

# The dynamics of avoidance goal regulation

Timothy Ballard<sup>1,2</sup>  · Gillian Yeo<sup>1</sup> · Jeffrey B. Vancouver<sup>3</sup> · Andrew Neal<sup>2</sup>

© Springer Science+Business Media, LLC 2017

**Abstract** An avoidance goal is an undesired state from which a person seeks to distance themselves. Though important for understanding behavior, avoidance goals have received less attention than approach goals. In this paper, we present a dynamic, formal model that provides a framework for describing and predicting the dynamics of avoidance goal regulation. We conduct a series of simulations to examine the dynamic pattern of behavior that emerges from the model when an avoidance goal is pursued in isolation and when an approach goal is also present. Two versions of the model were examined. In the first, the avoidance goal is regulated by a positive feedback loop. In the second, the avoidance goal is regulated by a negative feedback loop. We find that the positive feedback model produces a pattern of runaway behavior, even in a scenario where an approach goal is also present. By contrast, the negative feedback loop model produces a stable pattern of behavior that is more consistent with existing theory. The findings provide an important step toward theoretical parsimony by demonstrating that avoidance goal regulation, like approach goal regulation, can be understood using a negative feedback control system framework. We discuss new insights provided by this model and its potential to spark empirical research.

**Keywords** Avoidance · Goal regulation · Motivation · Formal theory · Multiple goals

Avoidance goal regulation involves acting to distance oneself from something undesired; whereas approach goal regulation involves acting to obtain something desired. Although the ability to avoid undesired outcomes is crucial for survival (Roskes et al. 2014), avoidance goals have received far less empirical and theoretical attention than approach goals within the self-regulation literature. For example, approach goal regulation has been represented in dynamic, formal models based on the negative feedback loop architecture found in control theory (Powers 1973; Scherbaum and Vancouver 2010; Vancouver et al. 2005, 2010, 2014; Vancouver and Scherbaum 2008). Moreover, these models have been evaluated based on data available from observational and experimental studies using longitudinal designs that enable researchers to examine the dynamics of the approach goal regulation process (e.g., Kernan and Lord 1990; Schmidt and DeShon 2007; Schmidt et al. 2009; Schmidt and Dolis 2009). This level of theoretical rigor has not occurred in the avoidance goal regulation context. Therefore, despite its importance, the avoidance goal regulation process is still not well understood.

To address this issue, we present a dynamic, formal theory that describes the process of avoidance goal regulation. The model is based on the same feedback control system architecture that has been used to understand approach goal regulation. We develop two versions of the model using this architecture—one that assumes avoidance goal regulation is regulated by a positive feedback loop and another that assumes the process is regulated by a negative feedback loop. We then present a set of simulations that show the dynamic pattern of behavior that each model generates over

---

✉ Timothy Ballard  
t.ballard@uq.edu.au

<sup>1</sup> UWA Business School, University of Western Australia, Crawley, Perth, WA 6009, Australia

<sup>2</sup> School of Psychology, University of Queensland, St Lucia, QLD 4072, Australia

<sup>3</sup> School of Psychology, Ohio University, Athens, OH 45701, USA

time. The positive feedback loop model produces a runaway pattern of avoidance behavior that accelerates over time. The negative feedback loop model produces a stable pattern of avoidance behavior that decelerates as the undesired outcome becomes more distant. We conclude that the behavior generated by the negative feedback model is the more plausible model of avoidance goal regulation. We then discuss the insights this model provides about goal regulation and its potential to be used for future theoretical development.

## Approach and avoidance goal regulation

Two broad theories have been developed that attempt to explain the underlying process through which avoidance goals influence behavior. Townsend and Busemeyer (1989) developed a framework derived from Lewin's Field Theory and research on goal gradient effects, mostly conducted with animals (e.g., Hull 1932). This framework assumes that motivation to approach or avoid is determined by motivational gradients, which describe the change in motivation as the goal becomes nearer. The framework proposes that motivation increases as the approach or avoidance goal becomes nearer, though this increase is said to be stronger for avoidance goals. A problem with this framework is that the empirical work that has examined how humans pursue goals—which has dealt mainly with approach goals—has shown that people tend to prioritize the goal with the largest discrepancy (Neal et al. 2017). This finding suggests that when pursuing approach goals, motivation is higher when the goal is farther away than when it is nearer. Thus, the empirical work on the pursuit of approach goals is not consistent with the goal gradient perspective.

An alternative theoretical approach that considers approach and avoidance goals can be found in the self-regulation literature. For example, Carver and Scheier (1998) developed a theory of self-regulation based on principles from control theory (e.g., Powers). Their theory assumes that the information processing architecture responsible for the regulation of approach and avoidance motivation can be described in terms of feedback loops. Specifically, approach behavior is regulated by discrepancy-reducing feedback loops. For example, a hunter-gatherer might go into the forest to forage for nuts and berries. The hunter-gatherer needs to gather enough to feed her family. If the current quantity of nuts and berries is less than desired, she will be motivated to continue foraging until she's obtained the desired level (i.e., the goal). The intensity of this motivation will depend on how large the discrepancy is between her current quantity and the desired level. Large discrepancies produce a more intense experience of motivation, because more foraging needs to be done in order to reach the desired quantity. This control system is a *negative feedback loop*

because the output of the system (i.e., gathering nuts and berries) opposes or counteracts the discrepancy (i.e., the amount of nuts and berries desired minus the amount gathered), bringing it closer to the goal (Lord and Levy 1994)<sup>1</sup>. As the discrepancy decreases, the motivation to keep moving closer tends to decrease because there is less need to do so. Importantly, this information processing system has been represented formally and been found to explain behavior in situations where people are pursuing multiple goals simultaneously (Vancouver et al. 2010, 2014).

In contrast, the architecture that explains avoidance goal motivation has received less attention. In general, Carver and Scheier (1998) suggest that avoidance behavior is regulated by discrepancy-amplifying loops. For example, our hunter-gatherer might be concerned about predators. Thus, the hunter-gatherer is looking for cues that may indicate whether a predator is nearby (e.g., bear tracks). In this case, the discrepancy that drives behavior is the proximity of the predator, which might be indicated by the age of the tracks. The feedback loop is discrepancy-amplifying, because the hunter-gatherer seeks to increase this discrepancy in order to avoid the predator.

Carver and Scheier (1998) describe this type of discrepancy-amplifying structure as a *positive feedback control system*. A positive feedback loop is one in which the output reinforces the input. Positive feedback loops are self-perpetuating or self-reinforcing. For example, one of the reasons why climate change is thought to be so dangerous is that it triggers positive feedback loops (Castro de la Guardia et al. 2015; Lewandowsky et al. 2015). In one such loop, increases in global temperatures cause polar ice caps to melt, reducing the amount of sunlight that gets reflected by these ice caps back into space, producing further increases in global temperatures. For this reason, positive feedback loops are unbounded and often produce behavior that proceeds without limit (i.e., are unstable). Carver and Scheier recognized that this is a problem for any self-regulating system. They addressed the issue of positive feedback in avoidance goal regulation by arguing that approach goals act to constrain avoidance behavior. When one is regulating an avoidance goal, an approach goal is identified that facilitates and provides an end point for the avoidance process. For example, the hunter-gatherer may decide that a bear track that is 6 h old is sufficient for her to feel safe. The desired distance then becomes an approach goal that constrains the avoidance behavior.

Although avoidance goal regulation is often described as a positive feedback control system, it is also possible for

<sup>1</sup> Formally, a feedback loop is negative if there is an odd number of links between the variables in the loop. Otherwise, the loop is positive (Richardson 1991).

discrepancy-amplifying loops to be regulated by negative feedback control systems. Recent empirical and theoretical work on multiple-goal pursuit has shown that motivation to prioritize avoidance goals increases as the undesired state becomes closer (Ballard et al. in press; Ballard et al. 2016a, b). These findings suggest that avoidance goal regulation may be able to be explained using a negative feedback control system architecture, because acting on avoidance motivation reduces the need for further avoidance over time. For example, acting on the urge to flee the bear by running in the other direction will increase the amount of distance between the bear and the hunter-gatherer, which will lead to a reduction in the intensity of her urge to flee.

The negative feedback model supports a more conceptually parsimonious theory of goal pursuit by unifying approach and avoidance goal regulation under a common architecture (i.e., the negative feedback loop). However, the dynamic properties of this model have not been fully explored. In the current paper, we present a formalized model of avoidance goal regulation. This enables us to simulate the model and examine the dynamic behavior that emerges from it over time. Specifically, we examine two versions of the model—one in which avoidance goal regulation is a negative feedback process and one in which it is a positive feedback process—and compare their results.

### Formalizing the model of goal regulation

To refine our theoretical understanding of avoidance goal regulation and provide a framework for predicting the dynamics of this process, we present a formal model. Avoidance goal regulation, by definition, must be modeled as a discrepancy-amplifying feedback loop. However, as discussed above, discrepancy-amplifying behavior can be produced by either positive or negative feedback control systems. We therefore create two versions of the avoidance goal regulation model—one in which the process is regulated by a positive feedback loop and one in which it is regulated by a negative feedback loop. Our models of avoidance goal regulation are based on the same broad framework that has been used to formally model approach goal regulation (e.g., Scherbaum and Vancouver 2010; Vancouver et al. 2005, 2010, 2014), which is Powers' (1973) control theory. Here, we describe how this framework can be generalized to accommodate both approach and avoidance goals.

#### Input

The first component of the feedback loop, as shown in Fig. 1, is the input function. The input function represents the processes that create the individual's perception ( $p$ ) of the current state of the variable that the system is regulating. In

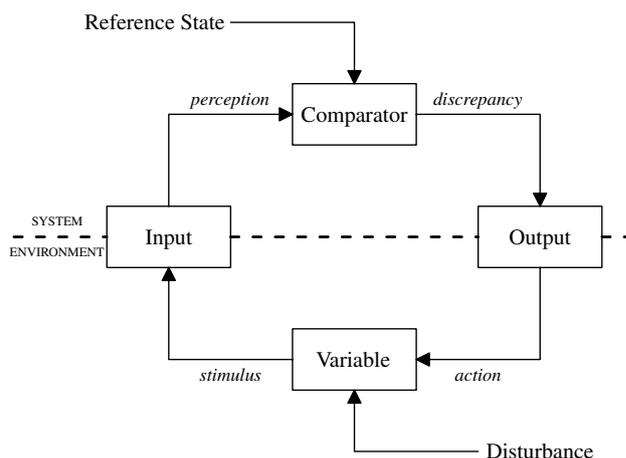


Fig. 1 A generic feedback control system

the example of the hunter-gatherer, the input function may simply translate the amount of nuts or berries he or she has gathered into a perception of the total quantity, or it may include a mental model that allows the hunter-gatherer to estimate the amount of nuts and berries that will be edible given that some of the fruit may spoil. In the case where there is no bias or error in this perception, one can define the perception as follows:

$$p(t) = v(t), \tag{1}$$

where  $v(t)$  represents the current state of the variable at time  $t$ .

#### Comparator

The comparator determines the discrepancy ( $d$ ) between the current state and the desired or undesired reference state. The discrepancy for an approach goal can be calculated by subtracting the current state of the variable from the desired reference state:

$$d(t) = p' - p(t), \tag{2a}$$

where  $p'$  represents the reference state<sup>2</sup>. In the approach case, the person is motivated to eliminate the discrepancy. In the avoidance case, by contrast, the person is often concerned with keeping the current state of a variable above some undesired level (e.g., keeping the age of the tracks greater than a minimum value). In this case, the discrepancy can be calculated as:

<sup>2</sup> Some models (e.g., Vancouver et al. 2005, 2008) have included a condition in the comparator function that restricts the value of the discrepancy from taking on negative values. We omit that condition for simplicity. In the simulations presented below, adding this condition makes no difference to the results.

$$d(t) = p(t) - p'. \quad (2b)$$

### Output

The intensity of motivation to engage in the behavior aimed at reducing or amplifying the discrepancy is typically assumed to be proportional to the size of the discrepancy. This motivation is referred to as the output ( $o$ ). The impact of the discrepancy may also be influenced by individual differences or situational factors that influence one's sensitivity to the discrepancy (Carver and Scheier 1982; Hyland 1987). For example, the hunter-gatherer might be more sensitive to the discrepancy between her current and desired quantities of nuts and berries if her family has not eaten in days and is at risk of starvation. In this case, discrepancies of the same size may yield a greater output than they would otherwise (Schmidt and DeShon 2007). For approach goals, the output function can be defined as follows:

$$o(t) = kd(t), \quad (3a)$$

where  $k$  is a gain parameter that represents the individual's sensitivity to the discrepancy. In this way, gain represents the importance of the goal to the individual (Hyland 1987).

For avoidance goals, the intensity of avoidance motivation is higher when the variable that is being monitored is closer to the undesired reference state (i.e., when the discrepancy is lower), and lower when the variable is further from the undesired reference state. Because the output increases as the discrepancy decreases, an intercept parameter ( $b$ ) is required to represent the intensity of avoidance motivation when the current state reaches the undesired state. We define the output function of the negative feedback model as follows:

$$o(t) = b - kd(t). \quad (3b)$$

The intercept ( $b$ ) influences the height of the function. Higher values of  $b$  strengthen the overall avoidance motivation. Higher values of gain ( $k$ ) mean that avoidance motivation increases more rapidly as the discrepancy decreases.

A final function is needed to complete the feedback loop. This last function represents the processes that determine the new state of the variable. Generally, variables that are regulated have memory. That is, they retain their value over time. Thus, the new state is a function of the previous state, the effects of output (i.e., action) from the control system, the rate ( $r$ ) at which the output affects the current state, and environmental disturbances ( $D$ ). In our example, the rate represents the ease with which the hunter-gatherer can increase the age of the tracks (i.e., by distancing herself from the predator). The rate is influenced by factors such as the speed and direction of the predator's movement, and the speed at which the hunter-gatherer can move. Environmental disturbances are externally determined influences that produce a

change in the current state, for example, the hunter-gatherer may encounter fresh tracks from a new predator. Given all of the above, the new state of the variable can be defined as follows:

$$v(t + 1) = v(t) + ro(t) + D(t). \quad (4)$$

Under the model described above, both the approach and avoidance goal regulation systems are negative feedback control systems. They are negative feedback systems because the output of each system has an effect that reduces the subsequent need for or intensity of that output. As an example of the approach case, spending a day collecting nuts and berries will reduce the discrepancy between the hunter-gatherer's current quantity and her desired quantity, which reduces the need for her to collect more the next day. As an example of the avoidance case, the hunter who runs away from a predator will put more space in between him or herself and the threat, reducing the need to spend more time and energy moving away.

## Simulations

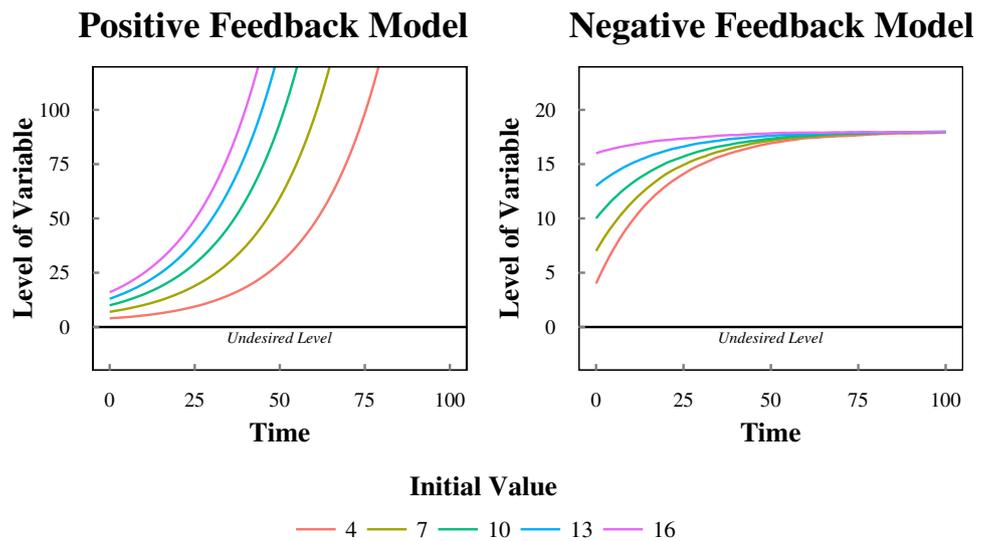
In this section, we simulate the negative feedback model to examine the dynamic behavior that emerges from it over time. As a comparison, we simulate a model which assumes that avoidance is regulated by a positive feedback loop, so that we can demonstrate how the sign of the feedback loop influences the predictions made by the model. The positive feedback model was constructed by swapping Eq. 3b in the avoidance model with Eq. 3a. Using Eq. 3a to describe the relationship between the discrepancy and the output of the avoidance system makes it a positive feedback loop, because it creates a self-perpetuating relationship between the discrepancy and the output. That is, as the discrepancy increases, the output also increases, and this output serves to increase the discrepancy further.

We first simulate the two models in a scenario where a single avoidance goal is being pursued in isolation, using the example of the hunter-gatherer trying to avoid encountering a predator in the woods. The purpose of this simulation is to establish the baseline pattern of behavior that each model produces when no other goals are present. We then simulate the two models in a more realistic scenario where there are multiple goals that need to be regulated. In this simulation, we examine the pattern of behavior that emerges when approach and avoidance goals are simultaneously active.

### Single-goal context

In our example, the variable that the avoidance system monitors is the age of the tracks. We set the undesired level of this variable to zero hours (i.e., the predator is in the immediate

**Fig. 2** Predicted level of the variable as a function of time and initial variable level for the positive and negative feedback models when regulating a single avoidance goal



vicinity), and the initial level to either 4, 7, 10, 13, or 16 h, so that we can examine the emergent behavior under different starting conditions. We set the rate to 0.1. The disturbance is a normally distributed random variable with a mean of  $-0.1$  and a standard deviation of 1. Thus, on average the disturbance has detrimental effects in the simulations, so that if the person does nothing, the variable will eventually reach the undesired level (e.g., if the hunter-gatherer takes no action to avoid contact with predators, she is likely to eventually encounter one). We set gain to 1 so that it has no effect on the model’s behavior. We set the intercept for the negative feedback model to 20, indicating that the avoidance output will equal 20 if the undesired state is reached. Because the disturbance is a random variable, we need to conduct multiple simulations to achieve reliable predictions. We therefore simulate each model 1000 times and average the predictions across simulations.

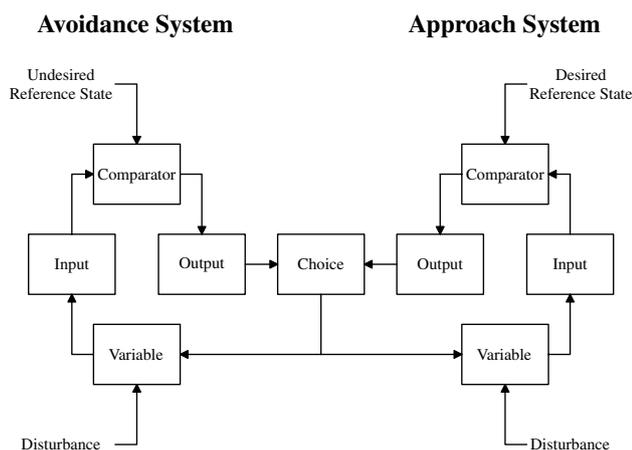
The two models produce different patterns of dynamic behavior (see Fig. 2). Both models are discrepancy-amplifying systems, and therefore both move the variable away from the undesired state. However, under the positive feedback model, the level of the variable increases at an accelerating rate, and rapidly approaches infinity. This behavior emerges because the intensity of avoidance behavior increases over time, as the level of the variable gets further away from the undesired level. The acceleration in the level of the variable happens regardless of the initial level, although the acceleration happens faster when the initial level is higher. Under the negative feedback model, the level of the variable increases at a decelerating rate, eventually leveling off. This behavior emerges because the intensity of avoidance behavior decreases over time, as the level of the variable gets further away from the undesired level. This deceleration happens regardless of the initial level of the variable. However, as

the initial level of the variable increases, the increase in the level of the variable over time becomes weaker.

**Multiple-goal context**

Having established a baseline for the behavior of each model when no other goals are present, we next demonstrate the emergent behavior of the models in a scenario where both approach and avoidance goals are present and the person needs to choose whether to approach or avoid. This type of scenario is important to consider because people rarely pursue goals in isolation, and are often striving for approach and avoidance goals simultaneously (Elliot and Sheldon 1997). To demonstrate the models’ behavior in a scenario where an approach goal is present, we construct two separate feedback systems—one that regulates the avoidance goal and another that regulates the approach goal—and assume that behavior is determined by the system that is more active (i.e., has the greater output) at any given point in time (e.g., Vancouver et al. 2010; see Fig. 3). This configuration is consistent with the notion that, although an individual may experience approach and avoidance motivations simultaneously, they often are only able to act on one at a time. For example, it is unlikely that our hunter-gatherer would stop to harvest nuts or berries (i.e., act on the approach goal), if she is running from a bear (i.e., acting on the avoidance goal). When the output from the avoidance feedback system is greater (e.g., because she believes that there is a bear close by), the system’s behavior is driven by the avoidance goal. When the output from the approach system is greater (e.g., because she does not have enough food to feed her family), the behavior is driven by the approach goal.

$$\text{If } o_{AP}(t) > o_{AV}(t), \tag{6}$$



**Fig. 3** A graphical representation of the process of regulating an avoidance goal in the presence of an approach goal

$$v_{AP}(t+1) = v_{AP}(t) + r_{o_{AP}}(t) + D_{AP}(t),$$

$$v_{AV}(t+1) = v_{AV}(t) + D_{AV}(t).$$

$$\text{If } o_{AP}(t) < o_{AV}(t),$$

$$v_{AP}(t+1) = v_{AP}(t) + D_{AP}(t),$$

$$v_{AV}(t+1) = v_{AV}(t) + r_{o_{AV}}(t) + D_{AV}(t).$$

When the outputs for the two systems are equal, the goal which drives behavior is randomly determined.

We simulate the multiple-goal context in a similar way to the single avoidance goal context. For the approach system, the variable being regulated was the quantity of nuts and berries (in kilograms). We set the initial level to 0 kg and the desired level to 10 kg. For the avoidance system, we once again set the initial level to either 4, 7, 10, 13, and 16 h, and the undesired level to 0 h. This manipulation enables us to examine scenarios in which the discrepancy for the avoidance system is smaller than, larger than, or equal to the discrepancy for the approach system. The rates, disturbances, gains, and the intercept (for the negative feedback model of the avoidance system) are identical to the previous simulations. Once again, we produce the patterns predicted by each model by simulating them 1000 times and averaging across simulations.

The results reveal that, as with the baseline scenario, when an approach goal is present, the dynamic pattern of behavior that emerges differs between the two models (see Fig. 4). Regardless of the initial level of the avoidance variable, the positive feedback model once again produces runaway avoidance behavior, with the level of the avoidance variable increasing at an accelerating rate and approaching infinity, and the desired state never being reached. However, the initial level determines when the avoidance behavior begins to accelerate. When the initial level of the avoidance

variable is greater than 10 (i.e., when the avoidance discrepancy is greater than the approach discrepancy), the avoidance system output starts off greater than the approach system output. As a result, the system quickly moves away from the undesired state, without making any progress toward the desired state.

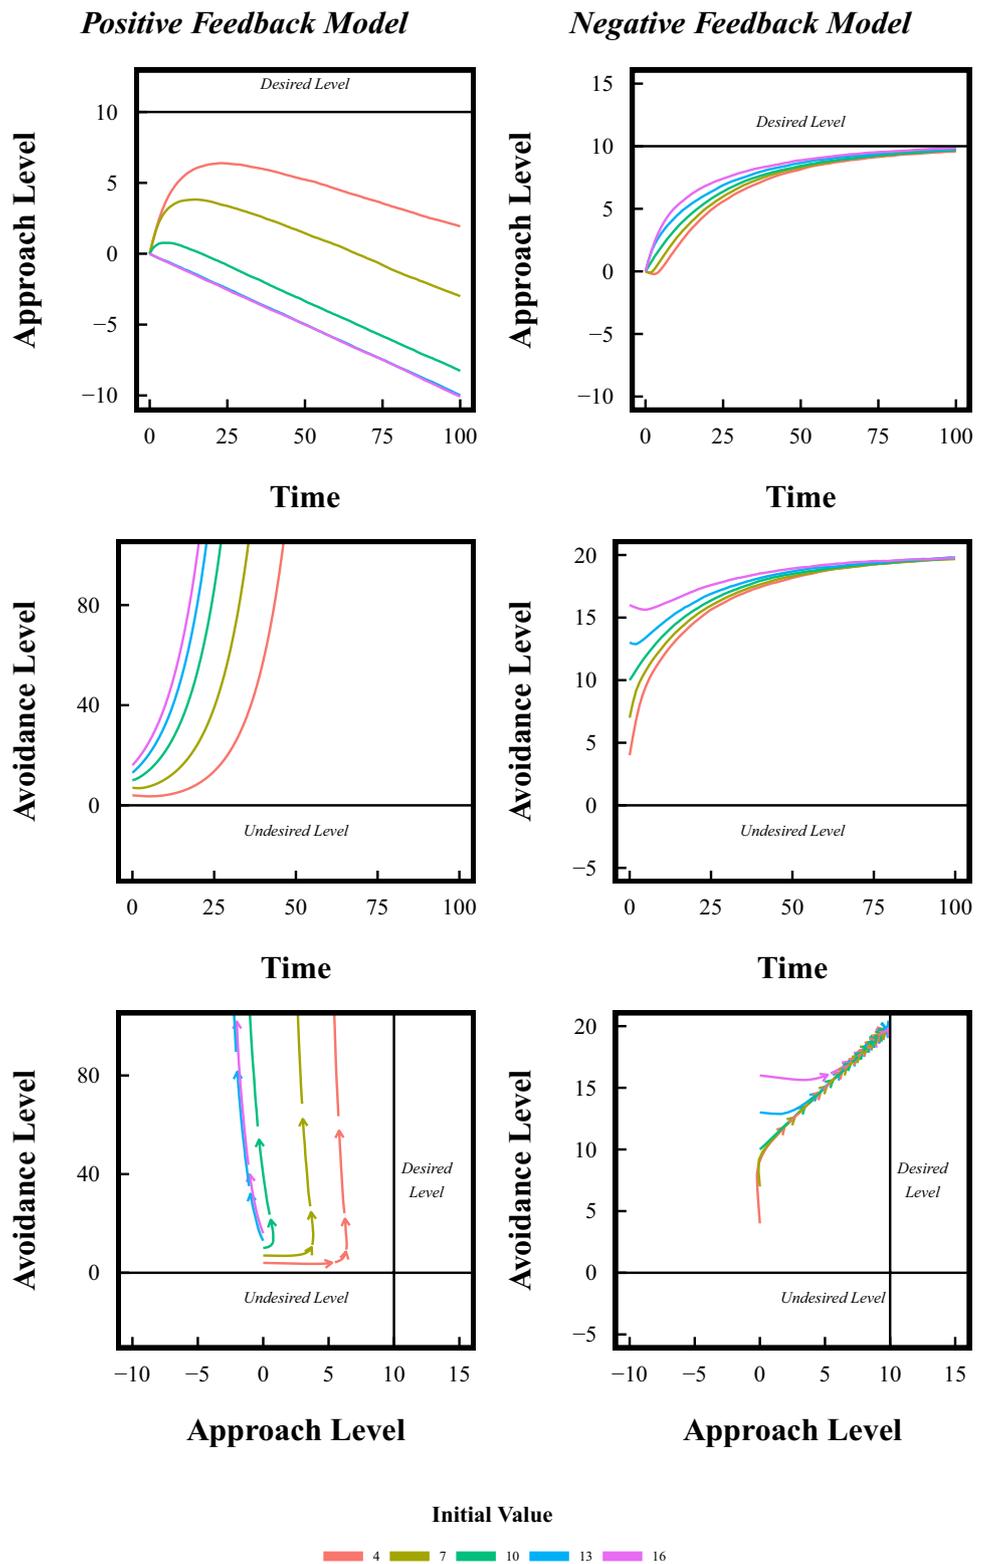
When the initial level of the avoidance variable is 10, the discrepancies and therefore the output for the approach and avoidance systems are equal. Under these conditions, avoidance behavior accelerates less quickly than when the initial level is greater than 10 and the system makes some progress toward the desired state. However, progress towards the desired state ceases relatively early in the simulation, as the avoidance behavior begins to accelerate. When the initial level of the avoidance variable is less than 10 (i.e., the discrepancy for the avoidance system is smaller than the discrepancy for the approach system), the approach system output starts off greater than the avoidance system output. Under these conditions, the model makes the most progress towards the desired state, but still fails to achieve it. As progress is made toward the desired state, the approach system output reduces enough that the avoidance output becomes stronger. At this point, the avoidance behavior accelerates and progress towards the desired state stagnates. In summary, the positive feedback model produces a pattern of behavior in which (a) the avoidance behavior never ceases, but rather accelerates over time, and (b) the desired state is never reached.

By contrast, the negative feedback model produces a pattern of behavior in which the avoidance behavior decreases over time and the desired state is reached, regardless of the initial level of the avoidance variable. When the initial level of the avoidance variable is greater than 10, the output for the approach system is stronger than the output for the avoidance system. As a result, the model begins making more progress towards the desired state than away from the undesired state. When the initial level of the variable is less than 10, the output for the avoidance system is stronger than the output for the approach system. As a result, the model begins making more progress away from the undesired state than towards the desired state. However, regardless of the initial level of the avoidance variable, the model quickly converges to a point where the outputs for the approach and avoidance systems were equal. When this point is reached, the model moves toward the desired state and away from the undesired state at the same rate until the desired state is reached.

## Discussion

Avoidance goal regulation is critical to our survival. Yet this process has not received the theoretical or empirical attention that it deserves (Neal et al. 2017). To enhance

**Fig. 4** Predicted levels of the approach and avoidance variables as a function of time and initial avoidance variable level for the positive and negative feedback models when regulating an avoidance goal in the presence of an approach goal



understanding of this process and provide a framework for generating predictions about its dynamics, we presented a formal model of avoidance goal regulation. We examined

the dynamic behavior that emerged from two versions of the model. One version assumed that avoidance goals are regulated by a positive feedback loop, whereas the other assumed

that avoidance goals are regulated by a negative feedback loop. The dynamic pattern of behavior that emerged from the positive feedback model showed the intensity of avoidance behavior (i.e., the output) accelerating over time, regardless of whether or not an approach goal was present. The dynamic pattern of behavior that emerged from the negative feedback loop shows the intensity of avoidance behavior decreasing over time, as the undesired state becomes more distant. We argue that the emergent behavior of the negative feedback model was more consistent with the dynamic process described by existing theory (e.g., Carver and Scheier 1998), in which the avoidance behavior reduces in intensity over time as the undesired state becomes more distant. We conclude that the negative feedback loop model is the more plausible account of avoidance goal regulation.

### Contributions to the goal regulation literature

The dynamic, formal model presented in this paper can help facilitate theory building regarding avoidance goal regulation in much the same way that these models have accelerated progress in the approach goal regulation literature. One advantage of a dynamic, formal model of avoidance goal regulation is that it can make transparent the process theorists are describing. A second advantage is that the model can be simulated, and used to demonstrate the phenomena that emerge as the interaction between system and environment plays out over time. We hope that this model will be used as a framework for generating predictions about how the avoidance goal regulation plays out over time that can be tested in future research.

We have also provided insights about the relationship between the discrepancies between current and undesired states and the intensity of avoidance behavior. Existing theory recognizes that avoidance goal regulation involves discrepancy-amplifying behavior, in which the person is motivated to engage in behaviors that increase the size of the discrepancy. However, it does not address how the intensity of avoidance behavior changes as a function of the size of the discrepancy. Understanding this relationship is critical for understanding how avoidance goal regulation plays out over time because, as we have shown, it has a significant impact on the dynamic behavior that emerges from the model. It is the relationship between the discrepancy and the intensity of the behavior that determines whether the feedback loop is positive or negative. Assuming that the discrepancy/intensity relationship is positive, such that avoidance behavior intensifies as the undesired state becomes more distant, equates to an assumption that avoidance goals are regulated by a positive feedback loop. When a model containing this assumption is simulated, the avoidance behavior produced by the model accelerates over time. Assuming the discrepancy/intensity relationship is negative, such that avoidance

behavior reduces in intensity as the undesired state becomes more distant, equates to an assumption that avoidance goals are regulated by a negative feedback loop. When a model containing this assumption is simulated, the avoidance behavior tapers off over time. We have shown that a model that assumes the avoidance behavior reduces in intensity as the undesired state becomes more distant, and is therefore a negative feedback loop, provides a more plausible account of avoidance goal regulation.

Our work suggests that both approach and avoidance goal regulation can be understood using the principles of negative feedback control. This provides an important step towards a parsimonious theory of goal regulation, because it unifies approach and avoidance within a common framework. The unifying principle is that a person's motivation to engage in goal-directed behaviors is higher when the current state is relatively unfavorable (i.e., when a desired state is far away or when an undesired state is near) and lower when the current state is relatively favorable (i.e., when a desired state is near or when an undesired state is far away). Consequently, both types of goal regulation processes should stabilize over time. When regulating an approach goal, the intensity of approach behavior should reduce in intensity over time as the individual gets closer to the desired state, and eventually stabilize when the desired state is reached. When regulating an avoidance goal, the intensity of avoidance behavior should reduce in intensity over time as the individual gets further away from the undesired state, and eventually stabilize when the individual is far enough away that the motivation to avoid is negligible. Thus, the negative feedback model of avoidance is parsimonious because it does not require an approach goal to constrain the runaway avoidance behavior. Instead, the avoidance behavior stabilizes on its own as the undesired state becomes sufficiently distant.

### Additional considerations and avenues for future research

In order to implement the models, it was necessary to make certain assumptions about how goal regulation plays out over time. One assumption is that the output of the loop (i.e., the intensity of the motivation) is related to the size of the discrepancy. Although this assumption is consistent with previous models of approach goal regulation (e.g., Vancouver 2008; Vancouver and Purl 2017), there are discrepancy-reducing systems in which the output is not related to the size of the discrepancy. One example is the thermostat, where the same output is triggered (i.e., heating) whenever a discrepancy between the current and desired temperature exists. Discrepancy-reducing systems will tend to converge around the reference value regardless of whether the discrepancy/output relationship is positive, negative, or non-existent (e.g., in the thermostat example). However, as we

have shown, the dynamic behavior of discrepancy-amplifying systems differs markedly depending on the nature of this relationship. It is therefore important to consider the relationship between the discrepancy and the intensity of avoidance motivation when understanding and predicting the dynamics of avoidance goal regulation.

When constructing the simulation for the multiple-goal context, we also made the assumption that people have to choose between goals, because only one could be acted on at a time. An alternative approach could be to assume that behavior could be directed toward both goals simultaneously, and that the degree to which the behavior affects one goal or another is determined by the relative outputs of the two systems. However, we believe that the range of situations to which this alternative model applies is narrower. In many cases, there is some constraint in the environment or the person's resources that prevents both goals being acted on simultaneously. For example, the hunter-gatherer will likely find it difficult to forage for nuts and berries whilst she is fleeing the predator. Likewise, an academic striving to increase his or her research productivity (an approach goal), whilst striving to avoid poor teaching ratings (an avoidance goal) will have to decide whether to spend time working on a paper or preparing for a lecture, because he/she cannot do both at once. It is important to note though that the alternative approach described above produces the same basic pattern of dynamic behavior, in that the positive feedback loop model produces runaway avoidance behavior, whilst the negative feedback model stabilizes over time.

We limited our focus to the theory and research that derived from control theory (Powers 1973), because it has produced a rigorous framework for formal theorizing. However, there are a number of other theories that are relevant to avoidance goal regulation. For example, Elliot and Sheldon (1997) address the relationship between avoidance goal regulation and well-being, and suggest that regulating avoidance goals has undesirable consequences that extend beyond the goal itself to affect one's overall level of personal adjustment. Scholer and Higgins (2013) propose that there are different types of avoidance goals that people can pursue. Promotion focused individuals may be more motivated to regulate avoidance goals by approach "mismatches" to the undesired state, whereas prevention focused individuals may do so by avoiding "matches" to the undesired state. These individual differences can be incorporated into the model by assuming that they the importance of different types of goals (i.e., via the gain parameter).

Future research should seek to empirically test the predictions of the model presented in this paper. As modelled in this research, avoidance goal regulation is a dynamic process that unfolds over time. Thus, an appropriate empirical test of the model's predictions requires multiple measurements of (a) discrepancy and (b) an indicator of motivation such as

effort, time, or resource allocation, whilst participants pursue an avoidance goal. This test can be achieved using a variety of empirical methods. For example, some research has used diary studies which elicit repeated self-report measures of the relevant constructs (Eddington et al. 2012; Louro et al. 2007; Righetti et al. 2010). Other research has used laboratory experiments that obtain a behavioral measure (Ballard et al. in press; Ballard et al. 2016a; Schmidt and DeShon 2007; Schmidt and Dolis 2009; Schmidt et al. 2009). Ideally, a combination of approaches should be used to ensure methodological generalizability.

## Conclusion

Although our ability to survive and thrive as a species depends critically on avoidance goal regulation, the avoidance goal regulation process has received little theoretical or empirical attention. In an effort to accelerate progress in the field, we have presented a dynamic, formal model of avoidance goal regulation and examined the pattern of behavior that emerges when different versions of this model are simulated. The findings have improved our understanding of how avoidance goals regulate behavior by showing that avoidance goal regulation can be understood using negative feedback control architecture. This model provides an important step towards theoretical parsimony by unifying approach and avoidance goal regulation under a common architecture. It also provides a useful framework for generating predictions about the dynamics of avoidance goal regulation that can be used in future research. We hope that these new insights and tools will provide a spark that ignites future theoretical and empirical endeavors.

**Acknowledgements** This work was supported by an Australian Research Council Grant DP120100852, awarded to Neal (CI), Yeo (CI), Zacher (CI), Vancouver, & Schmidt.

## References

- Ballard, T., Farrell, S., & Neal, A. (in press). Quantifying the psychological value of goal achievement. *Psychonomic Bulletin & Review*.
- Ballard, T., Yeo, G., Loft, S., Vancouver, J. B., & Neal, A. (2016a). An integrative, formal model of motivation and decision making: The MGPM\*. *Journal of Applied Psychology*, *101*, 1240–1265. doi:[10.1037/apl0000121](https://doi.org/10.1037/apl0000121).
- Ballard, T., Yeo, G., Neal, A., & Farrell, S. (2016b). Departures from optimality when pursuing multiple approach or avoidance goals. *Journal of Applied Psychology*, *101*, 1056–1066. doi:[10.1037/apl0000082](https://doi.org/10.1037/apl0000082).
- Carver, C. S., & Scheier, M. F. (1982). Control theory: A useful conceptual framework for personality-social, clinical, and health psychology. *Psychological Bulletin*, *92*, 111–135.

- Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. New York: Cambridge University Press.
- Castro de la Guardia, L., Hu, X., & Myers, P. G. (2015). Greenland Ice Sheet melt increases Baffin Bay heat content on the west Greenland shelf. *Geophysical Research Letters*, *42*, 4922–4930. doi:10.1002/2015GL064626.
- Eddington, K. M., Majestic, C., & Silvia, P. J. (2012). Contrasting regulatory focus and reinforcement sensitivity: A daily diary study of goal pursuit and emotion. *Personality and Individual Differences*, *53*, 335–340. doi:10.1016/j.paid.2012.04.001.
- Elliot, A. J., & Sheldon, K. M. (1997). Avoidance achievement motivation: A personal goals analysis. *Journal of Personality and Social Psychology*, *73*, 171–185.
- Hull, C. L. (1932). The goal gradient hypothesis and maze learning. *Psychological Review*, *39*, 25–43.
- Hyland, M. E. (1987). Control theory interpretation of psychological mechanisms of depression: Comparison and integration of several theories. *Psychological Bulletin*, *102*, 109–121.
- Kernan, M. C., & Lord, R. G. (1990). Effects of valence, expectancies, and goal-performance discrepancies in single and multiple goal environments. *Journal of Applied Psychology*, *75*, 194–203.
- Lewandowsky, S., Ballard, T., & Pancost, R. D. (2015). Uncertainty as knowledge. *Philosophical Transactions A*, *373*, 1–11.
- Lord, R. G., & Levy, P. E. (1994). Moving from cognition to action: A control theory perspective. *Applied Psychology: An International Review*, *43*, 335–398. doi:10.1111/j.1464-0597.1994.tb00828.x.
- Louro, M. J., Pieters, R., & Zeelenberg, M. (2007). Dynamics of multiple-goal pursuit. *Journal of Personality and Social Psychology*, *93*, 174–193. doi:10.1037/0022-3514.93.2.174.
- Neal, A., Ballard, T., & Vancouver, J. B. (2017). Dynamic self-regulation and multiple-goal pursuit. *Annual Review of Organizational Psychology and Organizational Behavior*, *4*, 410–423. doi:10.1146/annurev-orgpsych-032516-113156.
- Powers, W. T. (1973). *Behavior: The control of perception*. Chicago: Aldine.
- Richardson, G. P. (1991). *Feedback thought in social science and systems theory*. Philadelphia: University of Pennsylvania Press.
- Righetti, F., Rusbult, C., & Finkenauer, C. (2010). Regulatory focus and the Michelangelo Phenomenon: How close partners promote one another's ideal selves. *Journal of Experimental Social Psychology*, *46*, 972–985. doi:10.1016/j.jesp.2010.06.001.
- Roskes, M., Elliot, A. J., & De Dreu, C. K. W. (2014). Why is avoidance motivation problematic, and what can be done about it? *Current Directions in Psychological Science*, *23*, 133–138. doi:10.1177/0963721414524224.
- 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. doi:10.1111/j.1559-1816.2010.00656.x.
- 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*, 928–941. doi:10.1037/0021-9010.92.4.928.
- Schmidt, A. M., & Dolis, C. M. (2009). Something's got to give: The effects of dual-goal difficulty, goal progress, and expectancies on resource allocation. *Journal of Applied Psychology*, *94*, 678–691. doi:10.1037/a0014945.
- Schmidt, A. M., Dolis, C. M., & Tolli, A. P. (2009). A matter of time: Individual differences, contextual dynamics, and goal progress effects on multiple-goal self-regulation. *Journal of Applied Psychology*, *94*, 692–709. doi:10.1037/a0015012.
- Scholer, A. A., & Higgins, E. T. (2013). Dodging monsters and dancing with dreams: Success and failure at different levels of approach and avoidance. *Emotion Review*, *5*, 254–258. doi:10.1177/1754073913477506.
- Townsend, J. T., & Busemeyer, J. R. (1989). Approach-avoidance: Return to dynamic decision behavior. In *Cognitive Processes: Tulane Flowerree Symposium On Cognition*.
- Vancouver, J. B. (2008). Integrating self-regulation theories of work motivation into a dynamic process theory. *Human Resource Management Review*, *18*, 1–18. doi:10.1016/j.hrmr.2008.02.001.
- Vancouver, J. B., & Purl, J. D. (2017). A computational model of self-efficacy's various effects on performance: Moving the debate forward. *Journal of Applied Psychology*, *102*, 599–616.
- 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. doi:10.1177/1094428104271998.
- Vancouver, J. B., & Scherbaum, C. A. (2008). Do we self-regulate actions or perceptions? A test of two computational models. *Computational and Mathematical Organization Theory*, *14*, 1–22. doi:10.1007/s10588-008-9021-7.
- Vancouver, J. B., Tamanini, K. B., & Yoder, R. J. (2008). Using dynamic computational models to reconnect theory and research: Socialization by the proactive newcomer as example. *Journal of Management*, *36*, 764–793. doi:10.1177/0149206308321550.
- 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. doi:10.1037/a0020628.
- Vancouver, J. B., Weinhardt, J. M., & Vigo, R. (2014). Change one can believe in: Adding learning to computational models of self-regulation. *Organizational Behavior and Human Decision Processes*, *124*, 56–74. doi:10.1016/j.obhdp.2013.12.002.