that handles the opportunistic behavior observed by DeShon and Rench (2009). The functions of the model are provided in Appendix B and the model can be downloaded from the first author's web site. Three conceptual differences distinguish this variation of the MGPM from the one presented in Fig. 2. First, the expectancies are associated with the actions needed to achieve the goals rather than the goals themselves. In particular, according to DeShon and Rench a property of the food and drink icons – their distance from the center of the screen – determined how long it took to reach each item. We used one (1) to represent the maximum distance one could travel in a trial (i.e., from the center of the screen to a corner of the screen), assuming one traveled in a straight line. Thus, if an item

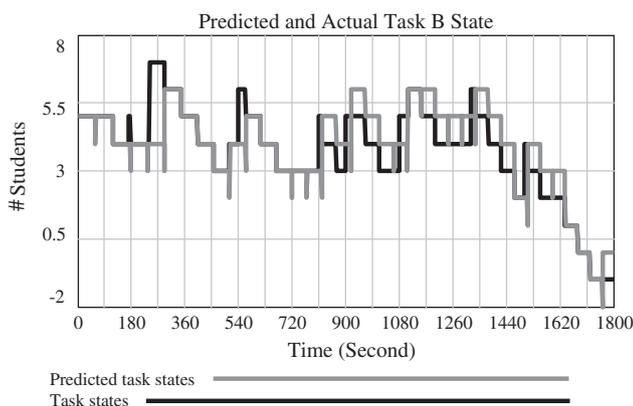


Fig. 5. Predicted and actual Task B state.

were near the corner of the screen, we could represent that distance as 0.9. In the DeShon and Rench paradigm, the distances differed randomly for each trial. In our simulation, we randomly drew from a uniform distribution between 0.1 and 0.9 to represent the distance for each item on each trial. The expectancy agent in the model weights this distance by the expected lag, which is a belief in the time it takes to traverse the screen. Later, we describe the model that develops this belief; however, for now we set expected lag to 10, which means the simulated individual believes it will take 10 s to traverse to the corner of the screen. When multiplied by the furthest item distance simulated, the simulated individual believes it will take 9 s to reach the item. This represents the *expectancy agents'* perception.

The second difference between this model and the original MGPM is that the referent for the expectancy agents, which is a resource limit, does not come from a time agent. Instead, it is a constant reflecting the time of a trial, which was 10 s. Given these values, expectancies never drop to zero because predicted time to reach the furthest item never exceeded 10 s.

The third difference between this model and the original MGPM is that three choices are possible. The choice agent could pick doing nothing (i.e., it outputs zero, which initiates no action) if the expected utility of both choices are zero (or identical, which is unlikely). Otherwise, positive outputs mean choosing food and negative outputs mean choosing drink – choices that affect the states of hunger or thirst, respectively, because the choice results in acting to obtain and consume the food and drink, respectively.

This simple representation of the MGPM uses trial as the time step and only models the choice process (i.e., the model did not represent within-trial behavior). A simulation of the model revealed that choices were a function of the relative opportunities associated with the food and drink items on each trial, as well as the relative difference in discrepancies in the hunger and thirst agents. This is because the opportunity was reflected in the expectancy agents' output and discrepancies in the hunger and thirst agents' output, all of which the choice agent used in selecting an item. In this way, the model accounted for the first observation that DeShon and Rench (2009) claimed was beyond existing self-regulation models.

To this basic model, we added the learning agents depicted in Fig. 2 (see Appendix C or the downloadable model from the first author's website) and an agent (i.e., the *action agent*) that acted on the discrepancy between the cursor and the chosen target (i.e., food or drink). The action agent replaced the schedule agent because traversing the computer screen, rather than creating schedules, was the action needed to achieve/maintain the higher-level goals (i.e., hunger and thirst). The expected disturbances learning agent was essentially identical to the agent represented in the first model and Fig. 2. However, the expected lag learning agent required a more sophisticated input function and memory support similar to that used for the expected disturbance-learning agent. That is, in addition to storing a perception of the time to reach the target (i.e., *perception of lag*), the individual would likely remember the distance traversed as well. This is because the action agent reduces discrepancies of varying degrees across trials, depending on the distance of the chosen item from the center of the screen. Thus, the discrepancy at the beginning of the action sequence was remembered and used to scale the expected lag (i.e., the input lag function within the learning agent was *expected lag times memory of discrepancy*), allowing expected lag to represent the belief in the time it takes to reduce a unit of discrepancy.¹¹

We operationalized this new model with seconds instead of trial as its time step so it could learn the lag. We also made the rate of action a positive, nonlinear function of (a) the degree of hunger or thirst, depending on the choice of food or drink, and (b) the distance to the target while pursuing it (see Appendix C). The first factor, which affected the gain of the action agent, made the vigor of the model's action (i.e., degree of output) a positive function of need. It is also consistent with an effect found by DeShon and Rench (2009). The second factor is a common element of psychomotor control theory models (Jagacinski & Flach, 2003; Powers, 1973). This created some variety in the rate of movement to the food or drink targets across and within trials that the expected lag learning agent did not account for.

A simulation of the full model of the DeShon and Rench (2009) paradigm showed several interesting findings. First, we found that the expected lag could be reasonably estimated and used for making opportunistic choices (i.e., targets were chosen that were relatively closer even though the states the targets improved were relatively more satiated at the time). Moreover, the model would pursue food and drink even when hunger and thirst states were beyond their goal levels. DeShon and Rench felt these two phenomena were beyond any model of self-regulation. Yet, the degree to which many of the DeShon and Rench participants exhibited the pursuit of food and drink beyond the goal level was substantially larger than our model exhibited using the parameter ranges described in Table 2. On the other hand, we could get to that level of pursuit if we set the expected disturbance weight to around 20, indicating a desire to buffer 20 trials of anticipated disturbances. Interestingly, the simulated individual under this condition reached the hunger and thirst limits built into the simulation/paradigm, which created some miscalibration in the expected disturbance belief. This finding is the result of the nonlinearity of the hunger and thirst states' functions, which our simple internal (i.e., mental) models could not handle. Meanwhile, one individual did not exhibit this pursuit-beyond-the-goal behavior, which we could easily produce by setting the expected disturbance weight or learning rate to zero (when initial expected disturbance was zero).

One of the DeShon and Rench participants also exhibited no opportunistic behavior. This we found more difficult to remove from our models behavior. Indeed, the only way to remove opportunistic behavior from the model was to set the initial expected lag to zero. However, this also resulted in the model never choosing food or drink. However, this issue could be resolved by lowering the resolution of the expectancy agents. That is, the expectancy agents processed information at 5 degrees of significance. A less fine-grained resolution more frequently equalized the expectancy component, thus reducing the opportunistic behavior. At the gross-est level of resolution the expectancy agent would output only zero (i.e., target cannot be reached) or one (target can be reached), which eliminated the opportunistic behavior. We discuss these findings and other issues in the next section.

Discussion

A sign of the maturity of a science is the degree to which it can formally account for the dynamics of the phenomena in its domain. For decades, organizational behaviorists have been trying to take time and dynamics more seriously, though with little success (Mitchell & James, 2001). Motivation and self-regulation, in particular, are dynamic phenomena that oblige dynamic theories (Dalal & Hulin, 2008; DeShon, 2012; Lord et al., 2010). Vancouver et al. (2010) began to fill the theoretical gap with a dynamic model of multiple-goal pursuit. However, the dynamics of that model were limited to the information and environment with which the

¹¹ In the earlier model, the discrepancy reduced was always one, making this element unnecessary.

individual interacted. Here, we added a formal description of how persons change, not just the signals passing through them, as they interact with the environment. Moreover, we integrated this description with the MGPM. Together, the theoretical elements represent mechanisms for individuals to develop their own internal representations of the environment that can be used to facilitate self-regulation. We tested the model's behavior with simulations, sensitivity analysis, and experimentation. We also assessed the model against the observations from two published studies. Below, we discuss what further work is needed to assess the current model, to extend the scope of phenomena such a model might address, and to apply the model.

Empirical extensions

As with informal theories, researchers need to test computational models with data (Busemeyer & Diederich, 2010; Davis et al., 2007; Vancouver & Weinhardt, 2012). Typically, this involves assessing the ability of the model to predict the behavior of the entities the model presumably represents, particularly relative to alternative models or explanations. In the case of dynamic models or theories, the predictions can involve trajectories of behavior or the association of variables over time. Relatively little of this type of data exists and we used what we could find (i.e., DeShon & Rench, 2009; Schmidt & DeShon, 2007), but additional data is needed to further assess the model. To motivate that data collection, we describe some protocols that could challenge the model.

Expected lag

One of the central elements of the MGPM not yet tested but further elaborated upon here is the dynamic expectancy and expected lag concepts. In particular, the present model represents change in expectancy beliefs arising somewhat as a function of changes in expected lag, which arise somewhat as a function of change in actual lag and bias perceptions of the passage of time. Although there is some evidence that expectancies play an important role in dynamic, multiple-goal pursuit (e.g., Louro et al., 2007; Schmidt & Dolis, 2009), can change over time, and can be miscalibrated (e.g., Vancouver, Thompson, & Williams, 2001), the notion that the miscalibration is a function of the misperception of the passage of time has not been examined. It would be interesting to develop a protocol whereby participants were in conditions known to affect the perception of the passage of time and see if that affected behavior that was dependent on the level of expected lag belief. For example, Gable and Poole (2012) found that affective experiences during approach motivation influenced perceptions of time. Specifically they found that individuals engaged in approach motivation, particularly during pleasant experiences, perceive time as passing more quickly. Furthermore, in a study on the relationship between visual and temporal perception, Vigo and Zeigler (submitted for publication) found that the structural complexity of a visual stimulus determines time estimates. If some work contexts tend to have these kinds of conditions, one might expect greater miscalibration effects for expected lags and expectancies. More generally, meta-analysis of time perceptions by Block and Zakay (1997) found such perceptions are typically underestimated. Our model implies, and research supports (Roy, Christenfeld, & McKenzie, 2005) the notion that this underestimation problem might be a primary source of the prevalence of overconfidence (Dunning et al., 2004) and the planning fallacy (Buehler, Griffin, & Ross, 1994).

Another interesting study might involve manipulating initial levels of expected lag using a paradigm that would be sensitive to the changing levels of expected lag over time. The model would predict that the initial manipulation might play a role early on, but it would soon cease to matter as the individual learned the lag from interacting with the task. It would also predict that if the

conditions that determined lag in a performance environment were different from the conditions in a training environment, misperceptions of expected lag might undermine training transfer (Goldstein & Ford, 2002).

Expected disturbances

The major conceptual contribution of the present model regards learning contextual contingencies (e.g., disturbances) and the notion that what is learned is used during subsequent goal striving. Thus, it would be important to determine if the nature of an individual's history with disturbances predicts task engagement beyond the goal level (i.e., buffering). That is, if undisturbed, the model would predict engaging in other tasks once the focal task was achieved. On the other hand, the model would predict that participants in a condition where disturbances move a variable away from a goal would engage in behavior beyond the goal. Perhaps most interestingly, the model also predicts that participants in a condition where disturbances move a variable toward a goal are more likely to direct resources elsewhere (i.e., choose other goals). These effects are likely to influence multitasking behavior. For example, if a salesperson is able to predict a coming slump in sales, he or she may aggressively reach out to past customers to buffer against projected future revenue loss. However, if customers are calling in on their own, one might focus on other tasks like learning about a new product.

We also modeled two types of disturbances. In the Schmidt and DeShon (2007) protocol, the disturbances were of one direction, but their occurrence was unpredictable (i.e., there was uncertainty). The learning mechanism we described was able to acquire a reasonable representation of this type of disturbance when the learning rate was low. On the other hand, in the DeShon and Rench (2009) protocol the disturbances were not only in one direction, but completely regular. Interestingly, our model had no trouble representing this disturbance until the variable hit a ceiling. It is not clear if individuals would also have trouble maintaining a belief about the disturbance, developed an understanding of the constraints of the value of the variable, or no longer update understandings about the lag when the limit is hit. These latter two possibilities have been the topic of cognitive models of learning (Anderson, 1995), which might be worth incorporating into the current model. For example, Anderson (1995) describes a simple model where the learning rate is divided by the number of learning trials, meaning that learning rate would drop over time. A question would be how generalizable such a model would be to the kind of learning contexts of interested to applied researchers (i.e., ones where the nature of the environment is constantly capable of changing). For example, Phillips, Neal, and Vancouver (2012) found a learning rate of 0.5 provided optimal fit with data for a model that represented learning performance capacity. Other contexts suggest learning rates of one (Wang & Arbib, 1993). Further explorations of learning rates for these kinds of models are sorely needed.

One might speculate that the differences in the types of disturbances also accounts for the apparent value the DeShon and Rench (2009) participants placed on the expected disturbance – a value reflected in the high expected disturbance weight needed to account for many of the participants' goal-pursuit behavior. For their part, DeShon and Rench (2009) were partial to Simon's (1956) simple rules model of goal-directed behavior as a way to explain the pursuing-beyond-goal behavior. According to Simon, having a rule to collect food and/or drink whenever available might be a simple rule one would use to navigate the environment. Such a rule might be reasonable if (a) states are frequently disturbed from ideal, (b) the opportunities for improving the states were scarce (e.g., food was scarce), and (c) the organism could store excess. Alternatively, the last two conditions might justify a high value for the expected

disturbance weight given the expected disturbances represents the first condition. The problem with this interpretation is that the second condition did not apply in the [DeShon and Rench \(2009\)](#) paradigm. Indeed, given the large expected disturbance weight needed to fit many of the participants' data, we are inclined to believe that what is more likely is that the lack of costs associated with pursuing food or drink accounts for the excessive goal-directed behavior. That is, the value of minima and maxima goals is that resources can be directed elsewhere when the goals are achieved. Moreover, there were no risks associated with goal pursuit (e.g., exposure to predators). Neither of these conditions held (e.g., there was nothing else for particular to do but to pursue food and drink). It would be interesting to examine these notions empirically.

Finally, another classic [Simon \(1955\)](#) paper gave us the clue regarding the level of resolution explanation to the lack of opportunistic behavior problem. Specifically, Simon suggested that individuals apply simple all-or-none payoff rules when making choices. The modeling approach here suggests that this might be the case if information (i.e., stimuli) is fuzzy or it takes work to obtain precise perceptions used in downstream processes. Of course, such speculation requires empirical confirmation.

Theoretical extensions

When [Vancouver et al. \(2010\)](#) introduced the MGPM, they noted that it integrated motivational theories of goal striving and goal choice. Here we added a learning component, providing a link to cognitive computational models of learning and more formally representing processes described in comprehensive theories of human behavior (e.g., social cognitive theory; [Bandura, 1986](#)). Regarding the link to computational cognitive models, we have provided a bridge between levels of analysis typically not crossed (cf., [Gibson et al., 1997](#)). Regarding the representation of processes in comprehensive theory, we have explicitly added a way to represent changes to the person. Below, we describe how to develop these important contributions further.

Link to cognitive computational models

In recent years, computational modelers have come to believe that computational approaches that cross levels of analysis will be important for moving theorizing forward ([Anderson et al., 2004](#); [Boden, 2008](#)). Focusing on models of cognitive processes, [Sun, Coward, and Zenzen \(2005\)](#) described four levels of analysis (i.e., social/cultural; psychological; componential; and physiological). Social/cultural models include models of collective behavior and models of how individual humans interact with their physical and sociocultural environments. Psychological models include models of behaviors, beliefs, emotion, and knowledge. [Vancouver \(2005\)](#) referred to psychological models as system-level theories where humans are the system of interest. Most verbal theories in psychology are at this level. Component or subsystem theories describe how components or subsystems within individuals interact with one another to produce observable psychological phenomena. Finally, physiological theories describe the biological substrate of the process, or what [Marr \(1982\)](#) called the implementation level.

Typically, cognitive models of learning are either component theories or blend component and physiological elements. For example, the delta-learning rule described here is found in animal learning theories ([Rescorla & Wagner, 1972](#)), which inspired, in combination with neurology (i.e., the physiological level), component connectionists models ([Thomas & McClelland, 2008](#)). The delta-learning rule has also been incorporated in symbolic processing models (e.g., ACT-R, [Fu & Anderson, 2006](#)) that are component models, which are developing links to the physiological substrates of behavior ([Anderson et al., 2004](#)). Meanwhile, perceptual control

theory (PCT; [Powers, 1973](#)), which is the basis of the MGPM, is a cybernetic/systems theory ([Vancouver, 2000](#)) that crosses social/cultural, psychological, and componential levels. Specifically, PCT includes the social/cultural level by highlighting the influence of the physical and social environment as the feedback loop passes through it. Indeed, the primary functional role of control systems is to dampen or counteract disturbances to variables of interest – disturbances coming from external sources. Moreover, computational models of control systems can include environmental constructs that can affect the actions and stimuli coming from and going into the human, respectively (e.g., [Vancouver et al., 2010](#)). Meanwhile, psychological constructs like goals and perceptions are key to understanding human control process ([Vancouver et al., 2005](#)). However, the principle level of analysis for control theory is the component level. In particular, PCT describes the possible organization of multiple control sub-systems (i.e., agents) within human systems as well as how this organization can account for various human behavioral phenomena like learning and motivation. However, it also describes learning as a function of system-level error and a person change process that is random. In the model presented here, the error is specific to a subsystem and the change represents a gradient descent toward a minimum error ([Anderson, 1995](#)). Meanwhile, until recently cognitive models have tended to ignore feedback processes that extend beyond the human system, focusing more on feedback among the components within the human system.

Linking to comprehensive theories of human behavior

One of the advantages of linking to cognitive computational models of learning is that it facilitates the further integration of cognitive models into more formal representations of comprehensive theories of behavior. That is, in addition to adding learning to a model of goal striving and goal choice, we need to add elements to capture more fully the complexity of human behavior, particularly in complex, dynamic environments like those found in work settings. For example, the modeling scenarios thus far examined involved one-to-one mapping of actions to goals. That is, one and only one behavior or choice positively affects one and only one of the goals sought. In the first instantiation, the expectancies regarding achieving each goal were relevant and used by the choice agent. In the second instantiation, the expectancies regarding each action were relevant and used. However, in many cases, multiple actions might have varying effects, both in terms of sign and degree, on numerous variables relevant to multiple goals. To understand how these conditionals might be learned and used when self-regulating will require more elaborate modeling and empirical scenarios. We see this as an issue of scaling up the current models, which will likely be facilitated by more fully integrating cognitive models of learning and decision making (e.g., [Busemeyer & Johnson, 2008](#); [Vigo and Doan \(submitted for publication\)](#)), but it is not clear how straightforward this is likely to be.

Indeed, issues like the role of attention and affect during self-regulation are likely to be particularly complex. Fortunately, some issues (e.g., affect) have been approached using more informal theorizing within the self-regulation literature. For example, [Carver and Scheier \(1990\)](#) have speculated that affect arises not from discrepancies per se, but from deviations between predictions of the velocity of discrepancy change and actual velocity. Interestingly, these deviations are what the discrepancies in the expected lag learning agent represents. This is why we suggested emotion might also be an output from this agent (see [Fig. 2](#)). A question is then how this signal subsequently influences processes within the individual. If emotion is information ([Schwarz, 1990](#)), then we should be able to model how that information is used. Alternatively, most theories of job satisfaction suggest that the attitude,

which has an affective component, is a direct function of discrepancies from values (e.g., Locke, 1976), needs (e.g., Porter, 1962), or other standards that might be modified based on experience (e.g., Hulin, 1991; Thibault & Kelley, 1959). These theories were the basis of a simple model of job satisfaction and well-being that Vancouver and Weinhardt (2012) developed as an example of computational modeling. In that model, the discrepancies themselves, not the difference in anticipated change and observed change, were the source of the affective experience. Finally, the more dynamic affective events theory (Weiss & Cropanzano, 1996) ties affect to events that might frustrate or facilitate goal achievement. However, is it not clear whether the affect arises because the events create changes in discrepancies or changes in the rate of discrepancy reduction that deviate from expected rate (or some other standard). Working out these distinctions computationally and incorporating them into the present model will likely facilitate theory testing given the dynamics and other complexities (i.e., nonlinearities) involved.

Besides elaborating on the current model by incorporating learning, the modeling approach represented here is likely to facilitate theory integration and development more generally. Weinhardt and Vancouver (2012) provide several examples where computational modeling might facilitate understanding in various areas of organization psychology. Here, we focus on examples where the current model of multiple dynamics (i.e., changing person, environment, and behavior) may be especially relevant. In particular, the relatively recent focus on dynamic performance (Sonnetag & Frese, 2012) will require dynamic theories of behavior. Interestingly, in the past, these theories have tended to be either changing-subject or changing-task models (Keil & Cortina, 2001). Given that the model presented here describes changes in both the person and the environment, we agree with Keil and Cortina that these earlier models are likely complementary. Meanwhile, other dynamic processes are likely to be cyclical, which presents certain challenges conceptually (Sonnetag & Frese, 2012), empirically (Hanges & Wang, 2012), and statistically (DeShon, 2012). Computational models are likely to be especially useful conceptually because human's reasoning about dynamic variables and problems is often incorrect (Cronin, Gonzalez, & Serman, 2009) and there could be emergent properties within the system that are difficult to account for without simulations of computational theories (Kozlowski, Chao, Grand, Braun, & Kuljanin, 2013).

Indeed, the self-regulation literature includes an example where complex dynamics are likely undermining comprehensive integrative efforts, for which the modeling presented here might provide some reconciliation. Specifically, Bandura, 2012; Bandura & Locke, 2003) and Vancouver (2005, 2012) are debating the role of self-efficacy beliefs within models of human behavior. For Bandura self-efficacy is the central construct determining all motivated behavior. For Vancouver self-efficacy can play a positive, negative, or null role in determining motivated behavior, but often appears to play a positive role because it incorporates knowledge of capabilities obtained from observing past performance and via other spurious effects. Sitzmann and Yeo (2013) addressed this controversy in a recent meta-analysis of longitudinal studies with repeated measures of self-efficacy and performance. The findings indicated that positive correlations between self-efficacy and performance at the between-person level are much more likely to be because of past performances' influence on self-efficacy than self-efficacy's influence on performance. Specifically, they found that within-person correlations of self-efficacy lagged behind performance were often positive (i.e., implying self-efficacy positively influences performance) until trend or past performance was controlled. Once these factors were controlled, the effect of self-efficacy on performance was null, though they found evidence of

moderators that might make self-efficacy's influence positive or negative as well.

Although not addressing all the issues regarding self-efficacy (e.g., the negative effects for self-efficacy are presumed to arise when performance feedback is ambiguous or absent [Schmidt & DeShon, 2010; Vancouver, 2005]), which is not a context represented in our models), our models can shed a great deal of light on Sitzmann and Yeo's (2013) findings. Specifically, as noted above, expected lag is a central belief influencing expectancies and one's self-efficacy for task accomplishment. We would also argue that in the first model, expected lag is the essential component of one's self-efficacy belief regarding scheduling – the subtask to the task of providing schedules for students in the two schools. Moreover, scheduling performance would likely be measured in terms of number of schedules correctly completed, which would be a function of the actual lag in completing schedules. In our model, actual lag was unaffected by motivational processes within the model, and in particular, it was unaffected by expected lag or expectancies. On the other hand, expected lag was a positive function of actual lag, changing as actual lag changed. Thus, our model is consistent with the meta-analytic findings.¹²

Yet, the apparent null effect for self-efficacy (or expected lag) on scheduling performance described above tells, as Bandura (1997) puts it, only half the story. One of the main reasons for the modeling effort engaged in here was to reveal how expectancy beliefs like self-efficacy arise. If these beliefs were of no relevance for motivation or human behavior, it seems that we have been wasting the readers' time. Yet, these beliefs, our model argues, are used to make decisions regarding goal pursuit. In particular, at the task level (i.e., above the scheduling level), beliefs about the inability to reach the goals (i.e., zero students needing schedules in both schools) leads to the abandonment of one of those goals, undermining performance on that task. This non-spurious, positive effect represents a critical role for self-efficacy beliefs relating to the persistence of behavior (Pinder, 2008). Moreover, recent research (e.g., Vancouver, More, & Yoder, 2008) implicates expectancy and self-efficacy beliefs as positive influences on accepting a goal in the first place. In this way, Bandura's (1997) foundational role for self-efficacy may still have merit and reconcile with the control theory perspective used here. Moreover, our models provide clues regarding the moderators of the self-efficacy-to-performance relationship (e.g., the level of performance considered; the choice context).

Finally, one can use the modeling approach taken here to understand meso-level phenomena like team processes. We mentioned above that the computational approach we have taken here includes the social/culture level (Sun et al., 2005) by explicitly modeling the individual's environment. However, the social/culture level also includes models of multiple individuals interacting with each other. This would be a useful extension. In particular, DeShon, Kozlowski, Schmidt, Milner, and Wiechmann (2004) have a theory of team performance based on a conceptual structure similar to ours. In particular, DeShon and colleagues sought to understand

¹² Importantly, our simulations reveal that correlating expected lag at time $t-1$ with actual lag at time t would produce a negative coefficient if examined while expected lag was approaching actual lag from a position of overconfidence (i.e., expected lag was lower than actual lag). This can be seen in the first two-thirds of Fig. 3, where the sign of the trend lines differ between expected and actual lag. On the other hand, this same analysis will give a positive coefficient after expected lag has caught up to actual lag if actual lag is changing because the sign of the trend lines for both expected and actual lag will be the same (see last third of Fig. 3, where both lines are slightly decreasing). We suspect that this latter situation is more common given the protocols used in the primary studies examined by Sitzmann and Yeo (2013). Indeed, this is likely why controlling for trend or past performance generally turned otherwise positive coefficients representing the lagged effect of self-efficacy on performance from positive to null.

how multiple-goals and feedback influence self-regulatory processes at the individual and team level. However, no one has represented the theory computationally, nor does it address mental model formation, which research has shown to be important for team functioning (Bell, Kozlowski, & Blawath, 2012). Yet, the complexities arising from multiple-goal pursuit at the team level are likely to be substantially more difficult to predict as compared to the single individual.

Practical applications

Beyond model testing and extending, there is the issue of application. For example, the description of the development process of expected lag beliefs can inform interventions designed to improve planning (e.g., reduce the planning fallacy). Above we talked about factors that might bias perceptions of the passage of time and how those might result in planning errors, but another aspect of the model is the demarcation used to determine learning events and how that demarcation might interact with time perception bias and disturbance expectations. Specifically, individuals using more abstract or higher-level events (e.g., the time from inception to acceptance of a paper) are more likely to misperceive lags and subtract away (i.e., overly discount) disturbances or times engaged in off-task activities. On the other hand, those who use more concrete or lower-level events (i.e., subtasks like the analysis of the data from a study) are likely to not recognize the degree to which off-task events interfere with the ability to combine the subtasks into a completed project (Bishop & Trout, 2008). Understanding how individuals demarcate tasks and understand the disturbances to them, both directly and via off-task conflicting goals, should improve attempts to mitigate the problem (Lovallo & Kahneman, 2003).

The issue of event demarcation is also relevant during goal striving, not just during planning. This has been most clearly observed in team research, where the halfway point seems a common demarcation point (Gersick, 1988). In particular, teams are more likely to change their approach to the team task if they believe their current pace is insufficient for reaching the goal by the deadline (Gersick, 1989). The halfway point appears to represent a learning event end that is used to determine the team's pace (i.e., lag) and assess that pace against the time left, as described in the MGPM. Yet, leaders or organizations might want to be more active in determining when teams (or individuals) develop their expected lags and use them in planning or restructuring their approach in the time left.

More generally, adding learning components to models of organizational behavior should help facilitate the development of accurate mental models of the person, team, and context (i.e., capacity and context beliefs; Ford, 1992) by employees, teams, and leaders. Besides improving task and team performance, accurate mental models might mitigate perceptions of injustice. For example, Gilliland and Anderson (2011) note that several elements of injustice perception are based on the acquisition of scarce resources. Specifically, the acquisition of resources beyond need can lead to perceptions of greed (Gilliland & Anderson, 2011). The current model provides an explanation regarding why such acquisition might be motivated (e.g., uncertain disturbances). Alternatively, extensions of the model might support interventions that improve the ability of others to see why someone appears to be acquiring beyond need.

Conclusion

Comprehensive theories of human behavior should integrate explanations of action, thinking, learning, and emotion (Vancouver, 2008). Presumably, the field would prefer a formal

(i.e., mathematical) and parsimonious theory (Forster, 2000), and specific emphasis on dynamics is likely an important future direction (Kozlowski, 2012). In this paper, we described a computational model based on a simple structural concept to represent the processes (i.e., dynamics) of action, thinking (i.e., decision making), and learning. The parsimony arises from the repetition of the single architecture found in the simple structure. Yet, using this structure, we describe how individuals might choose among goals and actions based on current states of key variables as well as the understandings they construct to predict future states or estimate current states. Via this approach, we seek to move the field beyond static models and verbal theories by formally integrating the promise of computational cognitive modeling, the dynamics of motivation, and the conceptualization of the interaction of person, behavior, and environment.

Acknowledgement

This work reported here was partially supported by National Science Foundation Grant SES-0851764.

Appendix A

average disturbance = 1.03
 bias = 1
 change in rate = 4.8e-006
 completed schedule = 1
 deadline = 1800
 disturb comparator = task states – predicted task states
 disturb output = disturb comparator * learning gain
 emotion = –lag output
 expectancy A comparator = MAX (time output – expectancy A input, 0)
 expectancy A input = expected lag * task A comparator
 expectancy A output = expectancy A comparator
 expectancy B comparator = MAX (time output – expectancy B input, 0)
 expectancy B input = expected lag * task B comparator
 expectancy B output = expectancy B comparator
 expected disturbance weight = 1
 expected disturbances = INTEG(disturb output * learning rate, initial expected disturbances)
 expected lag = INTEG(lag output * learning rate, initial expected lag)
 expected lag in minutes = expected lag/60
 incentive sensitivity = 1
 initial expected disturbances = 0
 initial expected lag = 30
 initial line lengths = 5
 initial rate = 0.0212
 lag comparator = perception of lag – lag input
 lag output = lag comparator * learning gain
 learning gain = IF THEN ELSE (schedule comparator <= 0, 1, 0)
 learning rate = 0.04
 memory of task states = INTEG(IF THEN ELSE (schedule input = 0, 1, 0) * (task states – memory of task states), initial line lengths * 2)
 perception of lag = INTEG(TIME STEP * bias – (perception of lag – 1) * IF THEN ELSE (schedule input = 0, 1, 0), 0)
 lag input = expected lag
 predicted task states = memory of task states + expected disturbances
 schedule comparator = IF THEN ELSE (completed

(continued on next page)

schedule – schedule input > 0, 1, 0)
 schedule done = IF THEN ELSE (Schedule status <= 1, 0, 1)
 schedule input = Schedule status
 schedule output = schedule comparator
 Schedule status = INTEG(scheduling rate * schedule
 output – Schedule status * schedule done, 0)
 scheduling rate = INTEG(change in rate, initial rate)
 Task = task choice output
 task A comparator = MAX (task A input – task A goal, 0)
 task A disturbance = PULSE (170, 1) + PULSE (240,
 1) * 3 + PULSE (510, 1) + PULSE (540, 1) * 2 + PULSE (810,
 1) * 2 + PULSE (900, 1) * 2 + PULSE (1080, 1) * 2 + PULSE
 (1110, 1) + PULSE (1320, 1) * 2 + PULSE (1500,
 1) * 2 + PULSE (1595, 1)
 task A gain = 1 + incentive sensitivity * task A incentive
 task A goal = 0
 task A incentive = 1
 task A input = Task A state + expected disturbance
 weight * expected disturbances
 task A output = task A comparator * task A gain
 Task A state = INTEG(–1 * schedule done * Task + task A
 disturbance, initial line lengths)
 task B comparator = MAX (task B input – task B goal, 0)
 task B disturbance = PULSE (60, 1) * 2 + PULSE (90, 1) + PULSE
 (330, 1) + PULSE (450, 1) + PULSE (630, 1) + PULSE (660,
 1) * 3 + PULSE (840, 1) + PULSE (960, 1) * 3 + PULSE (1260,
 1) + PULSE (1350, 1) * 2 + PULSE (1440, 1) * 2 + PULSE
 (1623, 1)
 task B gain = 1 + incentive sensitivity * task B incentive
 task B goal = 0
 task B incentive = 1
 task B input = Task B state + expected disturbance
 weight * expected disturbances
 task B output = task B comparator * task B gain
 Task B state = INTEG(–1 * schedule done * (1 – Task) + task B
 disturbance, initial line lengths)
 task choice comparator = IF THEN ELSE (task choice
 input – expectancy B output * task B output >= 0, 1, 0)
 task choice input = expectancy A output * task A output
 task choice output = task choice comparator
 task states = task A input + task B input
 time comparator = MAX (deadline – time input, 0)
 time gain = 1
 time input = Time
 time output = time gain * time comparator

Appendix B

choice agent = (hunger agent * hunger expectancy
 agent – thirst agent * thirst expectancy agent)
 distance to drink = RANDOM UNIFORM (0.1, 0.9, 110)
 distance to food = RANDOM UNIFORM (0.1, 0.9, 112)
 disturbances = 1
 expected lag = 10
 goals = 75
 hunger = INTEG(obtain food – disturbances – MAX
 (hunger – 100, 0), 50)
 hunger agent = MAX (goals – hunger, 0)
 hunger expectancy agent = MAX (resource limit – distance to
 food * expected lag, 0)
 obtain drink = IF THEN ELSE (choice agent < 0, 5, 0)
 obtain food = IF THEN ELSE (choice agent > 0, 5, 0)

resource limit = 10
 thirst = INTEG(obtain drink – disturbances – MAX
 (thirst – 100, 0), 50)
 thirst agent = MAX (goals – thirst, 0)
 thirst expectancy agent = MAX (resource limit – distance to
 drink * expected lag, 0)

Appendix C

action agent = IF THEN ELSE (distance to target – target
 goal > 0, (1 – 1/EXP (distance to target * IF THEN ELSE
 (choice agent > 0, hunger agent, IF THEN ELSE (choice
 agent < 0, thirst agent, 0)))) * 0.3 + 0.07, 0)
 bias = 1
 choice agent = INTEG((hunger agent * hunger expectancy
 agent – thirst agent * thirst expectancy agent) * trial
 begin – choice agent * reached target, 0)
 delayed dist to target = DELAY FIXED (distance to target, 1, 0)
 distance to drink = INTEG((RANDOM UNIFORM (0.1, 0.8,
 110) – distance to drink) * reached target, RANDOM
 UNIFORM (0.1, 0.8, 110))
 distance to food = INTEG((RANDOM UNIFORM (0.1, 0.8,
 112) – distance to food) * reached target, RANDOM
 UNIFORM (0.1, 0.8, 112))
 distance to target = INTEG((IF THEN ELSE (choice agent > 0,
 distance to food, IF THEN ELSE (choice agent < 0, distance to
 drink, 0)) – distance to target) * make choice – action
 agent – IF THEN ELSE (distance to target < 0, distance to
 target, 0), 0)
 disturbances = 0.5
 expected disturbance weight = 1
 expected disturbances = INTEG(learning agent for
 disturbances * learning rate, initial expected disturbances)
 expected lag = INTEG(learning rate * learning agent for lag,
 initial expected lag)
 goal = 75
 hunger = INTEG(IF THEN ELSE (choice agent > 0, reached
 target * 20, 0) – disturbances – MAX (hunger – 100, 0), 50)
 hunger agent = MAX (goal – (hunger + expected disturbance
 weight * expected disturbances), 0)
 hunger expectancy agent = MAX (resource limit – distance to
 food * expected lag, 0)
 initial expected disturbances = 0
 initial expected lag = 1
 internal states = hunger + thirst
 learning agent for disturbances = (internal states – (memory
 of states + expected disturbances)) * make choice
 learning agent for lag = (perception of lag – expected
 lag * memory of discrepancy) * reached target
 learning rate = 0.04
 make choice = DELAY FIXED (trial begin, 1, 0)
 memory of discrepancy = INTEG(distance to target * record
 distance – memory of discrepancy * trial begin, 0)
 memory of states = INTEG((internal states – memory of
 states) * make choice + reached target * 5, 100)
 perception of lag = INTEG(bias * TIME STEP – perception of
 lag * record distance, 0)
 reached target = IF THEN ELSE (distance to target <= 0:
 AND::NOT: delayed dist to target <= 0, 1, 0)
 record distance = DELAY FIXED (make choice, 1, 0)
 resource limit = 10
 target goal = 0

thirst = INTEG(IF THEN ELSE (choice agent < 0, reached target * 20, 0) – disturbances – MAX (thirst – 100, 0), 50)
 thirst agent = MAX (goal – (thirst + expected disturbance weight * expected disturbances), 0)
 thirst expectancy agent = MAX (resource limit – distance to drink * expected lag, 0)
 trial begin = IF THEN ELSE (MODULO (Time, 10) = 0, 1, 0)

References

- Adner, R., Polos, L., Ryall, M., & Sorenson, O. (2006). Call for papers: Academy of management review special topic forum: Formal approaches to management theory. *The Academy of Management Review*, 31, 775–776.
- Anderson, J. A. (1995). *An introduction to neural networks*. Cambridge, MA: MIT Press.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4), 1036–1060.
- Ashby, W. R. (1958). Requisite variety and its implications for the control of complex systems. *Cybernetica*, 1(2), 83–99.
- Atkinson, J. W., & Birch, D. (1970). *The dynamics of action*. New York, NY: Wiley.
- Austin, J. T., & Vancouver, J. B. (1996). Goal constructs in psychology: Structure, process, and content. *Psychological Bulletin*, 120(3), 338–375.
- Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84, 191–215.
- Bandura, A. (1986). *Social foundations of thought and action*. Englewood Cliffs, NJ: Prentice-Hall.
- Bandura, A. (1997). *Self-efficacy: The exercise of control*. New York: Freeman.
- Bandura, A. (2012). On the functional properties of perceived self-efficacy revisited. *Journal of Management*, 38, 9–44.
- Bandura, A., & Locke, E. A. (2003). Negative self-efficacy and goal effects revisited. *Journal of Applied Psychology*, 88(1), 87–99.
- Bell, B. S., Kozlowski, S. W. J., & Blawath, S. (2012). Team learning: A review and integration. In S. W. J. Kozlowski (Ed.), *The Oxford handbook of industrial and organizational psychology*. New York: Oxford University Press.
- Bishop, M. A., & Trout, J. D. (2008). Strategic reliabilism: A naturalistic approach to epistemology. *Philosophy Compass*, 3(5), 1049–1056.
- Block, R. A., & Zakay, D. (1997). Prospective and retrospective duration judgments: A meta-analytic review. *Psychonomic Bulletin & Review*, 4, 184–197.
- Boden, M. A. (2008). An evaluation of computational modeling in cognitive science. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 667–683). New York, NY: Cambridge University Press.
- Boekaerts, M., Maes, S., & Karoly, P. (2005). Self-regulation across domains of applied psychology: Is there an emerging consensus. *Applied Psychology: An International Review*, 54, 149–154.
- Buehler, R., Griffin, D., & Ross, M. (1994). Exploring the “planning fallacy”: Why people underestimate their task completion times. *Journal of Personality and Social Psychology*, 67(3), 366–381.
- Bussemeyer, J. R., & Diederich, A. (2010). *Cognitive modeling*. Los Angeles, CA: Sage.
- Bussemeyer, J. R., & Johnson, J. G. (2008). Microprocess models of decision making. In R. Sun (Ed.), *Cambridge Handbook of Computational Psychology* (pp. 302–321). New York: Cambridge University Press.
- Cannon, W. B. (1932). *The wisdom of the body*. New York: W.W. Norton.
- Carver, C. S., & Scheier, M. F. (1981). *Attention and self-regulation: A control theory approach to human behavior*. New York: Springer-Verlag.
- 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.
- Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. New York, NY: Cambridge University Press.
- Cronin, M. A., Gonzalez, C., & Sterman, J. D. (2009). Why don't well-educated adults understand accumulation? A challenge to researchers, educators, and citizens. *Organizational Behavior and Human Decision Processes*, 108(1), 116–130.
- Cropanzano, R., James, K., & Citera, M. A. (1993). A goal hierarchy model of personality, motivation, and leadership. In L. L. Cummings & B. M. Staw (Eds.), *Research in organizational behavior* (Vol. 15, pp. 267–322). Greenwich, CT: JAI Press.
- Dalal, R. S., & Hulin, C. L. (2008). Motivation for what? A multivariate dynamic perspective of the criterion. In R. Kanfer, G. Chen, & R. D. Pritchard (Eds.), *Work motivation: Past, present, and future* (pp. 63–101). New York, NY: Routledge.
- Davis, J. P., Eisenhardt, K. M., & Bingham, C. B. (2007). Developing theory through simulation methods. *The Academy of Management Review*, 32(2), 480–499.
- DeShon, R. P., & Gillespie, J. Z. (2005). A motivated action theory account of goal orientation. *Journal of Applied Psychology*, 90, 1096–1127.
- DeShon, R. P. (2012). Multivariate dynamics in organizational science. In S. W. J. Kozlowski (Ed.), *The Oxford handbook of industrial and organizational psychology*. New York: Oxford University Press.
- DeShon, R. P., Kozlowski, S. W. J., Schmidt, A. M., Milner, K. A., & Wiechmann, D. (2004). A multiple-goal, multilevel model of feedback effects on the regulation of individual and team performance. *Journal of Applied Psychology*, 89, 1035–1056.
- DeShon, R. P., & Rench, T. A. (2009). Clarifying the notion of self-regulation in organizational behavior. In G. P. Hodgkinson & J. K. Ford (Eds.), *International review of industrial and organizational psychology* (Vol. 24, pp. 217–248). Wiley-Blackwell.
- Diefendorff, J. M., & Chandler, M. M. (2011). Motivating employees. In S. Zedeck (Ed.), *APA handbook of industrial and organizational psychology* (Vol. 3, pp. 65–135). Washington, DC: American Psychological Association.
- Dunning, D., Heath, C., & Suls, J. M. (2004). Flawed self-assessment: Implications for health, education, and the workplace. *Psychological Science in the Public Interest*, 5(3), 69–106.
- Edwards, J. R. (1992). A cybernetic theory of stress, coping, and well-being in organizations. *Academy of Management Review*, 17, 238–274.
- Fagot, J., Kruschke, J. K., Depy, D., & Vauclair, J. (1998). Associative learning in humans (*Homo sapiens*) and baboons (*Papio papio*): Species differences in learned attention to visual features. *Animal Cognition*, 1, 123–133.
- Farrell, S., & Lewandowsky, S. (2010). Computational models as aids to better reasoning in psychology. *Current Directions in Psychological Science*, 19, 329–335.
- Ford, M. E. (1992). *Motivating humans: Goals, emotions, and personal agency beliefs*. Thousand Oaks, CA: Sage Publications.
- Forster, M. R. (2000). Key concepts in model selection: Performance and generalizability. *Journal of Mathematical Psychology*, 44, 205–231.
- Fu, W.-T., & Anderson, J. R. (2006). From recurrent choice to skill learning: A reinforcement-learning model. *Journal of Experimental Psychology: General*, 135(2), 184–206.
- Gable, P. A., & Poole, B. D. (2012). Time flies when you're having approach-motivated fun: Effects of motivational intensity on time perception. *Psychological Science*, 23(8), 879–886.
- Gersick, C. J. (1988). Time and transition in work teams: Toward a new model of group development. *Academy of Management Journal*, 31(1), 9–41.
- Gersick, C. J. (1989). Marking time: Predictable transitions in task groups. *Academy of Management Journal*, 32(2), 274–309.
- Gibson, F. P., Fichman, M., & Plaut, D. C. (1997). Learning in dynamic decision tasks: Computational model and empirical evidence. *Organizational Behavior and Human Decision Processes*, 71, 1–35.
- Gilliland, S. W., & Anderson, J. (2011). Perceptions of greed: A distributive justice model. In S. W. Gilliland, D. D. Steiner, & D. P. Skarlicki (Eds.), *Emerging perspectives on organizational justice and ethics* (pp. 137–166). Charlotte, NC: Information Age Publishing.
- Gluck, M. A., & Bower, G. H. (1988). Evaluating an adaptive network model of human learning. *Journal of Memory & Language*, 27, 166–195.
- Goldstein, I. L., & Ford, J. K. (2002). *Training in organizations: Needs assessment, development, and evaluation* (4th ed.). Belmont, CA: Wadsworth.
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27(3), 377–396.
- Hanges, P. J., & Wang, M. (2012). Seeking the holy grail in organizational psychology: Establishing causality through research designs. In S. W. J. Kozlowski (Ed.), *The Oxford handbook of industrial and organizational psychology*. New York: Oxford University Press.
- Hulin, C. L. (1991). Adaptation, persistence, and commitment in organizations. In M. D. Dunnette (Ed.), *Handbook of industrial and organizational psychology* (Vol. 2, 2nd ed., pp. 445–505). Palo Alto, CA: Consulting Psychologists Press.
- Hutter, M., & Legg, S. (2008). Temporal difference updating without a learning rate. *Advances in Neural Information Processing Systems*, 20, 705–712.
- Hyland, M. E. (1988). Motivational control theory: An integrative framework. *Journal of Personality and Social Psychology*, 55, 642–651.
- Ioannou, P. A., & Sun, J. (1996). *Robust adaptive control*. Upper Saddle River, NJ: Prentice Hall.
- Jagacinski, R. J., & Flach, J. M. (2003). *Control theory for humans: Quantitative approaches to modeling performance*. Mahwah, NJ: Lawrence Erlbaum.
- Kanfer, R. (2012). Work motivation: Theory, practice, and future directions. In S. W. J. Kozlowski (Ed.), *The Oxford handbook of industrial and organizational psychology*. New York: Oxford University Press.
- Keil, C. T., & Cortina, J. M. (2001). Degradation of validity over time: A test and extension of Ackerman's model. *Psychological Bulletin*, 127(5), 673–697.
- Keith, N., & Frese, M. (2005). Self-regulation in error management training: Emotion control and metacognition as mediators of performance effects. *Journal of Applied Psychology*, 90, 677–691.
- Klein, H. J. (1989). An integrated control theory model of work motivation. *Academy of Management Review*, 14, 150–172.
- Klein, H. J., Austin, J. T., & Cooper, J. T. (2008). Goal choice and decision processes. In R. Kanfer, G. Chen, & R. D. Pritchard (Eds.), *Work motivation: Past, present, and future* (pp. 101–151). New York, NY: Routledge.
- Kozlowski, S. W. J., Chao, G. T., Grand, J. A., Braun, M. T., & Kuljanin, G. (2013). Advancing multilevel research design: Capturing the dynamics of emergence. *Organizational Research Methods*, 16, 581–615.
- Kozlowski, S. W. J. (2012). The nature of industrial and organizational psychology. In S. W. J. Kozlowski (Ed.), *The Oxford handbook of industrial and organizational psychology*. New York: Oxford University Press.
- Krantz, D. H., & Kunreuther, H. C. (2007). Goals and plans in decision making. *Judgment and Decision Making*, 2, 137–168.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22–44.
- Lewandowsky, S., & Farrell, S. (2011). *Computational modeling in cognition: Principles and practices*. Los Angeles, CA: Sage.
- Lewin, K. (1951). *Field theory in social science: Selected theoretical papers*. Oxford, UK: Harpers.
- Locke, E. A. (1976). The nature and causes of job satisfaction. In M. D. Dunnette (Ed.), *Handbook of industrial and organizational psychology* (pp. 1297–1349). Chicago, IL: Rand McNally.

- Locke, E. A., & Latham, G. P. (2004). What should we do about motivation theory? Six recommendations for the twenty-first century. *The Academy of Management Review*, 29(3), 388–403.
- Lord, R. G., Diefendorff, J. M., Schmidt, A. M., & Hall, R. J. (2010). Self-regulation at work. *Annual Review of Psychology*, 61, 548–568.
- Lord, R. G., & Levey, P. E. (1994). Moving from cognition to action: A control theory perspective. *Applied Psychology: An International Review*, 43(3), 335–367.
- Louro, M. J., Pieters, R., & Zeelenberg, M. (2007). Dynamics of multiple-goal pursuit. *Journal of Personality and Social Psychology*, 93, 174–193.
- Lovall, D., & Kahneman, D. (2003). Delusions of success: How optimism undermines executives' decisions. *Harvard Business Review*, 81, 56–63.
- March, J. G. (1996). Learning to be risk averse. *Psychological Review*, 103, 309–319.
- Marr, D. (1982). *Vision: A computational approach*. San Francisco: Freeman & Co.
- Miller, J. G. (1978). *Living systems*. New York: McGraw-Hill.
- Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York: Holt.
- Mitchell, T., Hopper, H., Daniels, D., George-Falvy, J., & James, L. (1994). Predicting self-efficacy and performance during skill acquisition. *Journal of Applied Psychology*, 79, 506–517.
- Mitchell, T. R., & James, L. R. (2001). Building better theory: Time and the specification of when things happen. *Academy of Management Review*, 26, 530–547.
- Pinder, C. C. (2008). *Work motivation in organizational behavior* (2nd ed.). Upper Saddle River, NJ: Psychology Press.
- Porter, L. W. (1962). Job attitudes in management: I. Perceived deficiencies in need fulfillment as a function of job level. *Journal of Applied Psychology*, 46, 375–384.
- Porter, L. W., & Lawler, E. E. (1968). *Managerial attitudes and performance*. Homewood, IL: Irwin.
- Powers, W. T. (1973). *Behavior: The control of perception*. Oxford, England: Aldine.
- Powers, W. T. (1978). Quantitative analysis of purposive systems: Some spade work at the foundations of scientific psychology. *Psychological Review*, 85(5), 417–435.
- Powers, W. T., Clark, R. K., & McFarland, R. I. (1960). A general feedback theory of human behavior: Part 1. *Perceptual and Motor Skills*, 11, 71–88.
- Ratcliff, R. (1990). Constraints imposed by learning and forgetting functions. *Psychological Review*, 97, 285–308.
- Read, S. J., Monroe, B. M., Brownstein, A. L., Yang, Y., Chopra, G., & Miller, L. C. (2010). A neural network model of the structure and dynamics of human personality. *Psychological Review*, 117, 61–92.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Current research and theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.
- Roy, M. M., Christenfeld, N. J. S., & McKenzie, C. R. M. (2005). Understanding the duration of future events: Memory incorrectly used or memory bias? *Psychological Bulletin*, 131, 738–756.
- Scherbaum, C. A., & Vancouver, J. B. (2010). If we produce discrepancies, then how? Testing a computational process model of positive goal revision. *Journal of Applied Social Psychology*, 40, 2201–2231.
- Schmajuk, N. A., & DiCarlo, J. J. (1991). A neural network approach to hippocampal function in classical conditioning. *Behavioral Neuroscience*, 105, 82–110.
- 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.
- Schmidt, A. M., & DeShon, R. P. (2010). The moderating effects of performance ambiguity on the relationship between self-efficacy and performance. *Journal of Applied Psychology*, 95, 572–581.
- 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, 678–691.
- Schwarz, N. (1990). Feelings as information: Informational and motivational functions of affective states. In T. E. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 527–561). New York, NY, US: Guilford Press.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
- Sitzmann, T., & Ely, K. (2011). A meta-analysis of self-regulated learning in work-related training and educational attainment: What we know and where we need to go. *Psychological Bulletin*, 137, 421–442.
- Sitzmann, T., & Yeo, G. (2013). A meta-analytic investigation of the within-person self-efficacy domain: Is self-efficacy a product of past performance or a driver of future performance? *Personnel Psychology*, 66, 531–568.
- Sonnetag, S., & Frese, M. (2012). Dynamic performance. In S. W. J. Kozlowski (Ed.), *The Oxford handbook of industrial and organizational psychology*. New York: Oxford University Press.
- Steel, P., & König, C. J. (2006). Integrating theories of motivation. *Academy of Management Review*, 31(4), 889–913.
- Sun, R., Coward, L. A., & Zenzen, M. J. (2005). On levels of cognitive modeling. *Philosophical Psychology*, 18(5), 613–637.
- Sun, R. (2008). Introduction to computational cognitive modeling. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 3–19). New York, NY: Cambridge University Press.
- Thibault, J. W., & Kelley, H. H. (1959). *The social psychology of groups*. New Brunswick: Transaction.
- Thomas, M. S. C., & McClelland, J. L. (2008). Connectionist models of cognition. In R. Sun (Ed.), *Cambridge Handbook of Computational Psychology* (pp. 23–58). New York: Cambridge University Press.
- Vancouver, J. B. (2000). Self-regulation in organizational settings: A tale of two paradigms. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of Self-Regulation* (pp. 303–341). San Diego, CA: Academic Press.
- Vancouver, J. B. (2005). The depth of history and explanation as benefit and bane for psychological control theories. *Journal of Applied Psychology*, 90, 38–52.
- Vancouver, J. B. (2008). Integrating self-regulation theories of work motivation into a dynamic process theory. *Human Resource Management Review*, 18, 1–18.
- Vancouver, J. B. (2012). Rhetorical reckoning: A response to Bandura. *Journal of Management*, 38, 465–474.
- Vancouver, J. B., & Day, D. V. (2005). Industrial and organization research on self-regulation: From constructs to applications. *Applied Psychology: An International Review*, 54, 155–185.
- Vancouver, J. B., More, K. M., & Yoder, R. J. (2008). Self-efficacy and resource allocation: Support for a nonmonotonic, discontinuous model. *Journal of Applied Psychology*, 93, 35–47.
- Vancouver, J. B., Putka, D. J., & Scherbaum, C. A. (2005). Testing a computational model of the goal-level effect: An example of a neglected methodology. *Organizational Research Methods*, 8, 100–127.
- Vancouver, J. B., & Scherbaum, C. A. (2008). Do we self-regulate actions or perceptions? A test of two computational models. *Computational and Mathematical Organizational Theory*, 14, 1–22.
- Vancouver, J. B., Thompson, C. M., & Williams, A. A. (2001). The changing signs in the relationships between self-efficacy, personal goals, and performance. *Journal of Applied Psychology*, 86, 605–620.
- Vancouver, J. B., & Weinhardt, J. M. (2012). Modeling the mind and the milieu: Computational Modeling for micro-level organizational researchers. *Organizational Research Methods*, 15, 602–623.
- 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, 985–1008.
- Vigo, R., & Doan, C., (2013). *The structure of choice* (submitted for publication).
- Vigo, R., & Zeigler, D., (2013). *Attention and complexity in temporal judgments* (submitted for publication).
- Vogel, D. D. (2005). A neural network model of memory and higher cognitive functions. *International Journal of Psychophysiology*, 55, 3–21.
- Vroom, V. H. (1964). *Work and motivation*. Oxford, England: Wiley.
- Wang, D., & Arbib, M. A. (1993). Timing and chunking in processing temporal order (1993). *IEEE Transactions on Systems, Man, and Cybernetics*, 23, 993–1009.
- Weinhardt, J. M., & Vancouver, J. B. (2012). Computational models and organizational psychology: Opportunities abound. *Organizational Psychology Review*, 2, 267–292.
- Weiss, H. M. (1990). Learning theory and industrial and organizational psychology. In M. D. Dunnette, & L. M. Hough (Eds.), *Handbook of industrial and organizational psychology* (Vol. 1, 1st ed., pp. 171–221). Palo Alto, CA.
- Weiss, H. M., & Cropanzano, R. (1996). Affective events theory: A theoretical discussion of the structure, causes and consequences of affective experiences at work. In B. M. Staw & L. L. Cummings (Eds.), *Research in organizational behavior: An annual series of analytical essays and critical reviews* (Vol. 18, pp. 1–74). Greenwich, CT: JAI Press.
- Widrow, G., & Hoff, M. (1960). Adaptive switching circuits. In *Western electronic show and convention record* (pp. 96–104). Institute of Radio Engineers (reprinted in Anderson, J., & Rosenfeld, E. (1988). *Neurocomputing: Foundations of research* (pp. 123–134). Cambridge, MA: MIT Press).
- Wills, A. J. (2005). *New directions in human associative learning*. Mahwah, NJ: Lawrence Erlbaum.
- Young, M. E., & Wasserman, E. A. (2005). Theories of learning. In K. Lamberts & R. L. Goldstone (Eds.), *The handbook of cognition* (pp. 161–182). London: Sage.