MYOPIC SEARCH AND TEMPORALLY DISTANT GOALS

Daniel Albert
LeBow College of Business
Drexel University
da692@drexel.edu

Anoop Menon
The Wharton School
University of Pennsylvania
armenon@wharton.upenn.edu

Draft version: April 23, 2020

Abstract

In this paper, we study the role of temporal myopia in implementing firms' long-term goals. In a formal model, we define temporal myopia as the extent to which a firm temporally discounts a decision's contribution to its long-term goal. We find that the degree of temporal myopia influences how sensitive firms are with respect to evaluating decisions that lead to change in the temporal distance to their long-term goal. Temporal distance changes occur when a decision shortens or lengthens the path to a long-term goal. The greater temporal myopia, the greater the relative sensitivity to temporal distance changes. In other words, for a temporally myopic firm, shortening the path to their goal becomes substantially more attractive if the closer the decision moves it to the goal. For a more patient firm, the same decision only marginally increases attractiveness as they are comparably more indifferent to a potentially change in temporal distance. This sensitivity to temporal distance changes explains when firms find and accept two types of decisions that relate to achieving long-term goals. The first type are 'stepping-stone' decisions, choices that lead to an immediate performance decline but shortens the temporal distance to the goal; the second type are 'strategic reject' decisions, choices that forgo immediate performance opportunities in order to not lengthen the temporal goal distance. The caveat to this, however, is to match the exploratory horizon, i.e., how distant of goals the firm discovers, with the degree of temporal myopia. This is further moderated by the complexity of the firm's task environment. Temporal myopic firms can perform relatively high when complexity is also high and temporal horizons are medium to short. Under medium complexity, slightly longer horizons become more favorable but require greater patience (lower temporal myopia). Finally, under low complexity, temporal myopia and optimal choice of temporal horizons are of little importance.

Introduction

An enduring question in the tradition of the behavioral theory of the firm (BToF) remains how firms can achieve temporally distant opportunities when their decision making is overwhelmingly myopic. In the BToF tradition, firms are conceptualized as complex systems of interdependent choices (Cyert and March 1963, Gavetti et al. 2012, Levinthal and March 1993). Because of interactions among choices, internally consistent solutions exist, for which any incremental changes will only lead to inferior performance. That is, internally consistent solutions differ along multiple decisions and therefore are "distant" from one another. Consequently, to discover novel opportunities, firms need to consider changes along multiple choices at the same time. This type of search effort, termed exploration, is well-documented in the organizational learning literature and found to be crucial for long-term performance and survival (He and Wong 2004, Katila and Ahuja 2002, Levinthal 1997, March 1991).

Interestingly, while firms may explore and identify distant solutions, the actual decision-making for the implementation choices is often incremental, sequentially, and happening over an extended time horizon (Ford and Ford 1994, Huy 2001, Nadler and Tushman 1989, Yi et al. 2016). For example, prior literature has stressed that the actual implementation of planned strategic reorientation often takes multiple years to implement (Lamont et al. 1994, Nadler and Tushman 1989, Chandler 1962). The choice of how and which decisions to implement at a time often emerges incrementally over time as a trial and error process itself and may even lead to a deviation or abortion of explored goals(Greenwood and Hinings 1988, Nadler and Tushman 1989, Yi et al. 2016). Simon (1947) notes that that organizational decision-making, while guided by a higher order goal or aspiration-level, is nevertheless boundedly rational and focused on the feedback and evaluation available to decisions taken today (Simon 1967). In characterizing exploration and exploitation, March (1991) stresses that exploration relates to the discovery of new possibility with longer time horizons and local search (exploitation) to the choices of implementation and execution. March stresses that exploration without local efforts will lack to build the required competencies necessary to successfully implement discovered opportunities. That is, the balance of exploration and exploitation that March emphasizes, is a balance of discovery of novel solutions realized via an intertemporal series of local efforts. The intertemporal distinction of exploration and exploitation may be understood as firms as long-term goal setting and local (or myopic) search for implementation.

Previous literature has paid little attention to studying the relationship between the exploration of long-term goals and the myopic search for implementation that follows. A particular void remains with respect to the role of *temporal* myopia in pursuing exploratory goals. Levinthal and March (1993) describe temporal myopia as a firm's tendency to disregard long-term consequences of today's decisions. This can mean that in the search for goal implementation, firms may discount decisions' implication with respect to the long-term goals. An observation that has been associated more broadly with corporate short-termism, which has been argued to lead to decision making that renders goal achievement unlikely in favor of short-term opportunities (Laverty 1996). The short-termism literature typically considers one-shot decisions with difficult to reverse consequences (Ghemawat 2016). However, studies in the field of credit assignment, familiar to students of the computer sciences as well as cognitive psychology (Fu and Anderson 2006, Sutton and Barto 1998), provide evidence that temporal discounting plays are more complex role in learning and navigating between decisions that are spaced out over time (Denrell et al. 2004, Holland 1995, Rahmandad 2008, Sutton and Barto 1998). This literature assumes a series of actions that play out over time. A widely cited example is playing the game of checkers, where a move taken at one point may be crucial for success (winning the game) several moves later (Samuel 1959, 1967). Research in this field has shown that discounting a moves chances of future success is crucial in assigning differences in temporal value (a positioning value) to the choices a player faces While discounting reduces the absolute value of the long-term that is attributed to a decision, it also contrasts the temporal distance to a goal (e.g., winning the game) due to a decision, i.e., whether a decision shortens or lengthens the path to the firm's goal. For example, a move that will reduce the time to win the game of checkers within the next two instead of three moves, will be valued stronger than a move that reduces the time-to possibly win from five to four moves. This effect of "temporal differencing" has been shown to be related to the level of discounting and crucial in propagating future value back to immediate decisions in form of an attributed positional value (Denrell et al. 2004). Greater temporal discounting of distant values emphasizes temporal changes in decisions taken today and can lead to different decision-making, accordingly.

The relative dearth in temporal myopia research and the conflicting arguments of the role of discounting in decision-making raise some important research questions. For example, how does temporal myopia affect choices that will bring the firm closer to is long-term goal? Is temporal myopia affecting outcomes differently for different time horizons of firms' long-term goals? And how are these relationships possibly different given the complexity of the problem to solve? To answer these questions, we investigate the conceptual relationship between myopic decision making and achieving different types of long-term goals. Specifically, we develop a formal model within which we investigate the long-term performance consequences of firms' level of temporal myopia, i.e. the extent to which they discount a decision's contribution to a long-term goal. To do so, we include firms' exploratory search for long-term goals, i.e., aspired solutions that will require a series of decisions to be implement within a given planning horizon. Lastly, we explore how the complexity of the task environment, i.e., the extent to which decisions' value depends on other decisions made, influences the effect of temporal myopia on long-term performance.

Our analysis reveals that discounting temporally distant goals is in fact a necessary condition in order to achieve long-term goals in complex decision problems. More precisely, discounting serves as a mechanism of propagating future value back to individual precursor decisions (Denrell et al. 2004, Fu and Anderson 2006). More importantly, we investigate the mechanism and consequential decision-making dynamics that evolve from it. Our model surfaces that temporal myopia determines firm's temporal differencing for decision making. That is, temporal differencing may be understood as a firm's sensitivity to changes in the temporal distance between the status quo and a desired long-term goal. A more temporally myopic decision maker is more sensitive in evaluating a decision change that shortens the temporal distance than a less myopic decision maker.

These differences in the sensitivity to decisions that affect the temporal distance to a goal influence how likely decision makers take two different, yet related types of decisions that are crucial in accomplishing long-term goals in complex environments. First, temporal myopia can help identifying and taking 'stepping-stone' decisions that have a lower immediate return but shorten the temporal distance to a firm's long-term goal. Second, we find that temporal myopia also helps identifying which highly attractive short-term decisions to reject preemptively to avoid getting trapped by a mediocre alternative along the way. Preemptive reject decisions speed up how quickly a firm can achieve its long-term goal, but can be substituted by stepping stone decisions, which will increase a firm's time to solution substantially, however. Stepping-stone decisions, in contrast, are essential in accomplishing long-term goals and cannot be substituted by preemptive rejects. In fact, not engaging in stepping stone decisions but reject decisions can lead to getting stuck before achieving an even mediocre local peak.

2. COMPLEX ORGANIZATIONS AND TEMPORAL EXPLORATION

2.1. Complex Organizations and Bounded Rational Search

A key motivation of the behavioral theory tradition stems from better understanding how organizations make decisions (Cyert and March 1963, March and Simon 1958). Challenging a traditional neoclassic economics view, the BToF emphasizes that organizations face a large number of choices for which internal, organization-idiosyncratic processes influence the decision-making. The choices faced by the organization are typically rich in interdependencies, that is, one decision taken will affect how useful another decision will or will not be. For example, the value of how much to invest in training of a firm's sales force may depend on the decision of how much product variety will be offered, and vice versa. This complexity in finding a set of "good" decisions, which does not simply allow to optimize single decisions but requires to balance trade-offs among choices, which constitutes the core of organizational decision-making. Simon's seminal work on bounded rational decision-making addresses the challenges that organizations face with this. Accordingly, decision-makers are limited in the number of decision alternatives they can consider at a given point in time. Further, firms are most likely to make incremental improvements to the way they currently solve a problem (Levinthal and March 1993, Simon 1987). When firms are limited to incremental adjustments in their search for better solutions and choices are highly complex (i.e., interact with many other choices), firms will end-up doing more of previously successful practices and less likely deviate from currently coherent solutions.

2.2. Exploratory search as long-term goal setting

The BToF emphasizes that truly novel and potentially better solutions may exist but are distant to the current solution. Distant solutions require a broader search scope, an organizational effort termed as exploration (March 1991). Exploration happens when firms consider changes to multiple of its decisions in place. Indeed, empirical studies have observed firms exploratory search outcomes. For example, Rosenkopf and Nerkar (2001) captured firm exploration as patent activity that would span technological and organizational knowledge domains. Katila and Ahuja (2002) studied firms exploratory search activities as the extent to which issued patents drew upon previously established firm knowledge or new domains of knowledge. He and Wong (2004) using survey data asked firms to which extent they entered new markets and introduced new products to capture their exploration strategy over the past three years. All these studies have in common that they observe the exploratory outcome rather than the exploratory search and implementation. The underlying search process that implements the new activities is difficult to observe empirically and thus has received little attention in prior studies. Even formal models have paid little attention to the search underlying the implementation of exploratory efforts, somewhat implying that once a firm has discovered a distant solution, it may implement it somehow quickly and synchronously. A notable exception is Yi et al. (2016) who have studied and modelled delayed implementation of previously made decisions. They find and theorize that delayed implementation can further advance exploratory efforts because premature lock-ins are possibly mitigated.

2.2.1. Temporal exploration horizon

While synchronous implementation of many decisions is desirable and arguably necessary, the boundedly rational lens underlying the BToF suggests that this is hardly feasible. Empirical evidence in the field of strategic reorientation supports this intuition and finds that such exploratory goals often are followed by long temporal horizons of implementation efforts (Chandler 1962, Nadler and Tushman 1989), and even quick responses are likely to take two to three years (Tushman and Romanelli 1985, Tushman and Rosenkopf 1996). The temporal process of the implementation of exploratory goals is boundedly ration, i.e., incremental rather than system-wide, sequential rather than synchronous, and focused on available feedback rather than uncertain outcomes (Cyert and March 1963). This extends to the BToF's conceptualization of exploratory search. March (1991) stresses that exploration comes with longer time horizons and greater uncertainty of whether exploratory targets can be accomplished due to a lack of immediate feedback. His formal model of learning that specifies exploration as the rate of local learning from the organizational knowledge stock, where slower learning will mitigate premature assimilation of believe systems. Accordingly, the discovery of exploratory solutions may form a new targeted state of the organization but its path to implementation will still be subject to its boundedly rational internal processes. Consequently, exploratory search ought to be considered as a plan, i.e., a temporal goal of desired implementation efforts. That is, discovered opportunities may form a long-term goal informing future decision making.

The broader the exploratory efforts of a firm, the greater the chances to discover novel but very distant opportunities. That is, a broader set of activity choices needs to be changed in order to implement a discovered opportunity. When the implementation is boundedly rational, i.e., sequential and incremental in nature, the breadth with which the firms explores and sets long-term goals consequently determines the temporal horizon for execution. That is, greater exploration resulting in longer horizons (March 1991). How broadly organizations explore and thus allow for longer time horizons for goal setting can depend on a series of factors, such as the organization's planning routines and practices and whether the firm is experience performance declines (Cyert and March 1963). Literature on organizations' time horizon (Reilly et al. 2016) have argued that firms vary with respect to their horizon based on industry investment requirements (Das 1987, Souder and Bromiley 2012), manager's temporal orientation (Das 1987), and established planning routines (Cyert and March 1963, Souder et al. 2016).

2.2.2 Goal implementation as temporally myopic process

While discovered distant solutions may constitute a new long-term goal of the organization, its path to implementation will still be subject to its boundedly rational internal processes. The firm's actual implementation of its long-term goal constitutes in itself a search process. Simon states that goals are "value premises that can serve as inputs to decisions" (Simon 1964:3). That is, a decision may be evaluated in terms of how it contributes to the goals the organization has put forward.

According to Simon (1964), goals are constraints in an organization's immediate local search efforts. Such constraints may include for a decision taken to meet short-term requirements (e.g., profit, survival) (Levinthal and March, 1993) as well as contribute to achieving the firm's long-term goal (Greve 2008). Both information (immediate performance and contribution value to the long-term goal) thus may be evaluated differently by organizations based on their decision-makers' temporal orientation (Das 1987). The trade-off in the evaluation between immediate and long-term value of a decision relates to the concept of temporal myopia (or short-term orientation) and has been described as firms largely ignoring or discounting the future value of a decision (Levinthal and March 1993, Souder et al. 2016).

However, there are few studies known to us in the tradition of the BToF that examine or at least discuss temporal myopia (Levinthal and March 1993, Miller 2002, Vuori and Huy 2016). These few studies somewhat share a conceptual focus on that part of temporal myopic decision making for which managers may in fact know or be able to anticipate long-term consequences but discount (ignore) them in favor of immediate performance gains. Levinthal and March (1993) refer to temporal myopia as organizations *ignoring* the long-term and Miller (2002) as a lack of managerial foresight, i.e., how much decision makers anticipate future consequences. That is, temporal myopia may be defined as the degree to which a decision maker discounts a decision's contribution to the long-term, such as a long-term goal. Hence, temporal myopia constitutes a managerial or organizational evaluation criteria, i.e., how much weight is given to known long-term consequences in the decision-making evaluation process. For example, Vuori and Huy (2016) found in an in-depth case analysis of the cellphone manufacturer Nokia, that managers pushed for new products using an outdated operating system to boost short-term sales instead of committing to a new system, knowing that the old system was not able to compete with other players' offerings over time.

In sum, the BToF may actually suggest that exploration of distant opportunities is goal setting process that is temporally distinct from the actual decision implementation. In that sense, exploration may be understood as a temporal horizon within goals are identified. The implementation, however, remains a boundedly rational search for actions that are incremental and spread out over time. When implementation is boundedly rational, especially when decisions are evaluated temporally myopically, it is unclear how and if such search can lead to successful goal implementation. Because firms decision-making processes are widely considered boundedly rational but exploration as an outcome has been well-document, we will study the relationship between the two more formally.

Insert Table 1 about here

----

3. Model

Environment. In Table 1, we provide key terminology used in this study, which will be discussed in this and the next sections. We follow prior work and conceptualize organizations as systems of interdependent activity choices (Miller 1992, Nelson and Winter 1982, Porter 1996, Rivkin 2000, Siggelkow 2011). Each organization faces choices in its industry for which the set of actual decisions will form its strategy (Porter 1996, Siggelkow 2001, Miller 1992). For example, a commercial airlines needs to decide whether to offer meals on board, whether to implement a hub and spoke route system, enter alliances with other airlines, whether to operate using primary or secondary airports and so forth. These choices are often highly interdependent, i.e., the value of one decision is influenced by other decisions made. For example, the decision to offer meals on board is more valuable if the airlines operates long-haul compared to short-haul routes.

In our model, we specify that each organization faces N distinct choices $s_1, s_2, \dots , s_N$. We further specify that each choice can take one of two states s_i {0,1} where each state refers to a distinct decision. For example, s_1 could refer to whether the organization offers meals on board (s_1=0) or not (s_1=1). Because each of the N choices has two states, an organization faces 2^N possible configurations. Each of the organizational 2^N possible configurations correspond to a unique performance feedback of its environment. Because we intend to model complexity in decisions, i.e., where decisions can influence one another, we utilize Kauffman's NK framework (1993) to specify these performance values. Consequently, we specify that each performance landscape value Vi is a combination of all N choices' individual performance values. The individual performance value $c_i(s_i, s_{-i})$ is a unique value assigned to c_i given its own state (0,1) and the states of its influencing *K* other decisions. For example, the value of the decision to offer limited services on board of an airlines may depend on other choice such as the routes offered (i.e., whether they are long haul or short haul) and the ticket price choice (e.g., high price versus low price). Each distinct contribution value is a randomly assigned value from a uniform distribution, $c_i(s_i, s_{-i}) \sim U[0,1]$. A particular organizational configuration of its N choices is mapped onto the corresponding performance as the average over its N decision's contribution values:

$$\Pi(s) = \frac{1}{N}[c_1(s_1, s_{-1}) + c_2(s_2, s_{-2}) + \cdots + c_N(s_N, s_{-N})]. \qquad \text{Equation (1)}$$

Prior studies have used the analogy of a performance landscape of hills and valleys, where the height of a hill is analogous to the height of performance – the higher the hill one climbs up, the better the performance (Levinthal 1997, Siggelkow 2001). 'Hill-climbing' has a long tradition in the behavioral literature (Holland 1975, March and Simon 1958). The similarity between organizational decisions can be interpreted as the distance between two locations on the landscape. Any configuration that is only different in one decision is termed to be within the local neighborhood of the organization's configuration. Thus, on a hill, there are no local improvements (any one-decision change will lead to a descent from the hill). The topology of such a landscape has been shown to be determined by the complexity of choices, that is, how rich the interactions are between decisions. In a landscape where each decision is independent from all other decisions, there is one optimum, a global peak. Such landscapes have been described as smooth because they gradually lead up to the global optimum (any one decision improvement improves the overall performance because decisions are independent).

Modelling boundedly rational search and long-term goals. Following previous formal modelling approaches in the tradition of the BToF, we specify that a boundedly rational actor is limited in the scope of decision changes she can explore at a given point in time. Specifically, in our model, firms randomly pick one decision each time period and assess the performance implications of changing this decision (e.g., from 0 to 1, or vice versa) (Levinthal 1997, Nelson and Winter 1982). Two additional parameters are introduced in our model to investigate the role of temporal myopia and long-term goals. First, we specify firms' temporal horizon within which they specify their long-term goal. Second, we specify firms' degree of temporal myopia with which they take a decision's contribution to achieving the long-term goal into consideration. We elaborate on both parameters below.

Temporal exploration horizon and long-term goals. The temporal planning horizon of a firm is here defined as a search radius with respect to temporal constraints. A firm will identify a long-term goal (configuration of its choices) it may achieve from its current configuration within a specified timeframe λ. Because firms are boundedly rational and are limited to one decision change per time period, the temporal horizon equates to the maximum number of decisions that can differ between the organization's current configuration and its long-term goal. More specifically, at the beginning of each time period, the highest performing value $\Pi_{l(t)}$ within a λ step distance (the firm's temporal horizon) to its current configuration constitutes the bases for the long-term positional value. To illustrate this, let us assume an organization currently has the decision configuration 00000000000000. Let us further assume this firm has a temporal horizon of λ =3, that is, the firm discovers the performance potential of configurations that are within reach of three time periods of decision-making. That is, three or less changes away. Examples of such configurations would include for example 000000000111, 101010000000, who are three changes away, and 000010000100, who is two changes away. The firm assigns the highest performance available within its temporal horizon as its long-term goal, if the performance is higher than the performance of its current configuration. The search for long-term goals is thus a greedy variant of previous studies' exploratory search, using a search radius (Siggelkow and Rivkin 2006)[1]. It is important to note that assigning a long-term goal does not change the firm's current configuration. It merely constitutes the basis as an identified goal that will further inform decision-making. The actual implementation of decisions will be discussed in the next section.

Temporal myopia. We follow prior research in the tradition of the BToF and model spatial search as incremental neighborhood search (Levinthal 1997, March and Simon 1958, Nelson and Winter 1982). Each time period a firm randomly picks one decision and considers a state change in this decision (from 0 to 1 or 1 to 0). The newly discovered decision and the firm's current configuration both are evaluated using the following specification and whichever value is larger will be adopted by the firm. $Q(s) = \Pi_s + \ (\Pi_{l(t)} * \gamma^{d_t})$, with $\gamma = 1 - m$.

[1] The greedy search of long-term goal is no necessary conditions for the implications of our model's results. A non-greedy goal-setting search will lower the absolute performance achieved but qualitatively lead to same results. More details available from the authors.

Q(s) is the firm's evaluation of configuration s, where Π_s constitutes the immediate performance of policy configuration s, $\Pi_{l(t)}$ constitutes the performance of the long-term goal selected at period *t.* The term γ constitutes how much a firm discounts the performance of the long-term goal raised to d_t. In line with prior research (Denrell et al. 2004, Sutton and Barto 1998), γ is a value between zero and one, where small values imply high discounting and high values imply little discounting. To explore the role of temporal myopia m, as the extent to which a firm discounts, we can write γ as equal to (1-m). While this leaves the equation unaffected, it allows us to describe temporal myopia (m) as low for small values and high for high values (e.g. 1):

$$Q(s) = \Pi_s + \left(\Pi_{l(t)} * (1-m)^{d_t}\right) \qquad \text{Equation (2).}$$

The term *d* is the minimum number of time periods it would take the firm to get implement all decisions needed to arrive at its long-term goal. Because a firm in our model makes one decision per time period, d equates to the hamming distance between configuration s and the long-term goal configuration. For example, if the current configuration *s*=[000000000000] and the long-term goal configuration *l*=[110000000000] then *d*=2. (Because it would at least take two time periods to implement this configuration). This specification is adapted from Sutton and Barton's (1998) credit assignment algorithm and represents a decisions positional value (i.e., how much does the decision contribute to achieving the long-term solution). Because our model does not specify how belief systems are formed over time, but instead how immediate feedback interacts with long-term positional value, we subsume the positional value in $(\Pi_{l(t)} * (1-m)^{d_t})$ as the objective future solution value of state Q(*s)* (cf. a related discussion in Denrell et al. (2004: 1370)[2]. That is, in our model of learning, when

[2] In that sense, our model equates to an objective ex-ante credit-assignment model (the objective positional values are given) where the agent expects any non-solution state in the future (one or

evaluating a focal decision change, the firm will consider the decision's immediate performance consequences (Π_s) and how this decision will change the temporal distance to the long-term goal (d). Thus, our model is in line with BToF arguments that a firm makes decisions based on short-term feedback (Cyert & March 1963), which here includes a signal about the temporal distance change to a long-term goal.

We provide a numerical illustration of the decision-making algorithm in Appendix 4.

4. RESULTS

We present our main results along the two dimensions of temporal decision making, (1) the extent to which a firm's are temporally myopic (m) and consequently the extent to which they discount the value of potential long-term solution and (2) the firm's temporal horizon within which the firm sets long-term goals. Unless specified otherwise, the parameter specifications for our core model are for the number of organizational choices, N=12, complexity, K=11 (high complexity), temporal horizon, λ=3 and temporal myopia m=[0,1] in increments of 0.1. We created 1,000 firms and distinct performance landscapes and report the respective means. Each firm was given up to 300 time periods for search and their performance was tracked until they reached a sticking point, that is, when the firm found a configuration from which it will not move to any other configuration. It is important to note that a sticking point may but must not coincide with a local peak. A sticking point is organization-specific and based on the evaluation rules (here discount factor and time horizon) applied by this firm. That is, a sticking point constitutes

more steps away) to yield no immediate reward. Thus, the information about the immediate reward for any state becomes only available to the firm when searching the actual decision space.

the steady-state of a firm, from where no additional changes in configuration will follow[3]. (See Rivkin and Siggelkow 2002, Siggelkow and Rivkin 2006 for an elaborate discussion of sticking points). The full end of simulation results (steady-state) are listed in Appendix 1.

Insert Figure 1 & 7 about here

4.1. Result summary

Figure 1 plots the average performance at the end of the simulation for firms of different temporal horizon against the firm's degree of temporal myopia. Longer temporal planning horizons perform worse on average than shorter horizons, except for very short planning horizons {λ=1}. Exhibiting some degree of temporal myopia (independent of horizon) beats the extremes no myopia (m=0) and full myopia (m=1). The larger the temporal horizon, the greater the effect of temporal myopia, i.e., small differences in temporal myopia will have stronger performance implications (as shown by the steeper inverted U-shaped curve for larger time horizons). We further find that in all scenarios, a moderately high degree of temporal myopia (m=0.4) to perform the highest albeit very high temporal myopia (m <0.5) perform increasingly worse with increasing temporal horizon. The end performance follows a quasi-inverted U-shape distribution. The highest overall performance is accomplished by firms with a moderate temporal horizon {λ=3} and temporal myopia of *m*=0.4. Complete disregard of long-term potential (*m*=1) as well as full accountancy of long-term potential (*m*=0) perform the worst and equally so.

[3] For a very small fraction of firms (0.2% of all firm experiments) a steady-state could not be reached because of a rare oscillation between two long-term goals. We considered these firms to have reached their steady state when they started the oscillation cycle for the first time.

4.2 Discounting as a mechanism of temporal differencing (sensitivity to temporal differences)

How does the degree of temporal myopia influence a firm's decisions taken during their search for long-term goals? Because all firms face the same performance landscape and have the same temporal horizon (λ=3) in our main model, the only parameter varying across firms is the degree of temporal myopia, i.e., the extent to which firms discount ((1-m) in equ. 2). That is, the degree of discounting must lead to systematically making different decisions in similar situations. In the following, we will investigate how temporal discounting affects decision-making in our model.

Insert Figure 2 about here

As we will explain below, a firm's temporal discounting directly influences the relative value-change between decisions of different temporal distance to a firm's goal. In Figure 2, we plot the effect of temporal differencing, i.e., the net change in the attributed long-term gain $(1-m)^{d-1} - (1-m)^{d}$ against the initial distance (d) to the long-term solution for a series of m values. While the absolute attributed value changes disproportionally with temporal distance as shown in Figure 2, the relative value-change between two temporal distances is constant for a given discount factor and can be calculated as the inverse of (1-m). For instance, a discount of m=0.7 has a value change of 3.3 with each step that brings the firm closer to the goal[4]. In contrast, the relative value for m=0.1 firms changes by a factor of 1.1. Consequently, the

[4] The relative positional value change for a given discount m (e.g., 0.7) follows $\frac{1}{(1-m)} = \frac{1}{(1-0.7)} = 3.3$.

attributed absolute value is mostly attributed to temporally close decisions and very little (in absolute terms) to more distance decisions for m=0.7 firms (high temporal myopia, cf. red line in Figure2). In contrast, firms with m=0.1 (low temporal myopia) attribute only slightly higher value to decisions that are temporally close and generally higher absolute value to more distant decisions (cf. orange line in Figure 2) compared to more myopic discount factor firms. Accordingly, the more temporally myopic a firm, ***the higher the relative value change*** and ***the lower the absolute attributed long-term value*** for a given decision. That is, there are systematic differences in how a positional value is attributed to a decision with regard to its distance leading to a long-term goal (Denrell et al. 2001). So how do these mechanisms play out in our model? Let us consider the example of a decision (s') that would bring the firm one step closer (i.e., temporal distance will be *d-1*) to its long-term solution Π_l but would lead to an immediate performance which is lower than its current performance, i.e., Π_s<$\Pi_{s'}$. Then, following from equation (2) the firm will only accept this decision if

$$\frac{\Pi_s-\Pi_{s'}}{\Pi_l} < (1-m)^{d-1} - (1-m)^d; with\ 0 \le m \le 1. \qquad \text{Equation (3)}$$

That is, the performance change as percentage of the long-term target must be smaller than the increase of the net discount term (the change in discounting due to moving temporally closer), which we refer to as temporal difference of decisions s and s$'$ in reference to the long-term goal. Consequently, whether a firm trades-off a long-term gain for a missed short-term gain depends on the difference in immediate performance together with the attributed temporal difference, i.e., the change in net discount due to a change in temporal distance to the long-term target. The greater the change in net discount term, the greater the value a firm is able to forgo in the short-run. In other words, temporal differencing reflects the attributed gain in the long-term solution due to moving closer (or is negative when moving farther away) and is a direct function of temporal distance and discount level.

To better understand how this insight helps explain decision making differences in our model, let us consider a concrete example. A firm with a discount factor m=0.1. Let us further assume that the firm's current performance is 0.7 and chose a long-term goal of a performance of 0.8, which is three decision steps away. Following equation (2) the firm's overall evaluation of the current position is $0.7 + (1-0.1)^3*0.8 = 1.283$. The firm will only consider taking a decision that moves it further away (4 instead of three steps) from its long-term goal if the immediate performance increase is greater than $((1-0.1)^3-(1-0.1)^4)*0.8=0.058$. The firm will take a stepping-stone decision, i.e., move one step closer to the long-term goal while experiencing an immediate performance decline, if the performance decline is smaller than $((1-0.1)^2-(1-0.1)^3)*0.8=0.065$. For the same scenario, a firm with a discount factor of m=0.7 will accept a the decision that lengthen the distance to the goal if the immediate reward increases by more than 0.015; it will accept an immediate decline of up to 0.050, if this decision moves it closer to the goal. Of course, these numbers change with the distances considered. In fact, the maximum a firm is willing to forgo in immediate performance (δ_{max}) can be calculated following from equation (3):

$$\delta_{max} = ((1-m)^{d-1} - (1-m)^{d}) * \Pi_{max} \qquad \text{Equation (4)}$$

Because we use relative performance value, measured against a landscape's global peak, the highest possible performance, Π_{max} in any given landscape equals 1. Consequently δ_{max} equals the net discount from equation 3, which we plotted in Figure 2.

Mechanism summary. Temporal myopia renders firms more sensitive to decisions that alter the temporal distance to their long-term goal. Greater sensitivity means that the change in relative long-term value, which is attributed to a decision that shortens (or lengthens) the goal distance, is larger the greater temporal myopia. The relative value change is constant and independent of time while the absolute value change is not (as can be seen in Figure 2, where absolute change flattens out the greater the distance). That is, a temporally myopic firm will attribute greater relative value-change to a decision that gets it closer to its goal but not necessarily greater absolute value.

Insert Figure 3 about here

4.3. Relationship between discounting and decision-type

We observe that the degree of temporal myopia affects the occurrence of firms taking two distinct types of decisions. One decision type relates to a choice that will shorten the distance to the long-term goal but have negative consequences to the immediate performance. We call this type, "stepping-stone" decisions. Another decision type we observe is a choice that would increase immediate performance but is not taken by the firm because it would further increase the temporal distance to the long-term goal. We call this type, "strategic reject" decisions. We analyze both types below.

4.3.1 Stepping-stone decisions. In a first step, we explore all new decisions a firm takes related to its degree of temporal myopia. Figure 3a shows the average number of decision changes accepted by a firm given its discount factor. This number is further decomposed into the share of decisions that are immediate performance-gains, which we term 'hill-climbing' (in light blue) and those decisions which lead to a performance decline but will bring the firm closer to their long-term target, i.e., 'stepping-stone' decisions shown in dark blue. Two observations can be made from Figure 3. The number of stepping-stone decisions increases with an increase in *m* and then declines again for higher values (i.e., when myopia is stronger). Further, the total number of accepted decision changes follows a similar distribution – medium-high levels of discounting accept the greatest number of decision changes. For example, a firm with m=0.3 accepts on average 4 decision changes of which 1.5 are stepping-stone decisions. A natural question that arises from these results is whether taking a stepping-stone decision is crucial in finding the long-term target. In a rugged landscape, a firm's current position will be close to or at a basin of attraction of a local peak, which on average will not coincide with the targeted long-term solution (assuming the temporal planning horizon is sufficiently large. We will discuss the role of the temporal horizon in section 4.4). In such conditions, the firm would need to bridge the 'valley' between the current basin of attraction and the basin of attraction of the long-term target. Hence, on average, a firm will require to break free from a basin of attraction by taking a decision with instantaneously lower performance (stepping-stone). To understand the long-term performance implications of taking stepping-stone decisions, we plot in Appendix 2 the proportion of performance achieved at the end of the simulation for all firms. This plot is decomposed into the share of firms that achieved a given performance with no, exactly one, or multiple stepping-stone decisions. The plots indicate that stepping-stone decisions are highly associated with achieving superior performance. For example, the distribution plot for m=0.9 firms (very strong discounting) shows quasi-normally distributed performance with the majority of firms taking no stepping-stone decisions. In contrast, firms with a moderate discount factor of m=0.4 achieve a performance heavily skewed to the high-end and those firms closer to the higher end mostly have taken one or multiple stepping-stone decisions.

4.3.2 Strategic rejections. While a stepping-stone decision is a choice where the firm leaves a basin-of attraction of an undesirable local peak to move closer to its target goal, we now explore strategic rejections of decisions that would lead a firm into a basin of attraction in the first place. That is, we investigate in the following how temporal myopia relates to preemptively rejecting early hill climbing. Figure 4 shows the average number of times a firm faces a choice to immediately improve its performance (hill-climbing) given its discount rate decomposed into the number of times such choices are rejected by the firm (in dark blue) and accepted (in light blue). Two trends in this Figure can be observed. The relative share of performance-positive choices that are rejected happens in a medium range of discount factors (m ~ 0.3 to 0.5). These firms also end up facing a greater number of performance-positive choices. For example, a firm with m=0.4 finds on average 8.1 performance-positive decisions of which it rejects 5.6.

Insert Figure 4 & 5 about here

4.3.3. Counterfactual (mechanism) analysis for stepping-stone and reject decisions. To better understand the role that stepping-stone and reject decisions play in the search for long-term goals, we ran two counterfactual analyses. One with the full model specifications but removed the ability of firms to take reject decisions, and one with the full model specifications but removed the ability of firms to take stepping-stone decisions. Figure 5 shows the end of simulation performance for temporal planning horizon=3 firms for the full, no-reject decisions, and the no stepping-stone decisions models. Either counter-factual firm performs worse than in our main model (except for m=0.1 firms). However, the counterfactual without reject-decisions performs only slightly worse than the main model. In contrast, the counterfactual without stepping-stone decisions performs substantially now performs as low as the benchmark, i.e., what a firm would achieve that does not take the long-term in consideration at all (m=1).

What explains these stark performance differences for stepping-stone and reject decisions? Table 2 shows end of simulation statistics for m=0.4 and λ=3 firms for the full model (column 1), the no-reject decisions model (column 2), and the no stepping-stone decisions model (column 3). The performance of the no-reject decisions model is quite similar to the full model. However, that does not mean that reject decisions are of no relevance. The first line in Table 2 shows that without the ability to preemptively reject early hill-climbing decisions, the firm's time to steady-state (i.e., the firm will make no new decisions) doubles (from 26.6 to 52.1 time periods). When firms cannot reject premature hill-climbing, they will climb up mediocre hills more frequently, and move away from their long-term goal. However, such firms still achieve an only slightly lower performance than their counter-parts in the full model because they will end up taking a greater number of stepping-stone decisions. The number of stepping-stone decisions increases from 1.6 (full model) to 4.7 (counterfactual model, column 2). Because of the lack of reject decisions the firm ends up climbing up and down mediocre hills in an effort to reach its long-term goal. This also increases the overall number of decision changes from 4.0 to 12.6. While traveling a longer path the firm also changes the long-term goal 4.4 times compared to 1.5 in the full mode. The average long-term goal performance declines because the initial goal more likely falls outside the firm's temporal planning horizon when it moves toward mediocre performance neighborhoods and away from previous goals. It is also worth noting that a lack of decision rejections decreases the probability of a firm to reach the long-term goal altogether (10 percentage points fewer firms reach their goal) because the risk of getting stuck on a local peak or sticking point increases with early hill-climbing.

Insert Table 2 about here

In the second counterfactual model (column 3), firms are not able to take stepping-stone decisions. The time to steady state cuts into less than half (10.8 time periods) but with a very poor performance achieved (benchmark performance of a fully myopic firm m=1). This suggests that stepping-stone decisions are a necessary for firms that take long-term decisions into account and cannot be substituted by reject decisions. In fact, the number of reject decisions declines to an average of 1.7. That is, such firms behave very similar to fully myopic firms (m=1) by hill-climbing early but may not even end-up on a mediocre hill and instead get stuck uphill due to an inefficient set of reject decisions. Only 66 percent of firms are on a local peak at the end of the simulation compared to 92 percent in the full model (statistics not shown here).

Mechanism summary. Strategic reject decisions improve efficiency in accomplishing long-term goals, i.e., they reduce the time to goal, but can (at least in part) be substituted by additional stepping-stone decisions. When stepping-stone decisions substitute strategic reject decisions, however, the time-to-goal needed increases drastically. In contrast, stepping-stone decisions are necessary for achieving long-term goals and cannot be substituted by reject decisions. In fact, not engaging in stepping-stone decisions but strategic reject decisions will lead to inferior performance than not to doing either decision type, because strategic reject decisions can hinder simple hill-climbing.

4.4. The role of temporal horizon

The temporal exploration horizon determines the performance of the long-term goal. The greater the horizon, the greater goals will be identified. This is most obvious in the extreme case of

λ=12, where the goal will always be the global peak (performance 100%). Therefore, the greater the horizon, the smaller the times the long-term goal may be updated while the firm searches over the landscape (this is confirmed in column 8 of Appendix 1). Shorter horizons update their goals more frequently as they accomplish one goal and identify another one nearby. In the extreme case of λ=1, the exploration horizon simply identifies the next optimal hill-climbing decision (equivalent to a greedy local search). With increasing λ (e.g., λ=3), updated long-term goals are distant enough away so that accomplishing one local peak will be followed by goal that requires the firm to bridge a 'valley' to get there but improves overall performance and if accomplished, will navigate firms away from local neighborhoods toward higher long-term goals.

How does the goal distance influence the firm's ability to accomplish these goals and their overall performance? From column 5, we can see that the long-term goal performance is greater the greater the horizon. However, the actual firm performance (column 4) is highest for moderately myopic firms with a rather limited time horizon (λ=3). This can be explained by more firms achieving their long-term goals when their temporal horizon is rather moderate than larger. For example, 61% of the best performing myopic firms (m=0.4) with a horizon of λ=3 achieve their long-term goal, while only 33% of the best performing firms (m=0.3) with a horizon of λ=12 accomplish this.

From Figure 1, we were able to discern that different temporal planning horizons coincided with somewhat different optimal degrees of temporal myopia. Because each temporal distance to a goal is associate with a particular discount factor that will maximize the immediate reward difference that can be forgone in order to move closer to the long-term solution, a suitable discount factor for one distance may be less optimal once the firm moves closer. That should lead to firms getting stuck more distantly to the long-term solution the weaker the discount factor, because they will decline stepping-stone decisions more likely. Figure 6 supports this intuition. In this figure, we plot the proportion of firms given the distance to the long-term solution at the end of the simulation for four different temporal horizons (λ={3,5,8,12}). Temporally myopic firms (larger m) getting stuck more likely at farther distances, which is associated with larger temporal horizons. Less myopic firms (lower m) get stuck at closer distances (five and fewer steps away from the long-term solution), which happens for all levels of temporal horizons. Consequently, less discounting is penalized in all scenarios even when such firms have initially an advantage for very distant temporal horizons (e.g., λ=12) but will get stuck once they move closer.

Mechanism summary. The longer the temporal planning horizon, the greater the chances a firm sets a very distant long-term goal. Such distant goals will render absolute value changes between decisions for temporally myopic firms very small, which will mostly make them reject stepping-stone decisions. In contrast, weaker temporal myopia will still attribute large enough absolute value-changes to decisions that move the firm closer to a very distant long-term goal.

Insert Figure 6 about here

4.5. The role of task complexity

In this section, we analyze how the degree of complexity in the task environment, i.e., the extent to which one choice depends other decisions, influences our findings. In our core model, we specified complexity to be high, i.e., each decision's value is influenced by the configuration of all other decisions. In Appendix 3, we show our results for medium (Panel I) and low complexity

(Panel II). A first observation is that the highest achieved performance increased compared to our core results. Because we measure performance against the global peak, the results indicate that with decreasing complexity, the chances for some firm specifications (along temporal myopia and temporal horizon) to achieve the global peak increase.

Further, with decreasing complexity, we observe that the highest performance achieved is associated with longer temporal horizons than under high complexity (our core model). Panel I shows that for medium complexity, firms with a temporal horizon of $\lambda=5$ now achieve the highest performance followed by horizons of $\lambda=8$ and 12. For low complexity (Panel II), the best performance is associated with temporal horizons of 5 and larger.

To explain these results associated with complexity, it is important to take into consideration the temporal discount as well. For medium complexity, the degree of temporal myopia associated with the highest performance is now lower than in our core results for high complexity.

The landscape becomes more rugged (many local peaks) with complexity, and less rugged or even smooth with one global peak, the lower complexity. In less rugged landscapes, peaks become less steep but have wider basins of attraction (i.e., hill-climbing areas). Because of the less steep nature of local peaks, more distant goals bear the potential to offset the immediate "gradual" decline. However, this is more likely when discount rates are weaker, so that the absolute change in long-term value is large enough. This effect can be observed in a relative increase in stepping stone decisions and strategic reject decisions in medium complex environments.

In the case of low complexity, the landscape has very few peaks and is rather smooth. That is, in such a landscape, even fully myopic search proves quite effective (m=1 leads to 0.97

performance), because local improvements contribute to a system improvement due to pooled rather than non-linear interactions between decisions taken (Thompson 1967). In such environments, the role of stepping stone and reject decisions is still crucial in accomplishing the long-term goal but the major part of performance accomplished is due to effective local (myopic) search in a smooth landscape. This becomes evident in the fact that the optimal level of discounting is very broad, i.e., for *m* between zero and 0.5. Because of very subtle performance differences between local peaks, a wide range of discount factors is effective in offsetting stepping-stone (or reject) decisions immediate but gradual performance declines in the otherwise very smooth landscape. Larger temporal planning horizon are needed to find a higher peak because in low complex environments, basins of attraction are relatively wide and distinct peaks are relatively far away.

Insert Figure 7 about here

In Figure 7 we distill the key results of this paper given the different complexity levels in a more effective way. Panel I (of Figure 7) shows the level of temporal myopia (y-axis) with which a firm accomplishes the highest performance given a particular temporal planning horizon (x-axis). The optimal level of temporal myopia is rather high for short to medium planning horizons and lower for longer horizons. However, the level of the optimal degree of temporal myopia flattens out around 0.2, which suggests that less myopia than that is not desirable in any of the scenarios in our study. For medium complex environments less discounting for moderately long planning horizons becomes more important than for high complex environments. For low complex environments, higher levels of temporal myopia are optimal for all horizon comparisons. Panel II shows the additional performance (y-axis) that a firm, with the optimal level of temporal myopia for a given temporal horizon, can expect compared to a firm that does not take the future into account at all (benchmark). The greatest gains can be expected for moderately short time horizon firms (assuming the optimal level of temporal myopia from Panel I). For low complex environments this finding is also true for longer horizons.

5. DISCUSSION

Temporal discounting as navigator

We have explored how the degree of temporal myopic search, i.e., the extent to which a decision is discounted regarding its contribution to achieving a long-term goal, influences a firm's performance over time. As with most variables one studies, there is not a single, linear relationship. This is what we find for temporal myopia, too. The novel insight of our study, however, is exactly that. The ongoing debate on short-termism and temporal myopia in particular has attributed a generally negative relationship between temporal myopia and long-term performance outcomes. The core argument has been that short-term and long-term decisions are often trade-offs, and one cannot be optimized without forgoing the other one. The argument further goes that discounting future value will reduce the absolute value associated with a long-term decision, so that it becomes unattractive to take. In our model, trade-offs are plenty when the environment is complex, i.e., a decision's value is influenced by many other decisions. We find that in trade-off rich environments, neither a very strong nor a very mild degree of temporal myopia perform well over time. However, we find that in such trade-off rich environments, stronger degrees of temporal myopia, compared to less complex environments, can in fact be helpful in accomplishing better long-term performance. The extent of temporal myopia can be understood as a decision-maker's sensitivity to changes in the temporal distance of future value. This has been termed impatience, impulsive, and short-term oriented, all terms have been used largely in relation to negative or undesirable outcomes, such health risks and a lack of long-term investments (Laverty 1996). However, the sensitivity to changes in temporal distance of future value helps explain when greater discounting may guide firms toward long-term goals. With greater discounting the relative change in attributed long-term value is larger than for milder discounting. Accordingly, decisions that move a firm closer are associated with a greater long-term value increase as compared to mild discounting and consequently will render firms more likely in taking such choices even if there are short-term performance penalties.

On long-term horizon and goal setting

How distant should organizations set their goals? When decision problems are complex in nature, i.e., choices interact with one another, distant, far-away solutions often provide superior performance over local solutions (Cyert and March 1963). One would expect that longer planning horizons therefore allow for higher performance goal. Our model suggests a somewhat puzzling insight that in high complex environments, firms will perform better if their planning horizons are limited and discount rates are high. We do find that greater planning horizons will lead to higher performing goals, but that these goals rarely are accomplished (even when temporal myopia is mild). In our model, firms accomplish goals that are moderately distant and then update toward a new moderately distant goal. Consequently, the long-term search is subdivided into multiple achievable goals over time rather than one very high goal upfront. However, the temporal planning horizon needs to be sufficiently long to set goals that navigate firms away from local neighborhoods. This is less of concern in less complex environments. The reason being that short-term and long-term performance goals align more often and the path toward solutions is less seeded with trade-off choices that would “trap” firms early on. Thus, in simple environments, setting long planning horizons can often be beneficial and efficient.

The importance of choosing “what not to do”

Prior literature on long-term strategies has highlighted the role of stepping-stone decisions, i.e., immediate performance declines for greater long-term returns (Flammer and Bansal 2016, Laverty 1996, Rahmandad 2008). We too find this to constitute a crucial if not necessary decision organizations need to take in order to move closer to their long-term goals. We have shown how temporal myopia may in fact help taking such decisions. However, we find another insightful decision type “strategic rejections”, which has found less attention in the literature. Strategic rejections are decisions that forgo the opportunity to improve performance instantaneously by maintaining the status quo. We find this type of decision to speed-up the time to a targeted goal. Many successful strategies have proactively forgone opportunities of quick returns that would slow down the path to long-term goals. For the marriage-oriented online dating service eHarmoney, the strategic rejection of accepting 20% of paying customers, which it deemed less serious about long-term relationships would forgo substantial immediate profits (Piskorski et al. 2008). But this strategic reject decision brought eHarmoney closer to building a reputation of having a pool of members who are serious about long-term relationships. Porter (1996) has emphasized that strategic decision making requires to choose what not to do. What he calls the growth trap may capture not taking strategic reject decisions in our model. By taking decisions that instantaneously improve revenues and performance, firms may move further away from their long-term goals. Our model shows that proactively rejecting such decisions is crucial. Furthermore, if straying away from the path toward the long-term goal, firms need to compensate by taking more stepping-stone decisions, i.e., they will have to take short-term performance penalties to get back on track. Such effects are described and often observed for firms that entered the growth trap and had to refocus to the core mission.

Conclusion

Temporal myopia, i.e., firms' tendency to discount a decision's consequences for achieving a long-term goal, has been associated with poor long-term performance. In a formal model we show that temporal myopia has not a simple good or bad influence on long-term performance but shows a non-trivial relationship with long-term goal achievement and performance. We have shown that temporal myopia to some degree is necessary to navigate search through complex decision problems as it sensitizes decision makers toward taking decisions that actually move them closer to their long-term goals over time. We hope to trigger a discussion of a new way of looking at discounting, instead of as merely thought as cost of time or capital but as a tool to calibrate sensitivity of what long-term decisions to reject, which ones to pursue in order to effectively navigate to a targeted long-term objective.

REFERENCES

Chandler AD (1962) *Strategy and Structure* (M.I.T. Press, Cambridge, Mass.).

Cyert RM, March JG (1963) *A behavioral theory of the firm* (Blackwell, Oxford).

Das TK (1987) Strategic planning and individual temporal orientation. *Strategic management journal* 8(2):203–209.

Denrell J, Fang C, Levinthal DA (2004) From T-Mazes to labyrinths: Learning from model-based feedback. *Manage. Sci.* 50(10):1366–1378.

Flammer C, Bansal P (2016) Does a Long-Term Orientation Create Value? Evidence from a Regression Discontinuity. *Strat. Mgmt. J.*:n/a-n/a.

Ford JD, Ford LW (1994) LOGICS OF IDENTITY, CONTRADICTION, AND ATTRACTION IN CHANGE. *Academy of Management Review* 19(4):756–785.

Fu WT, Anderson JR (2006) From recurrent choice to skill learning: A reinforcement-learning model. *Journal of experimental psychology: General* 135(2):184.

Gavetti G, Greve HR, Levinthal DA, Ocasio W (2012) The Behavioral Theory of the Firm: Assessment and Prospects. *Academy of Management Annals* 6(1):1–40.

Greenwood R, Hinings CR (1988) ORGANIZATIONAL DESIGN TYPES, TRACKS AND THE DYNAMICS OF STRATEGIC CHANGE. *Organization Studies* 9(3):293–316.

Greve HR (2008) A behavioral theory of firm growth: Sequential attention to size and performance goals. *Academy of Management Journal* 51(3):476–494.

He ZL, Wong PK (2004) Exploration vs. exploitation: An empirical test of the ambidexterity hypothesis. *Organization science* 15(4):481–494.

Holland JH (1975) *Adaptation in Natural and Artificial Systems* (University of Michigan Press, Ann Arbor, MI).

Holland JH (1995) *Hidden order* (Basic Books, New York).

Huy QN (2001) Time, temporal capability, and planned change. *Academy of Management Review* 26(4):601–623.

Katila R, Ahuja G (2002) Something Old, Something New: A Longitudinal Study of Search Behavior and New Product Introduction. *The Academy of Management Journal* 45(6):1183–1194.

Kauffman SA (1993) *The origins of order: Self-organization and selection in evolution* (Oxford University Press, New York).

Laverty KJ (1996) Economic "Short-Termism": The Debate, The Unresolved Issues, and The Implications for Management Practice and Research. *ACAD MANAGE REV* 21(3):825–860.

Levinthal DA (1997) Adaptation on rugged landscapes. *Manage. Sci.* 43(7):934–950.

Levinthal DA, March JG (1993) THE MYOPIA OF LEARNING. *Strategic Management Journal* 14:95–112.

March JG (1991) EXPLORATION AND EXPLOITATION IN ORGANIZATIONAL LEARNING. *Organization Science* 2(1):71–87.

March JG, Simon HA (1958) *Organizations* (Wiley, New York).

Miller D (1992) Environmental Fit Versus Internal Fit. *Organ Sci.* 3(2):159–178.

Miller KD (2002) Knowledge Inventories and Managerial Myopia. *Strategic Management Journal* 23(8):689–706.

Nadler DA, Tushman ML (1989) Organizational frame bending: Principles for managing reorientation. *Academy of Management Perspectives* 3(3):194–204.

Nelson RR, Winter SG (1982) *An Evolutionary Theory of Economic Change* (Harvard University Press, Cambridge, MA).

Piskorski MJ, Halaburda H, Smith T (2008) EHarmony.

Porter ME (1996) What is strategy? *Harv. Bus. Rev.* 74(6):61–78.

Rahmandad H (2008) Effect of delays on complexity of organizational learning. *Management Science* 54(7):1297–1312.

Reilly G, Souder D, Ranucci R (2016) Time Horizon of Investments in the Resource Allocation Process: Review and Framework for Next Steps. *Journal of Management* 42(5):1169–1194.

Rivkin JW (2000) Imitation of complex strategies. *Manage. Sci.* 46(6):824–844.

Rivkin JW, Siggelkow N (2002) Organizational sticking points on NK landscapes. *Complexity* 7(5):31–43.

Rosenkopf L, Nerkar A (2001) Beyond local search: boundary-spanning, exploration, and impact in the optical disk industry. *Strategic Management Journal* 22(4):287–306.

Samuel AL (1959) Some studies in machine learning using the game of checkers. *IBM Journal of research and development* 3(3):210–229.

Samuel AL (1967) Some studies in machine learning using the game of checkers. II—Recent progress. *IBM Journal of research and development* 11(6):601–617.

Siggelkow N (2001) Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne. *Academy of Management Journal* 44(4):838–857.

Siggelkow N (2011) Firms as Systems of Interdependent Choices. *Journal of Management Studies* 48(5):1126–1140.

Siggelkow N, Rivkin JW (2006) When exploration backfires: Unintended consequences of multilevel organizational search. *Academy of Management Journal* 49(4):779–795.

Simon HA (1947) *Administrative behavior: A study of decision-making processes in administrative organization* (Macmillan).

Simon HA (1964) On the concept of organizational goal. *Administrative science quarterly*:1–22.

Simon HA (1987) Bounded rationality. *The New Palgrave: utility and probability*:15–18.

Souder D, Bromiley P (2012) Explaining temporal orientation: Evidence from the durability of firms' capital investments. *Strat. Mgmt. J.* 33(5):550–569.

Souder D, Reilly G, Bromiley P, Mitchell S (2016) A Behavioral Understanding of Investment Horizon and Firm Performance. *Organization Science* 27(5):1202–1218.

Sutton RS, Barto AG (1998) *Reinforcement learning: An introduction* (MIT press Cambridge).

Tushman ML, Romanelli E (1985) Organizational evolution: A metamorphosis model of convergence and reorientation. Cummings LL, Staw BM, eds. *Research in organizational behavior*. (JAI Press, Greenwich, CT), 171–222.

Tushman ML, Rosenkopf L (1996) Executive succession, strategic reorientation and performance growth: A longitudinal study in the US cement industry. *Management Science* 42(7):939–953.

Vuori TO, Huy QN (2016) Distributed Attention and Shared Emotions in the Innovation Process: How Nokia Lost the Smartphone Battle. *Administrative Science Quarterly* 61(1):9–51.

Yi S, Knudsen T, Becker MC (2016) Inertia in Routines: A Hidden Source of Organizational Variation. *Organization Science* 27(3):782–800.

Figure 1. Average performance of firm temporal horizon depending on discount factor.

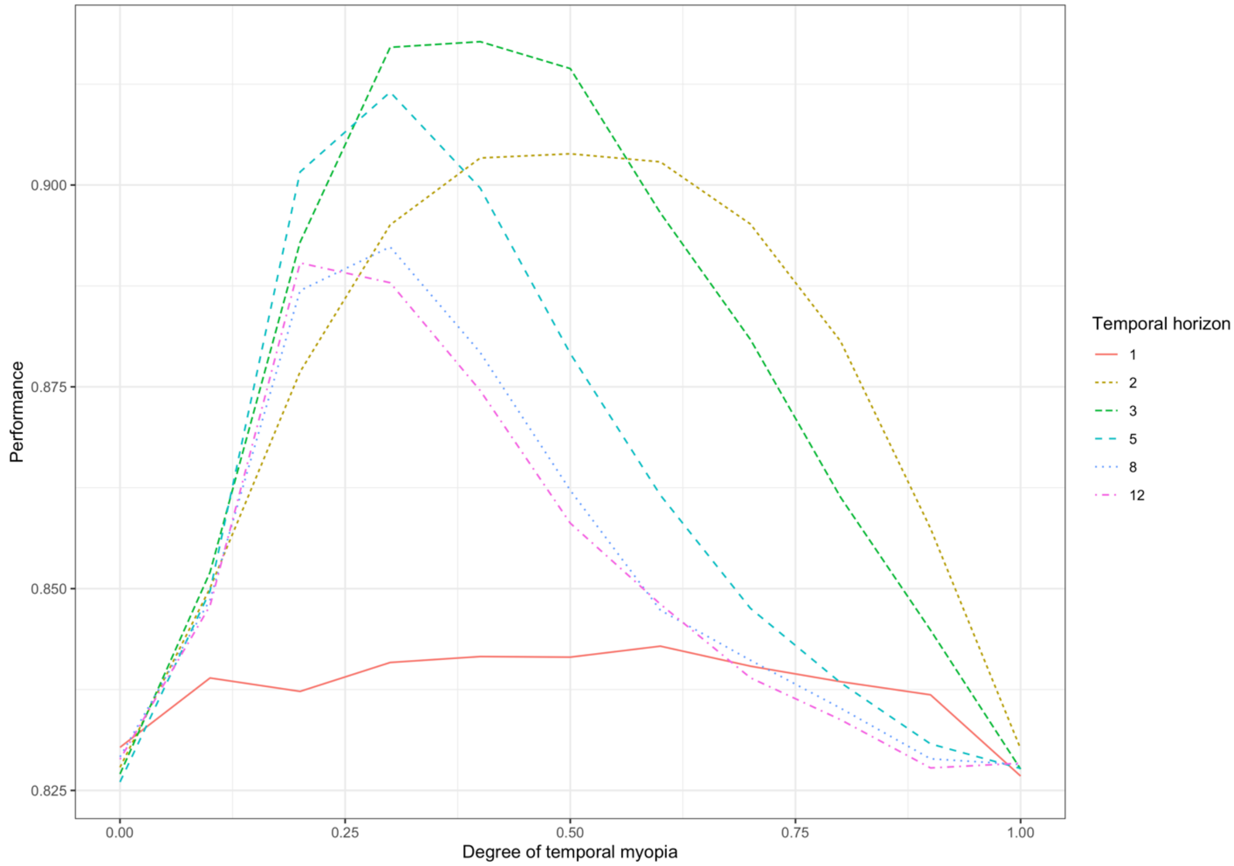

Note. Results are shown when firms have reached a steady state. Complexity level is set to high (K=N-1).

Table 1. Key terminology used in this study.

Construct	**Definition**
Temporal horizon	The maximum timeframe of decision changes considered.
Long-term goal	A set of decisions that can be accomplished within a firm's temporal planning horizon.
Temporal myopia	The degree to which a decision-maker discounts a temporally distant long-term goal.
Positional value	A decision's associated contribution in achieving a firm's long-term goal.
Stepping-stone decision	A choice that leads to an immediate performance decline but shortens the temporal distance to the long-term goal.
Strategic reject decision	A choice that forgoes an immediate performance opportunity in order to not lengthen the temporal distance to the long-term goal.

Figure 2. Discounting and temporal differencing (net discount change).

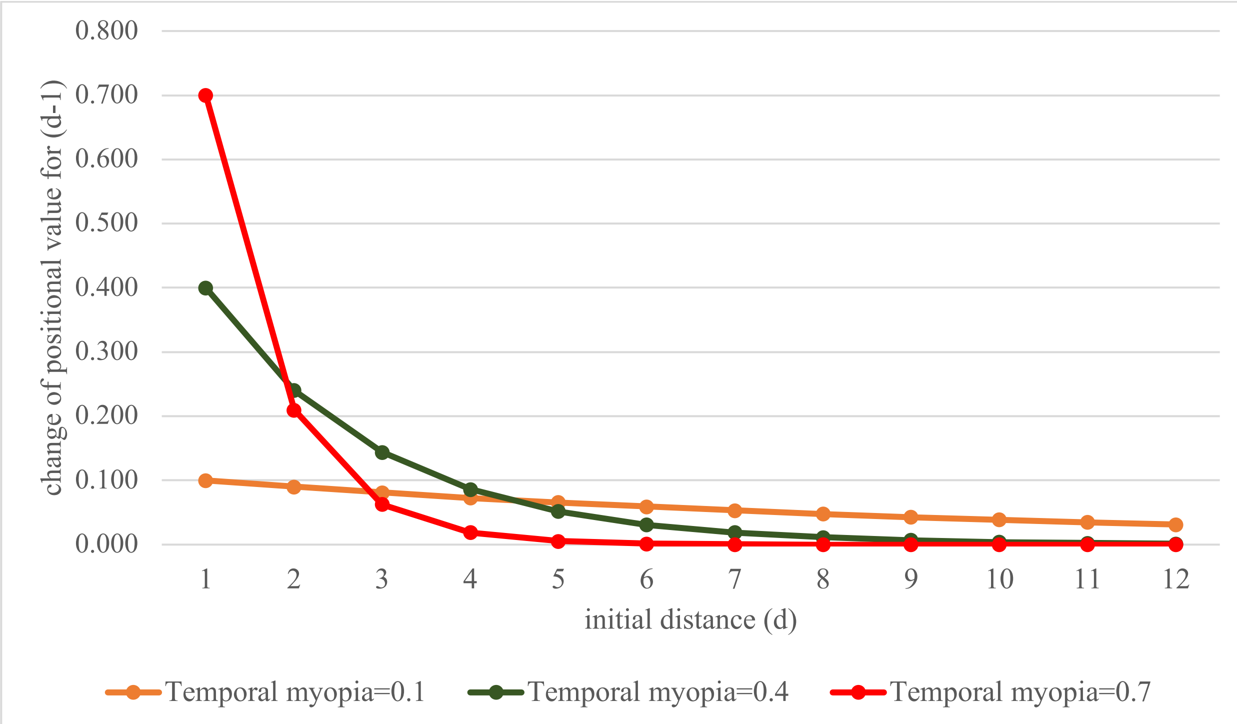

Note. The graph plots the change in the net discount term for a temporal distance reducing step (from d to d-1) following equation (3). For example, due to making a decision that moves a firm with temporal myopia=0.7 (red line) from two steps (*d*=2 on the x-axis) to only one step away from its targeted long-term goal, this firm attributes an additional 0.210 of the long-term value to this decision (y-axis). In contrast a firm with temporal myopia=0.1 attributes an additional 0.090 to the same decision.

Figure 3. Decision changes by degree of temporal myopia.

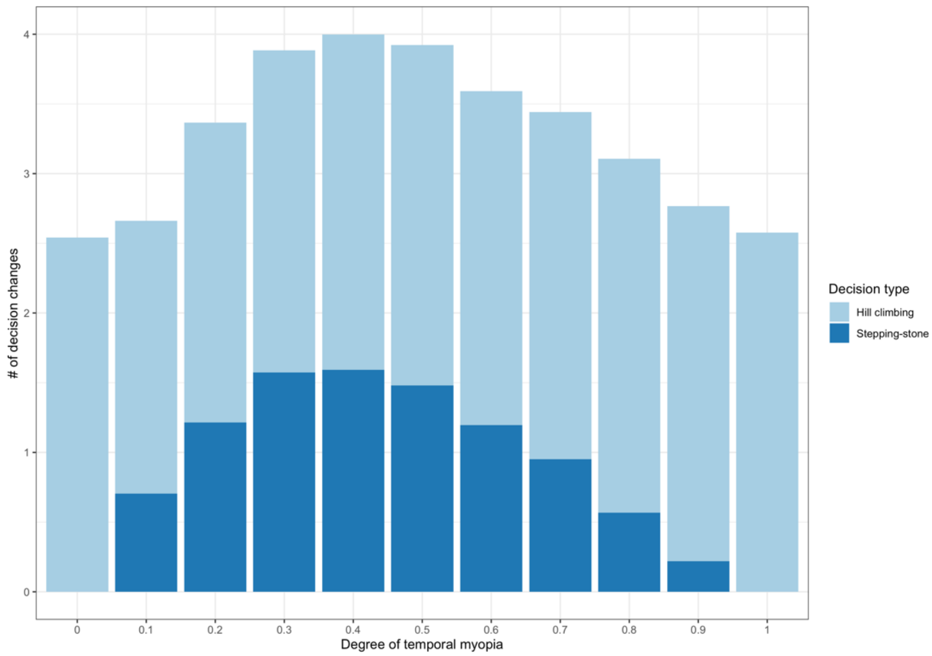

Note. The bar chart shows the average number of decision changes of a firm. This number is decomposed into the share of decisions that lead to immediate performance increase (light blue) and immediate performance decline (dark blue).

Figure 4. Number of choices that immediately improve performance.

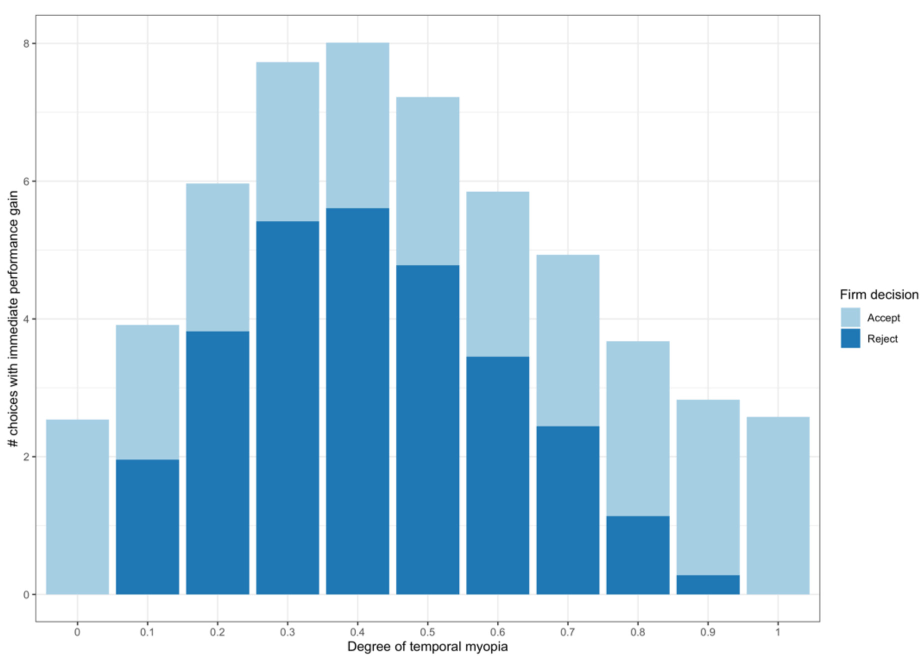

Note. The bar chart shows the average number of choices faced by a firm (until steady-state) that would immediate improve performance. This number is decomposed into the share of these choices that is accepted (light blue) and rejected (dark blue).

Figure 5. Counterfactual performance analysis for stepping-stone and reject decisions.

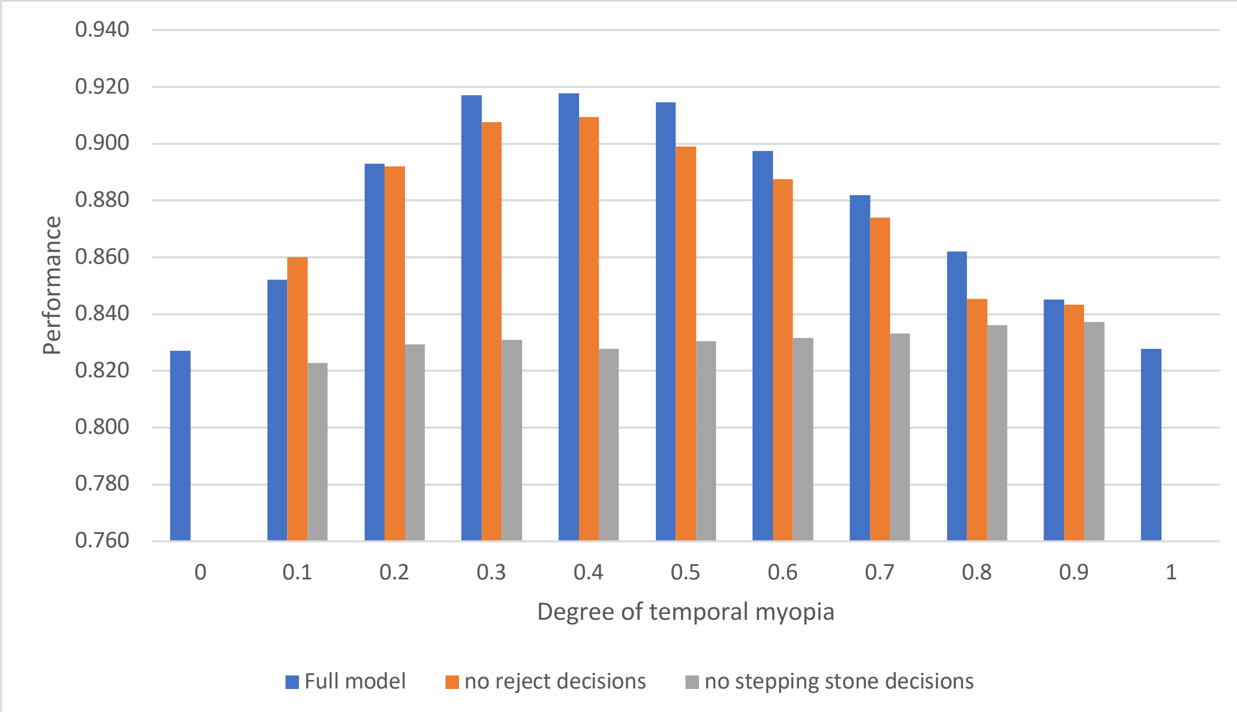

Note. The counterfactuals and the full model are for high complexity environments. Each counterfactual scenario is the average of 300 firms. No counterfactual models are shown for temporal myopia={0,1}, because neither of the two engages in reject or stepping-stone decisions in the full model.

Table 2. Counterfactual analysis for stepping-stone and reject decisions.

	Full model	No reject decisions	No stepping-stone decisions
Time to steady state	26.8	52.1	10.8
Performance	0.918	0.909	0.828
Long-term goal performance	0.951	0.948	0.939
Achieved long-term goal	0.61	0.51	0.17
Distance to long-term goal	1.1	1.40	2.2
# new long-term goals	1.5	4.4	0.8
# decision changes	4.0	12.6	1.5
# stepping-stone decisions	1.6	4.7	
Proportion of stepping-stone decisions	0.40	0.37	
Proportion of hill-climbing decisions	0.60	0.63	1.00
# rejected decisions	5.6		1.7

Note. Results shown for a firm with temporal myopia m=0.4 and λ =3 and high complexity. Each counterfactual scenario is the average of 300 firms. Full statistics available from the authors.

Figure 6. End of simulation distance to long-term goal distribution.

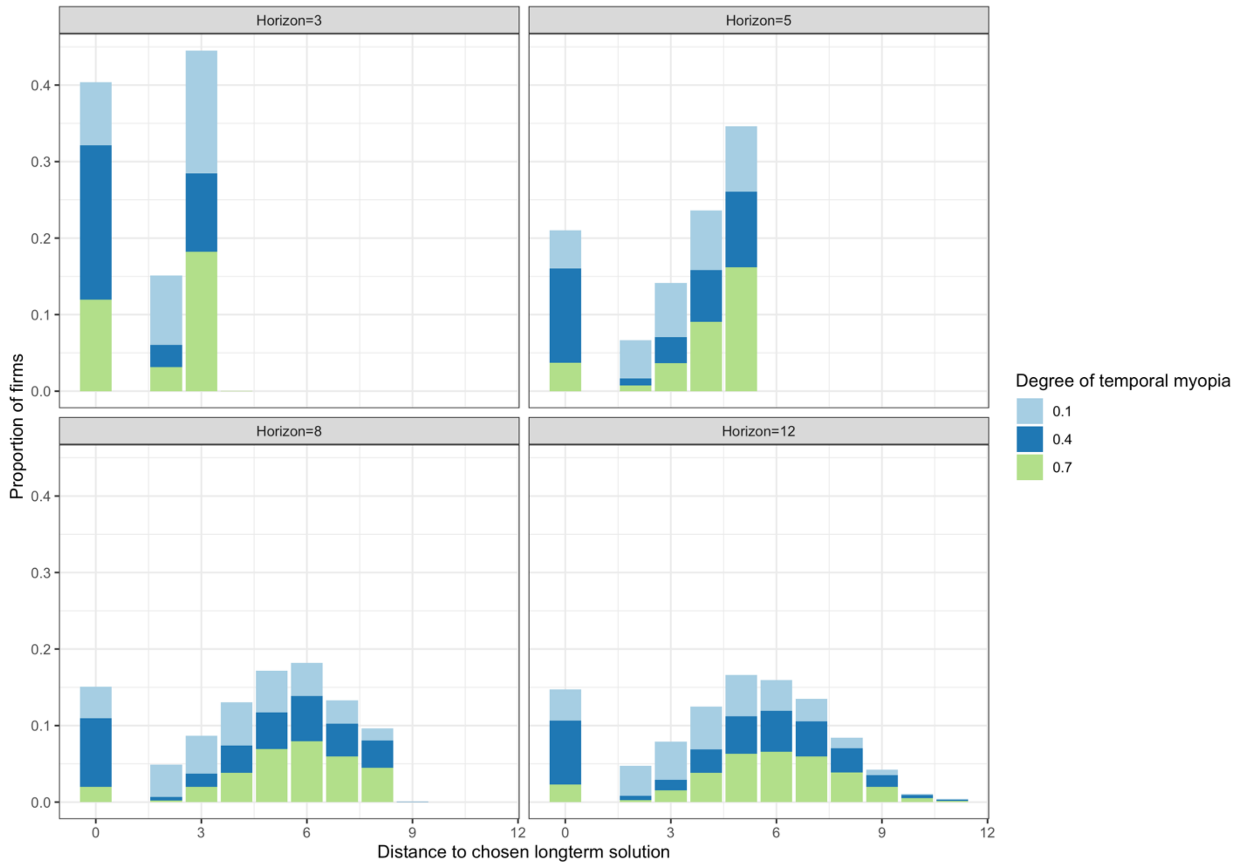

Figure 7. Result summary of optimal temporal myopia, temporal horizon and complexity.

Panel I. Optimal degree of temporal myopia per temporal horizon.

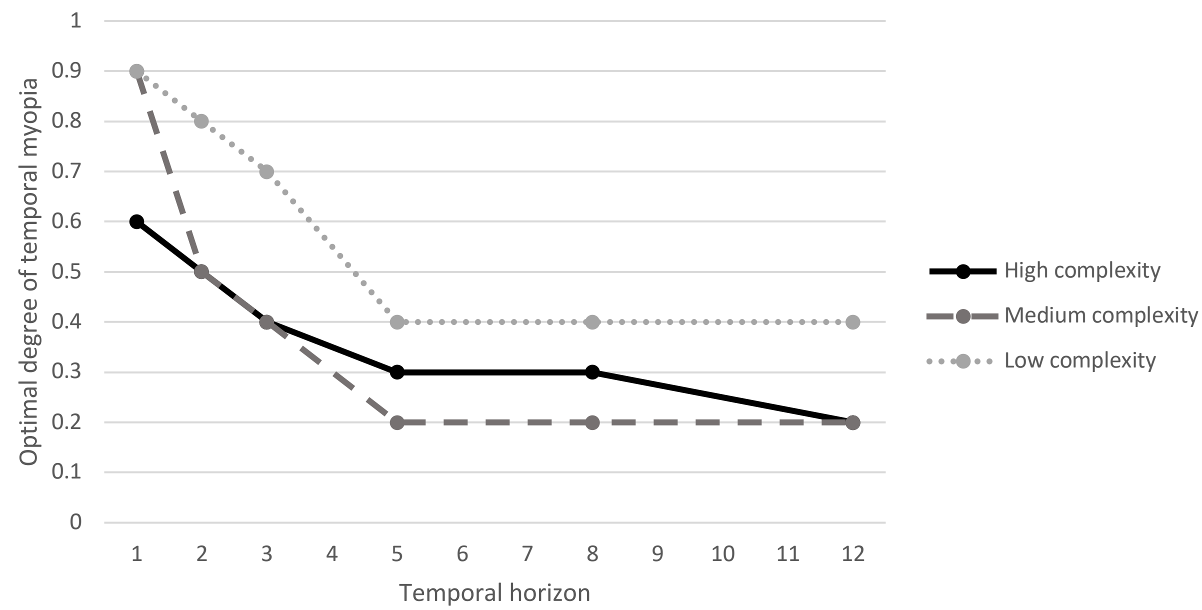

Panel II. Superior performance of optimal temporal myopia over benchmark

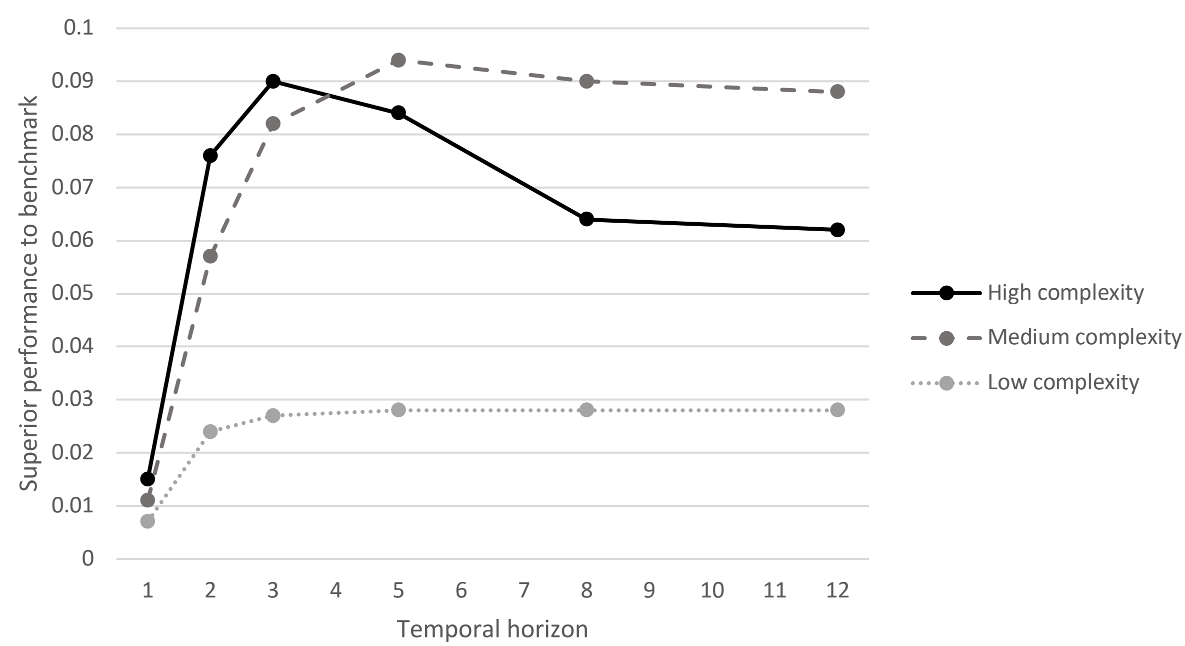

Note. Panel I shows the optimal degree of temporal myopia to achieve the highest performance for a given temporal horizon. If multiple levels of temporal myopia yield the same performance, the higher level is shown. Panel II shows the difference in performance between the optimal level of temporal myopia for a given temporal horizon and the benchmark performance, i.e., when firms do not account for long-term value at all.

Appendix 1. End of simulation results.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Row number	Temporal horizon (λ)	Temporal myopia (m)	Time to steady state	Performance	Long-term goal performance	Achieved long-term goal	Distance to long-term goal	# new long-term goals	# decision changes	# stepping-stone decisions	Proportion of stepping-stone decisions	Proportion of hill-climbing decisions	# rejected decisions	# of periods below best prior performance	Average performance until steady-state
1	1	0.0	15.0	0.830	0.830	1.00	0.0	1.9	2.6	0.0	0.00	1.00	0.0	0.0	0.761
2	1	0.1	16.8	0.839	0.839	1.00	0.0	1.1	2.0	0.0	0.00	1.00	1.7	0.9	0.756
3	1	0.2	17.3	0.837	0.837	1.00	0.0	0.9	1.7	0.0	0.00	1.00	2.8	1.9	0.743
4	1	0.3	17.6	0.841	0.841	1.00	0.0	0.8	1.7	0.0	0.00	1.00	3.3	2.5	0.739
5	1	0.4	18.0	0.842	0.842	1.00	0.0	0.7	1.6	0.0	0.00	1.00	3.6	3.2	0.734
6	1	0.5	18.5	0.842	0.842	1.00	0.0	0.7	1.6	0.0	0.00	1.00	3.9	3.3	0.733
7	1	0.6	18.4	0.843	0.843	1.00	0.0	0.7	1.6	0.0	0.00	1.00	3.8	2.9	0.733
8	1	0.7	17.7	0.840	0.840	1.00	0.0	0.7	1.6	0.0	0.00	1.00	3.3	2.7	0.740
9	1	0.8	17.5	0.838	0.838	1.00	0.0	0.9	1.7	0.0	0.00	1.00	2.6	1.9	0.745
10	1	0.9	17.2	0.837	0.837	1.00	0.0	1.1	2.0	0.0	0.00	1.00	1.8	1.0	0.755
11	1	1.0	14.6	0.827	0.827	1.00	0.0	1.9	2.6	0.0	0.00	1.00	0.0	0.0	0.759
12	2	0.0	14.2	0.828	0.888	0.26	1.5	2.0	2.5	0.0	0.00	1.00	0.0	0.0	0.763
13	2	0.1	14.6	0.850	0.900	0.41	1.2	1.2	2.2	0.4	0.18	0.82	1.8	3.0	0.752
14	2	0.2	19.3	0.877	0.907	0.61	0.8	1.2	2.5	0.7	0.29	0.71	3.6	7.0	0.741
15	2	0.3	23.8	0.895	0.911	0.75	0.5	1.3	3.0	1.0	0.33	0.67	5.4	10.3	0.734
16	2	0.4	24.7	0.903	0.913	0.83	0.3	1.3	3.1	1.0	0.34	0.66	5.7	10.8	0.732
17	2	0.5	25.1	0.904	0.912	0.85	0.3	1.5	3.2	1.1	0.34	0.66	5.6	10.6	0.735
18	2	0.6	25.0	0.903	0.911	0.83	0.3	1.6	3.3	1.1	0.34	0.66	5.2	10.5	0.739
19	2	0.7	23.8	0.895	0.909	0.74	0.5	1.9	3.4	1.1	0.31	0.69	4.1	9.2	0.744
20	2	0.8	21.0	0.881	0.904	0.62	0.8	1.9	3.2	0.8	0.26	0.74	2.4	6.3	0.753
21	2	0.9	19.0	0.857	0.900	0.42	1.2	2.1	3.0	0.5	0.17	0.83	1.0	3.5	0.761
22	2	1.0	14.8	0.830	0.887	0.26	1.5	2.1	2.6	0.0	0.00	1.00	0.0	0.0	0.762
23	3	0.0	14.6	0.827	0.933	0.06	2.6	1.8	2.5	0.0	0.00	1.00	0.0	0.0	0.759
24	3	0.1	17.6	0.852	0.943	0.25	2.0	1.1	2.7	0.7	0.26	0.74	2.0	5.5	0.750
25	3	0.2	22.7	0.893	0.949	0.47	1.4	1.1	3.4	1.2	0.36	0.64	3.8	9.8	0.743
26	3	0.3	26.6	0.917	0.951	0.60	1.1	1.3	3.9	1.6	0.40	0.60	5.4	13.0	0.735
27	3	0.4	26.8	0.918	0.951	0.61	1.1	1.5	4.0	1.6	0.40	0.60	5.6	13.4	0.735
28	3	0.5	24.7	0.915	0.950	0.56	1.2	1.6	3.9	1.5	0.38	0.62	4.8	11.3	0.739
29	3	0.6	22.5	0.898	0.945	0.46	1.5	1.7	3.6	1.2	0.33	0.67	3.5	8.9	0.747
30	3	0.7	21.2	0.882	0.943	0.37	1.8	1.9	3.5	1.0	0.28	0.72	2.5	7.1	0.753
31	3	0.8	19.1	0.862	0.939	0.25	2.1	1.9	3.1	0.6	0.18	0.82	1.2	4.2	0.759
32	3	0.9	16.8	0.845	0.935	0.15	2.4	1.9	2.8	0.2	0.08	0.92	0.3	1.5	0.763
33	3	1.0	14.6	0.828	0.933	0.06	2.6	1.8	2.6	0.0	0.00	1.00	0.0	0.0	0.759

Appendix 1 continued.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Row number	Temporal horizon	Temporal myopia	Time to steady state	Performance	Long-term goal performance	Achieved long-term goal	Distance to long-term goal	# new long-term goals	# decision changes	# stepping-stone decisions	Proportion of stepping-stone decisions	Proportion of hill-climbing decisions	# rejected decisions	# of periods below best prior performance	Average performance until steady-state
34	5	0.0	14.0	0.826	0.977	0.01	4.2	1.1	2.5	0.0	0.00	1.00	0.0	0.0	0.759
35	5	0.1	19.7	0.850	0.985	0.15	3.1	0.7	3.3	1.1	0.33	0.67	2.2	7.4	0.751
36	5	0.2	27.2	0.902	0.988	0.38	2.3	0.9	4.5	1.8	0.40	0.60	4.0	13.0	0.748
37	5	0.3	27.0	0.912	0.987	0.44	2.2	1.0	4.5	1.7	0.39	0.61	4.1	12.9	0.745
38	5	0.4	23.9	0.900	0.985	0.37	2.6	1.0	4.0	1.4	0.36	0.64	3.3	9.7	0.750
39	5	0.5	21.1	0.880	0.983	0.27	3.1	1.1	3.5	1.0	0.28	0.72	2.6	7.4	0.753
40	5	0.6	19.5	0.862	0.980	0.17	3.6	1.2	3.1	0.6	0.19	0.81	1.4	4.7	0.758
41	5	0.7	16.3	0.848	0.978	0.11	3.9	1.2	2.8	0.3	0.10	0.90	0.7	2.1	0.760
42	5	0.8	15.4	0.839	0.979	0.06	4.1	1.1	2.6	0.1	0.05	0.95	0.3	1.0	0.759
43	5	0.9	14.8	0.831	0.977	0.03	4.2	1.2	2.6	0.0	0.02	0.98	0.1	0.4	0.760
44	5	1.0	14.7	0.828	0.977	0.01	4.2	1.2	2.6	0.0	0.00	1.00	0.0	0.0	0.758
45	8	0.0	15.3	0.829	0.998	0.00	5.7	0.2	2.6	0.0	0.00	1.00	0.0	0.0	0.760
46	8	0.1	20.0	0.849	0.999	0.12	4.0	0.1	3.6	1.2	0.34	0.66	1.7	7.5	0.754
47	8	0.2	25.3	0.887	0.999	0.32	3.2	0.2	4.4	1.7	0.38	0.62	3.1	11.2	0.753
48	8	0.3	24.4	0.892	0.999	0.33	3.6	0.2	4.1	1.4	0.34	0.66	3.3	10.5	0.751
49	8	0.4	21.1	0.879	0.999	0.27	4.2	0.3	3.6	1.0	0.28	0.72	2.5	7.2	0.753
50	8	0.5	17.4	0.862	0.998	0.18	4.7	0.2	3.0	0.6	0.19	0.81	1.6	4.2	0.756
51	8	0.6	16.7	0.847	0.999	0.12	5.1	0.2	2.8	0.3	0.11	0.89	0.8	2.2	0.759
52	8	0.7	15.8	0.841	0.999	0.06	5.5	0.2	2.7	0.1	0.05	0.95	0.3	1.0	0.758
53	8	0.8	14.9	0.835	0.998	0.03	5.6	0.2	2.6	0.1	0.03	0.97	0.1	0.6	0.761
54	8	0.9	15.1	0.829	0.998	0.02	5.7	0.2	2.6	0.0	0.01	0.99	0.1	0.1	0.759
55	8	1.0	14.7	0.828	0.998	0.01	5.7	0.2	2.6	0.0	0.00	1.00	0.0	0.0	0.760
56	12	0.0	14.2	0.829	1.000	0.01	5.9	0.0	2.5	0.0	0.00	1.00	0.0	0.0	0.763
57	12	0.1	20.5	0.848	1.000	0.12	4.1	0.0	3.6	1.2	0.33	0.67	1.9	7.4	0.754
58	12	0.2	26.3	0.890	1.000	0.32	3.4	0.0	4.3	1.7	0.38	0.62	3.3	11.9	0.751
59	12	0.3	22.5	0.888	1.000	0.33	3.7	0.0	3.9	1.3	0.34	0.66	3.1	9.6	0.753
60	12	0.4	20.9	0.875	1.000	0.25	4.5	0.0	3.4	0.8	0.24	0.76	2.0	6.3	0.757
61	12	0.5	17.6	0.858	1.000	0.17	5.0	0.0	3.0	0.5	0.18	0.82	1.4	3.9	0.755
62	12	0.6	16.8	0.848	1.000	0.11	5.4	0.0	2.8	0.3	0.11	0.89	0.8	2.0	0.758
63	12	0.7	15.8	0.839	1.000	0.07	5.7	0.0	2.7	0.2	0.06	0.94	0.5	1.2	0.757
64	12	0.8	14.4	0.834	1.000	0.03	5.9	0.0	2.6	0.0	0.01	0.99	0.1	0.3	0.760
65	12	0.9	15.0	0.828	1.000	0.01	5.9	0.0	2.6	0.0	0.01	0.99	0.0	0.1	0.759
66	12	1.0	15.1	0.828	1.000	0.00	6.1	0.0	2.6	0.0	0.00	1.00	0.0	0.0	0.761

Explanation of columns.

Column 1: The maximum number of decision changes considered when choosing a long-term goal.

Column 2: The degree of temporal myopia, i.e., the factor γ with which the firm discounts the long-term goal.

Column 3: The average number of time periods needed to reach the steady-state, i.e., the firm will not accept any neighboring position over its current position.

Column 4: The firm's performance achieved by the end of the simulation.

Column 5: The performance of the long-term goal that has been chosen by the firm by the end of the simulation.

Column 6: The fraction of firms that have achieved their chosen long-term goal by the end of the simulation.

Column 7: The number of decision changes a firm is away from its long-term goal by the end of the simulation.

Column 8: The number of times a firm has selected a new long-term goal by the end of the simulation.

Column 9: The number of times a firm has accepted a decision change by the end of the simulation.

Column 10: The number of times a firm has accepted a decision that would lead to an immediate performance decline but brought the firm closer to their long-term goal.

Column 11: The fraction of accepted decision changes that led to an instantaneous decline in performance but brought the firm closer to their long-term goal.

Column 12: The proportion of accepted decision changes that led to an instantaneous increase in performance.

Column 13: The number of times a firm has rejected a decision that would lead to an immediate performance gain.

Column 14: The total number of time periods a firm has experienced a lower performance than the highest performance it has accomplished thus far.

Column 15: A firm's average performance over all time periods until it reaches a steady-state.

Appendix 2. End of simulation performance distribution decomposed into stepping-stone decision count.

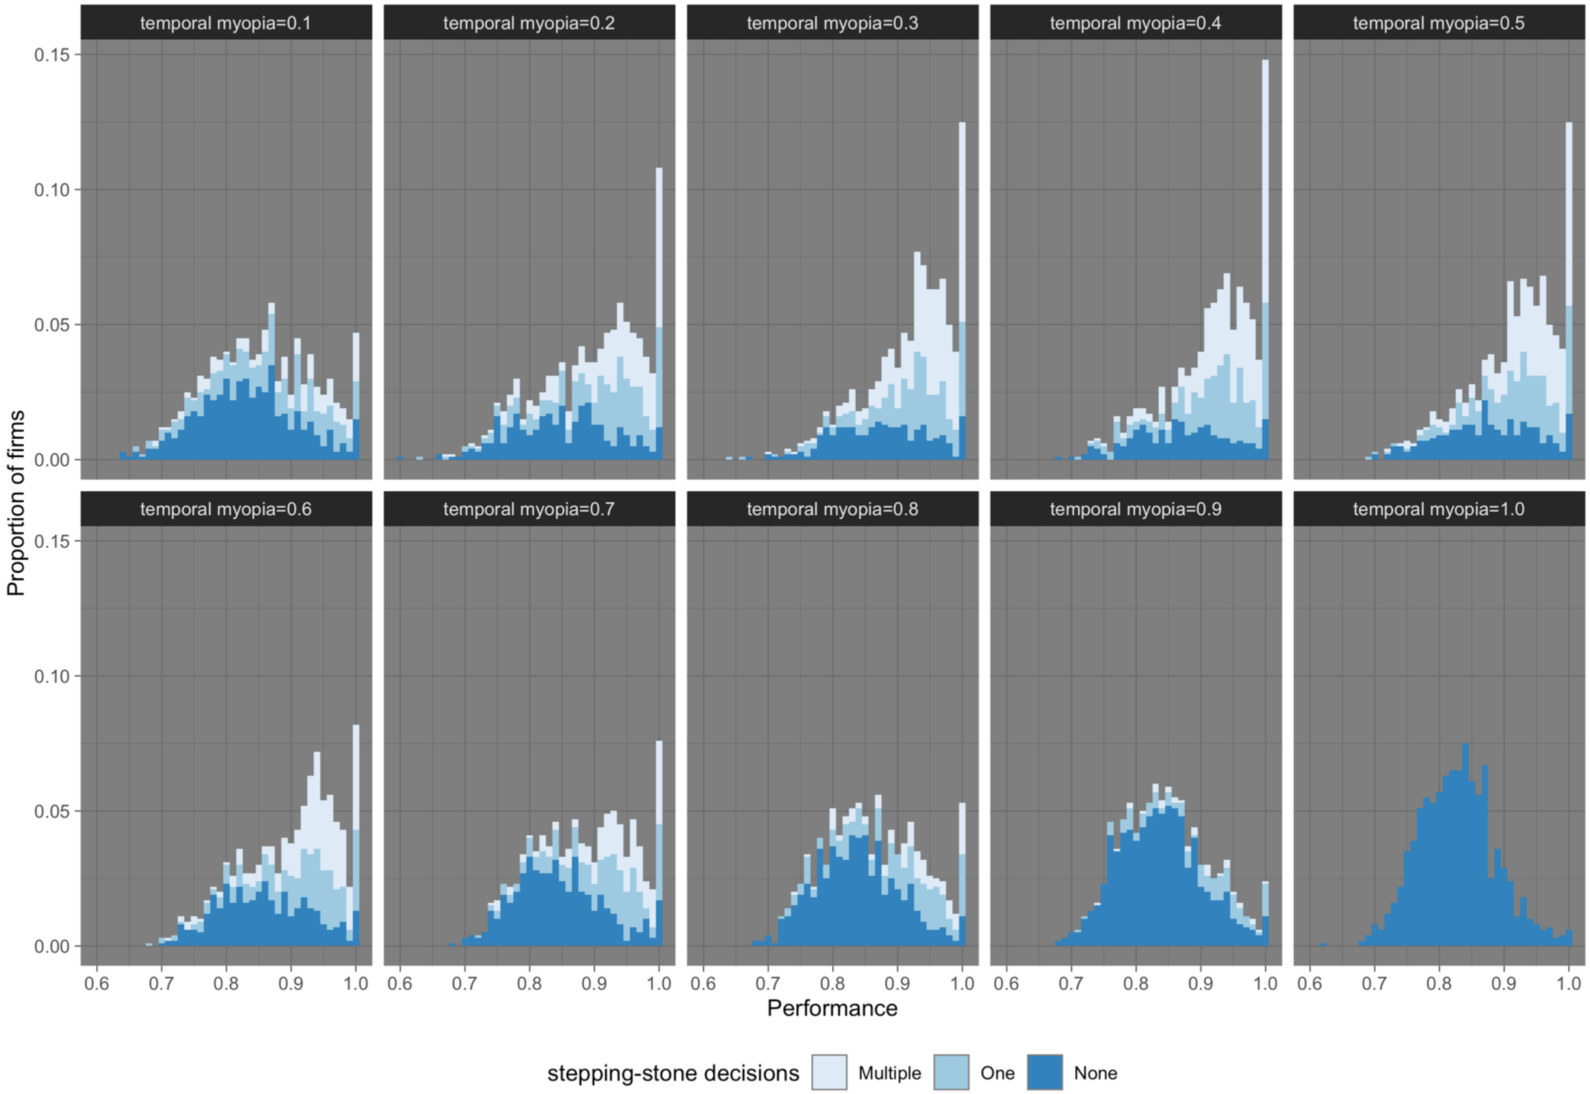

Appendix 3. Average performance of firm temporal horizon depending on discount factor for different levels of complexity.

Panel I: Environmental complexity medium (*K*=6)	Panel II: Environmental complexity low (*K*=1)
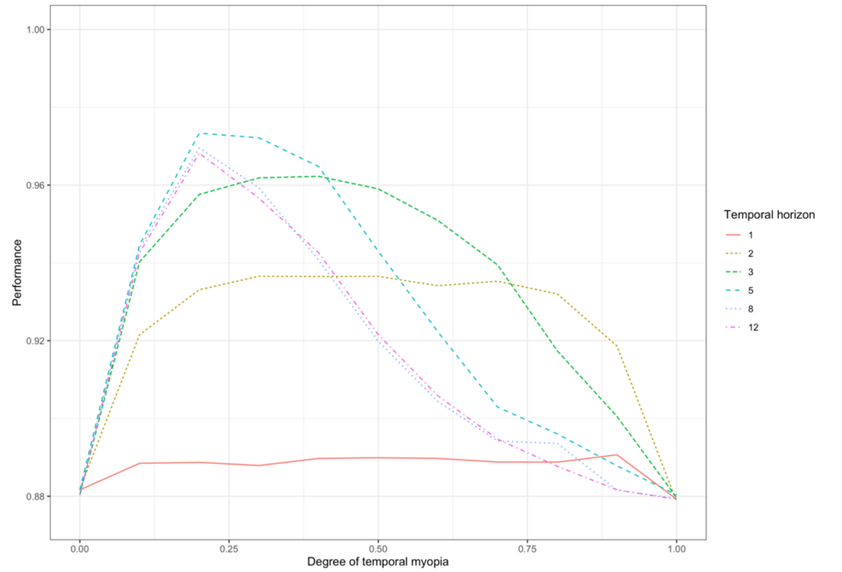	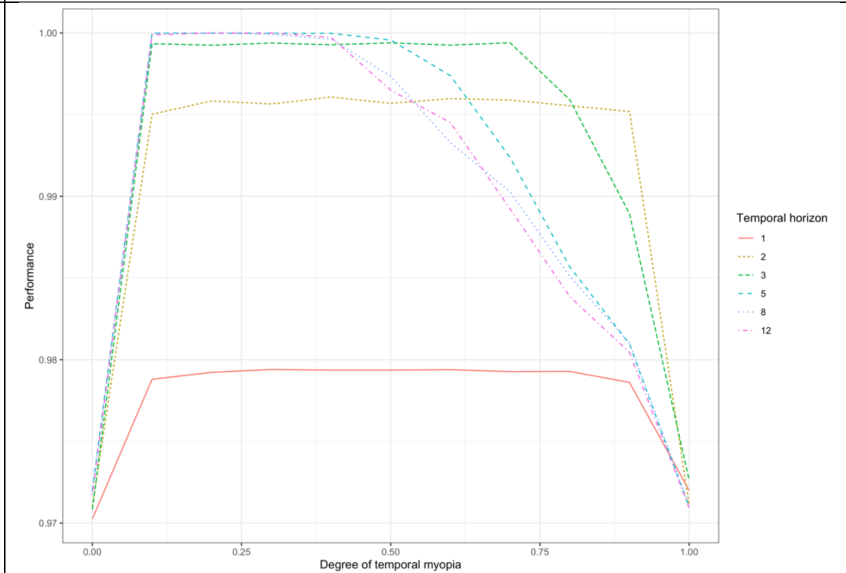

Note. Steady-state performance for medium environmental complexity (Panel I) and low environmental complexity (Panel II). Full statistics available from the authors.

Appendix 4. Illustration of decision algorithm

A firm with m =1 reflects the classic local search agent who attributes no additional value to a decisions positional contribution and only take into consideration the immediate performance. However, firms that also consider information about their configuration's positional value (m < 1) differ from those firms that are blind to positional values in their assessment of alternatives. In addition to the immediate reward, firms that don't fully discount (m<1) will take the discounted positional value into account. For example, let us assume a firm is comparing its current configuration s=[000000000000] against the configuration s'=[1000000000000]. Let us also assume Π_s =0.5 and $\Pi_{s'}$ =0.45. The firm's currently chosen long-term goal has a value of Π_l=0.7 and is currently three steps away from the firm's knowledge configuration (l=[111000000000000]). That is, the positional value for *s* is $(1\text{-}m)^3$*0.7. A firm which with *m*=0.1 (little myopia) will assign a positional value of 0.510 in addition to the immediate reward Π_s of the chosen configuration s. The decision that brings the firm one decision closer to the long-term goal, in form of configuration s', will be assigned a positional value of 0.567 in addition to the new configuration's immediate value $\Pi_{s'}$. Following equation (2), the organization will evaluate Q(s) = 0.5+0.510 = 1.01 and Q(s') = 0.45+0.567=1.017; and consequently, chooses to adopt configuration *s'*. In contrast, a firm with high myopia (e.g., *m*=0.9) will reject s' because of its evaluation of Q(s)=0.5+0.001=0.501 > Q(s')=0.45+0.07=0.457.