

Intermittent Animal Behavior: The Adjustment-Deployment Dilemma

Miguel Aguilera^{*,**}
University of Zaragoza

Manuel G. Bedia^{*,*}
University of Zaragoza

Francisco Seron^{*,*}
University of Zaragoza

Xabier E. Barandiaran[†]
University of the Basque Country

Abstract Intermittency is ubiquitous in animal behavior. We depict a coordination problem that is part of the more general structure of intermittent adaptation: the *adjustment-deployment dilemma*. It captures the intricate compromise between the time spent in adjusting a response and the time used to deploy it: The adjustment process improves fitness with time, but during deployment fitness of the solution decays as environmental conditions change. We provide a formal characterization of the dilemma, and solve it using computational methods. We find that the optimal solution always results in a high intermittency between adjustment and deployment around a non-maximal fitness value. Furthermore we show that this non-maximal fitness value is directly determined by the ratio between the exponential coefficient of the fitness increase during adjustment and that of its decay coefficient during deployment. We compare the model results with experimental data obtained from observation and measurement of intermittent behavior in animals. Among other phenomena, the model is able to predict the uneven distribution of average duration of search and motion phases found among various species such as fishes, birds, and lizards. Despite the complexity of the problem, it can be shown to be solved by relatively simple mechanisms. We find that a model of a single continuous-time recurrent neuron, with the same parametric configuration, is capable of solving the dilemma for a wide set of conditions. We finally hypothesize that many of the different patterns of intermittent behavior found in nature might respond to optimal solutions of complexified versions of the adjustment-deployment dilemma under different constraints.

Keywords

Intermittent behavior, adaptive behavior, adjustment-deployment dilemma, behavioral modeling

A version of this paper with color figures is available online at http://dx.doi.org/10.1162/artl_a_00133. Subscription required.

1 Introduction

Most models of biological behavior are based on steady state assumptions, considering that the processes governing organisms occur in a constant and sustained way. However, activity in living beings at many different levels often happens in bursts, in states of marginal instability in which pauses are alternated with brief activity. These forms of intermittency have been linked [17] with processes of adaptation to ever-changing and unpredictable environments, revealing the importance of continuous interaction with the world when dealing with unknown situations. More

* Contact author.

** Department of Informatics, University of Zaragoza, 50018 Zaragoza, Spain. E-mail: miguel.academic@maguilera.net (M.A.)

† IAS—Research Centre for Life, Mind and Society & Department of Philosophy & University School of Social Work, UPV/EHU, University of the Basque Country, Av. de Tolosa 54, 20018 Donostia, Gipuzkoa, Spain.

importantly, intermittency brings up the importance of timing and coordination in cognitive processes, connecting with dynamical perspectives that have had a mayor impact in behavioral and cognitive sciences [5, 13, 26] during the last two decades.

Intermittent locomotion is a widespread biological phenomenon. Many organisms' behavior (ranging from protozoans to mammals) is intermittent: They move, pause briefly, and move again. These pauses last from milliseconds to minutes, being part of a dynamical system by which organisms adjust their behavior to changing environments [17].

Despite the energy costs of acceleration and deceleration, a variety of benefits arise when pauses are alternated with action. Intermittent bounding and undulating flight modes in birds (which alternate periods of flapping with pauses where wings are either extended to permit gliding or held close to the body) save mechanical power compared to continuous flight over a broad range of speeds [21]. A similar effect takes place in fishes when *burst-coast* swimming [28]. Many species, when chasing a prey, alternate pauses and moves to stabilize their sensory field. Thus, while moves tend to be straight, both pursuits of a prey and changes of direction are initiated after pauses [14, 18, 25]. *Saltatory search* in foraging animals (from insects and lizards to mammals) minimizes the search time by alternating phases of fast motion and phases of intensive search [3, 8].

Additionally, intermittent behavior has benefits that are related not so much with locomotion dynamics as with the dynamics of sensorimotor processes such as attention to the visual field. For example, when examining the visual field, eye movement is not smooth but alternates rapid movements (saccades) with stable intervals (fixations) [22]. Other examples include primates pausing briefly while moving between trees in the canopy, the pauses being related to the requirement to identify a route for the next movement sequence [12]; or humans balancing a pole on their fingertips, displaying *on-off intermittency*, where most of the time the equilibration is stable, and corrective movements occur in quick bursts [10].

Broadly speaking, the nature of intermittent behavior can be considered to have two (non-exclusive, yet radically different) origins:

1. Intermittent behavior is just an epiphenomenon of embodied behavior; it is physical or dynamical constraints (e.g., the muscles needing to rest after some time of activation, or the refractory period of a neuron just after a spike, when it cannot fire again) that provoke intermittent behavior.
2. Intermittent behavior results from a strategy developed by organisms to face the challenge of dynamically adjusting their behavior to changing environments.

We may ask ourselves if the overabundance, and specific patterns, of intermittent modes of behavior in living beings are the result of a general strategy developed at different levels of biological organization to adapt to complex, ever-changing environments. This is the central question of this article. We do not deny that intermittent behavior is produced by physical or physiological constraints. But if these constraints appear so frequently, maybe it is because they play a fundamental role in organisms' adaptation. Neither do we suggest that all kinds of intermittency are caused by such a strategy. Most probably there are many instances of intermittent behavior that do not serve adaptive purposes. However, its presence over such a wide range of temporal scales and different types of organisms and forms of behavior suggests that there might be a more general reason for its existence than just the particular constraints that different organisms have to deal with. In this article we will try to show that there is in fact a general and wide adaptive problem, what we call the *adjustment-deployment dilemma* [1], for which intermittent behavior is an optimal solution.

We provide a formal characterization of the adjustment-deployment dilemma in which a system has to find an equilibrium between two complementary stages: adjustment of behavior to an environmental condition, and the deployment of that behavior. This dilemma captures the difficult compromise between the time spent in adjusting a response and the time used to deploy it: The adjustment process improves fitness with time, but does not directly benefit the

organism or contribute to the task goal until the adjusted response is deployed. However, during deployment the benefit of the adjustment phase starts to decay as the environment changes. As a result, if you spend too little time adjusting your behavior, the results of your action are poor, but if you spend too much during deployment, the result is no longer valid. As we shall see, the adjustment-deployment dilemma is able to offer a formal explanation of specific patterns of intermittent behavior in terms of adaptive efficiency. It shows how, under some conditions, the best response of a system (be it a neural ensemble, a control system, an organism, etc.) is to rapidly alternate between different modes of behavior.

In the next section, we review the assumptions we made in order to present a general minimal formal model of intermittent adaptation. Sections 2 and 3 introduce the model and its assumptions. Section 4 compares its results with experimental data. In Section 5 we present a minimal implementation that is able to solve the dilemma for a wide range of environment dynamics. Finally, Section 6 suggests some directions for future research, and Section 7 discusses the implications of the presented model.

2 The Adjustment-Deployment Dilemma: Characterization, Scope, and Modeling Assumptions

The challenge we are proposing is to present a model that can explain different kinds of adaptive intermittency within a common framework. The scope is meant to be general and multi-scale, covering interspecies differences as well as various types of intermittent behavior within the same organism. Ultimately, we expect it to be applicable to intermittent functioning of different components within a behavior generating mechanism (although, throughout the article, we will favor a behavioral agent-environment interpretation).

We consider that some of the previous approaches to intermittency (see [1] for a review) might have failed to build this common framework because they were typically limited to the study of a particular case of intermittent behavior. It is also worth noting that there have been few, if any, attempts to describe intermittency as a general systemic property that results from agent-environment adaptive feedback dynamics with a specific temporal structure. Descriptions of intermittent behavior in organisms are just concerned with the actions of the agent, typically dismissing the role of the environment and the agent's coupling with it. Part of the problem lies in the fact that a systemic approach brings forth many problems due to the great complexity of the many levels of interaction present in living beings. One aspect of this complexity is the fact that some adaptive properties are not shown for a single instance of behavior, but in a wider context of recurrent agent-environment interactions. Examples of adaptive and yet counter-intuitive (when studied in a single instance) forms of behavior are ubiquitous in nature—some examples include group foraging [23], host-parasite-predation interactions [16], and mating behavior [2]). We will favor a systemic approach, tackling the complexity of the task by reducing our model to a minimal expression that will still be able to capture the essence of the problem. Trying to reach the mathematical abstraction behind intermittent adaptation, we assume the simplest possible case of intermittency, where:

1. The system switches between only two possible behaviors. We do not consider situations where the system has a more complex repertoire of behaviors.
2. These behaviors cannot overlap in time; neither can there be a situation where the system is not performing any of them.
3. Transition times between states (modes of behavior) are considered to be small enough to be neglected.
4. Though many behaviors might adapt the system to its environment (adjustment) or improve the system's situation in that environment (deployment) at the same time and

in different degrees, we consider that the system cannot do these two things at these same time.

5. It is, in principle, possible to measure fitness of the agent-environment relationship at any given time, potential or virtual fitness during adjustment, and real or effective fitness during deployment.

Needless to say, these assumptions do not cover all the possible instances of intermittent behavior found in nature. In some cases behaviors are going to overlap in time, or transition times between behaviors will not be negligible. However, what we intend here is to make a first approximation, taking a minimal example that preserves the essence of the intermittent phenomena. Despite its simplicity, this minimal model can represent a lot of cases of intermittent behavior in animals. In order to illustrate the dilemma and its mathematical formulation, we shall make use of an example of the type of phenomena we are about to model: a predator having to decide whether to run chasing a prey or stop to stabilize its visual field and precisely locate the position of the prey (Figure 1).

In our model, we assume that an intermittent system generates a pattern that combines two mutually exclusive stages:

- *Adjustment* is a behavior that improves the position of the organism and increases its possibilities of achieving its goals by adjusting its possibilities for effective action in an specific task. In our example, adjustment will correspond to stopping during a pursuit to stabilize the visual field to localize the prey's position.
- *Deployment* is a behavior that takes advantage of the possibilities generated during the previous phase, executing an action that makes them effective (deploys them). Moving toward the chased prey would correspond to the deployment phase in our example.

The intermittency between adjustment and deployment is not a mere sequential ordering of phases of adjustment followed by phases of deployment, but poses a problem of functional coordination dynamics: How much time do I need to spend focusing on a prey before I move? What is the best ratio between stopping for sensory stabilization and moving during a pursuit? A correct dynamic equilibrium between adjustment and deployment is crucial in most cases and might change under different circumstances. We have coined the term *adjustment-deployment dilemma*




Figure 1. Representation of a toad moving intermittently while chasing a prey. The toad only can change its direction when it stops, due to the visual blurring that makes it blind while moving. This is an example of the adjustment-deployment dilemma, where the toad has to find a compromise between the time it spends moving toward the prey (deployment) and the time it spends unmoving, stabilizing its visual field and reorienting itself (adjustment).

to name a generic characterization of this problem. To our knowledge, no explicit theoretical, mathematical, or simulation approach has yet explicitly addressed it.

There are, however, several models that address the nature of intermittent behavior in specific contexts. Some quite interesting research in intermittent behavior has been developed in the field of intermittent search strategies, where kinetic models have been proposed in which intermittent behavior emerges as an optimal strategy for detecting prey [9, 11]. As well, intermittent behavior in gradient-climbing organisms based on sporadic cues and partial information has been modeled by the so-called *infotaxis* model, in which the searcher adopts a strategy of movement alternating exploration and exploitation phases, in order to maximize the expected rate of information gain [27]. As well, intermittent flight modes in birds have been modeled in terms of aerodynamic and energetic considerations [21].

Still, such models represent particular stances of intermittent behavior, which arise from particular constraints of the task the organism is facing. The adjustment-deployment model presented here aspires to be a general model of intermittent behavior, able to provide an explanation for the recurrence of intermittent solutions emerging as optimal strategies in a wide range of tasks. In the following sections, we try to define the simplest model that is able to capture the essence of different kinds of intermittent behavior, abstracting away some of the particularities of the specific task models mentioned above.

3 Formalization of the Adjustment-Deployment Dilemma

In order to explore the mathematical core of the adjustment-deployment dilemma, we have simplified the problem to its minimal form. In general terms we have a system adjusting its behavior (or solution to a problem) and then executing or deploying it. We can take as an example the case of a toad chasing a prey, having to alternate movement with pauses for stabilizing its visual field [18]. The toad cannot see while it is moving, because its visual field blurs. Thus, the toad has to stop for some instants to locate the position of the prey. In the absence of obstacles the toad moves toward the position where the prey was just the instant before the toad started to move. Prey velocity has no influence on the direction of the toad's movement. Also, while the toad is moving, it is not going to correct its course if the prey changes its position (it cannot perceive such a change). The distance the toad hops in a single bound depends on the initial separation between the toad and prey, and it is not altered if the prey vanishes or moves during the toad's approach. Both the distance moved and the direction of the toad are uncorrected by visual feedback until the toad stops its movement (Figure 1).

In terms of our adjustment-deployment dilemma, the toad has to alternate between a *move* state (deployment), where it can approach the prey, and a *stop* state (adjustment), in which it can stabilize the image it perceives and update the information about the prey's position. Thus, the toad has to find an equilibrium between how much time it spends adjusting its visualization of the prey and how much time it subsequently spends deploying a pursuit behavior. We also can see how the relative amount of time expended in either state is going to depend on the dynamics of the situation. When the prey moves slowly or when it is far away, the toad has less necessity of adjusting its behavior, and can move for longer amounts of time; when the prey moves fast or it is too close, the toad has to stop and adjust its orientation more frequently, having less time for effectively moving toward the prey.

More explicitly, we have expressed the model in a series of mathematical terms, which are seen in Table 1. We introduce them below.

3.1 Fitness

Fitness represents the mean ability of a system to maximize the chances of achieving its goals, that is, obtaining a successful solution for a given problem or situation. The fitness or quality of a solution at

Table I. Minimal intermittent behavioral model: Concepts.

Concept	Notation	Formulation	Description
Fitness	$f(t)$	Adjustment (virtual): $f(t) = 1 - e^{-t/\tau}$	Evolution of the virtual or effective fitness in relation to task goal during a behavioral phase (increases for adjustment and decays for deployment).
		Deployment (effective): $f(t) = e^{-t/\epsilon}$	
Choice	$\gamma(t)$	Adjustment: $\gamma(t) = \gamma_0$ Deployment: $\gamma(t) = \gamma_1$	Binary exclusive choice of a system over time between adjustment and deployment.
Performance	$\bar{p}(t)$	$\bar{p}(T) = \frac{1}{T} \int_0^T \gamma(t) \cdot f(t) dt$	Mean accumulated effective fitness during the task duration
Optimal solution	$f^*(\tau, \epsilon)$	$f^*(\tau, \epsilon) = \operatorname{argmax}_{\gamma(t)} \bar{p}(t)$	Optimal choice dynamics that maximizes mean effective fitness.

an instant t is denoted by a fitness function $f(t) \in [0, 1]$. Note that fitness here does not mean evolutionary or survival fitness directly. Rather it denotes fitness in relation to the task goal (e.g., following a prey), which shall in turn make a contribution to survival fitness (e.g., catching and eating the prey). We will assume that:

1. The system has an adjustment mechanism for improving its behavior with respect to the environment. We assume that the functional relation between the quality of a solution and time during adjustment is known, and we consider it to be a nonlinear function (the effort in obtaining better results grows in relative terms with time), and we assume it to be exponential, $f(t) = 1 - e^{-t/\tau}$, where τ is the adjustment speed.
2. We assume that the solution degrades throughout time due to environment changes during the deployment phase. Also taken as exponential is the functional dependence between quality of a solution and time: $f(t) = e^{-t/\epsilon}$, where ϵ stands for the degradation rate.

It has to be noted that although the fitness is generated during the adjustment phase, it can only be exploited during the deployment phase. To stress this fact, we will refer to the fitness value of the solution as *virtual* while the organism is adjusting and *effective* when the agent deploys the solution.

In the case of the toad, the fitness will correspond to the difference between the prey’s azimuth and the toad’s orientation. When the toad is pointed at the prey, the fitness value is 1, and it will decrease when the prey changes its position. In general, we are going to consider ϵ and τ as constants defining the dynamics of an environment. However, in Section 5 we will show how a system can adapt to an environment where it has to face different possible values of ϵ and τ .

At this point we have to make an important clarification. Even if we define $f(t)$ as a deterministic function, it does not mean that the environment we are modeling is predictable. On the contrary, the adjustment-deployment dilemma describes a situation where a system has to adapt to an uncertain and changing environment. We can know how fast the system is changing (how fast the prey moves away from the toad, or how fast the toad stabilizes its visual field), but we do not know *how* the system is changing, that is, we cannot predict the future position of the prey. Knowing the adjustment and degradation rates of the changes in the system does not mean that we can predict the state of the environment, just the state of the current fitness or adaptation level. Throughout the article, when we refer to the uncertainty or unpredictability of the environment, we shall be referring to this impossibility of predicting future states of the environment.

3.2 Choice

The resolution structure of the dilemma can be captured with a single variable denoted by $\gamma(t) \in \{\gamma_0, \gamma_1\}$, that is, as the binary exclusive choice of the system over time, with γ_0 representing adjustment and γ_1 deployment.

3.3 Formulation

Now, the following equations to describe the behavior of the system result from the previous formalization:

- *Adjustment:* $f(t) = 1 - e^{-t/\tau}, \gamma(t) = \gamma_0$
- *Deployment:* $f(t) = e^{-t/\varepsilon}, \gamma(t) = \gamma_1$

The structure of the dilemma can thus be reduced to finding the strategy (i.e., the form of $\gamma(t)$) that obtains the best results. However, the crucial point is that either adjustment or deployment requires a minimum duration to have an effect (e.g., a toad cannot perform half a jump). Therefore, after the value of $\gamma(t)$ switches it has to be maintained for a minimal time span. We will refer to the minimal adjustment and deployment periods as b_A and b_D respectively. Thus, the function $\gamma(t)$ can be reduced to a discrete sequence $\{\gamma_k\}$, where each value has to be maintained for its corresponding period, b_A or b_D . For example, each discrete value of γ_k could correspond to an action where the toad either performs a complete jump or stays still for a period of time.

3.4 Adjustment-Deployment Model

In order to compute the quality of the obtained results by a specific choice function $\gamma(t)$, we will specify the evolution of the fitness over time:

$$\dot{f}(t) = \begin{cases} \frac{1}{\tau}(1-f(t)), & \gamma(t) = \gamma_0 \\ -\frac{1}{\varepsilon}f(t), & \gamma(t) = \gamma_1 \end{cases} \tag{1}$$

The agent’s performance will be obtained by just integrating the fitness of the system during the deployment periods (as we said before, we assume that the system only can take advantage of its situation in the world during the deployment phase, where fitness becomes effective, unlike the virtual fitness obtained during adjustment periods). We will take $\gamma_0 = 0$ and $\gamma_1 = 1$, allowing us to combine both previous functions in a single equation describing the global behavior:

$$\dot{f}(t) = -\gamma(t) \cdot \frac{1}{\varepsilon}f(t) + (1-\gamma(t)) \cdot \frac{1}{\tau}(1-f(t)) \tag{2}$$

The quality of the obtained results will be defined by the performance (mean accumulated effective fitness) of the agent, $\bar{p}(T)$, evaluated in an interval $(0, T)$:

$$\bar{p}(T) = \frac{1}{T} \int_0^T \gamma(t) \cdot f(t) dt \tag{3}$$

Having the problem so defined, the best solution is the one that offers a maximum value of $\bar{p}(T)$. Yet, the mathematical solution to the problem is nontrivial. The cost of computing the evolution of $f(t)$ for every possible combination of values of $\gamma(t)$ is prohibitive. In order to circumvent this problem we have used the Bellman algorithm [7] for finding the optimal solution in a recursive way (Appendix 1 shows the mathematical derivation of the solution). The result is shown in Figure 2, where we can see the solution of the problem for given values of τ and ε . We see how the optimal




Figure 2. Representation of the optimal strategy for different situations: (a) fitness function for $\tau = 1, \epsilon = 1$, (b) fitness function for $\tau = 1, \epsilon = 0.25$, (c) fitness function for $\tau = 0.25, \epsilon = 1$. The dashed line represents the value of f_n^* .

strategy for solving the adjustment-deployment dilemma is not the one that maximizes the fitness at a given instant of time. Instead, the best solution is one that reaches a suboptimal solution and, instead of enhancing it, maintains it constant through time. Also, if the result is constrained to a limited time window, the final steps of the behavior exploit all the accumulated fitness in a final prolonged adjustment phase. Our results show that the global solution crucially depends on the accurate coordination of the agent-environment interactions.

3.5 Intermittent Adaptation: Maximizing Interactions with the Environment

What do the previous results mean? As seen in Figure 2, the obtained optimal strategy for solving the adjustment-deployment dilemma tends:

- not to maximize fitness, but to reach a intermediate value $f^*(\tau, \epsilon)$ which is kept until the process is about to end.
- to maximize the number of behavioral changes (i.e., the alternation between adjustment and deployment).

Let’s go back to the toad. We have an agent that has to act in a changing environment. Also, when the toad is performing a task in this environment (chasing the prey), the toad does not know how the environment is changing. However, we suppose that the toad knows (i.e., can adapt to) how fast the environment (the position of the prey) is changing, so the toad can have a measure of how long a time it can move until the direction of its movement is no longer valid. In an intuitive first approach to the problem, we could think that what the toad has to do is just change its orientation until it is pointing toward the prey, and start moving until it reaches a point where its orientation is no longer valid.

Nevertheless, our result shows that this intuitive view of the problem is not right, at least not always. If the movement of the prey is fast enough compared with the time it takes the toad to perceive the position of the prey, the optimal solution to the problem implies that the toad does not have to perceive the exact position of the prey. Instead, a less accurate but faster-to-obtain orientation (e.g., not waiting for its visual field to stabilize completely) is going to be a better choice. From that moment, fast alternations between movement and orientation will result in the optimal strategy. Also, the precision of the toad’s orientation is going to be kept fixed at a suboptimal point. The value of this point of optimal behavior at any instant of time, $f^*(\tau, \epsilon)$, is going to depend on the relation between τ and ϵ . Appendix 2 shows that, under some assumptions, the value of $f^*(\tau, \epsilon)$ can be computed by the following equation (Figure 3):

$$f^*(\tau, \epsilon) = \frac{1}{1 + \sqrt{\frac{\tau}{\epsilon}}} \tag{4}$$

In a nutshell, the optimal solution to the adjustment-deployment dilemma can be captured under the following dictum: When the environment changes, the best behavior is the one that maximizes the number of interactions with the world, the optimal fitness level being determined by the ratio between the organism’s adjustment speed and the environmental rate of change. The timing between interactions (i.e., the transitions between adjustment and deployment in the oscillations) is determined by physiological and environmental constraints (for a list of timings in different animals, see [17]). The last part of the conclusion is especially interesting, since it adds a new condition for adaptation by means of intermittent behavior. According to this result, adaptation to the environment is not always going to require well-adjusted solutions. Instead, suboptimal solutions combined in an intermittent way will be the best strategy to cope with changing environments. As well, this suboptimal fitness value is not going to be determined by either the agent or the environment alone, but it is going to be a result of the dynamical coupling between both.

4 Comparison with Experimental Data

An interesting result of the presented model is that $f^*(\tau, \epsilon)$ determines the amount of time that a system spends in adjustment and deployment. Specifically, under some conditions (see Appendix 2), the relative time spent in deployment, r_{dep} , is going to be equal to the optimal fitness value:

$$r_{dep} = \frac{1}{T} \int_0^T \gamma(t) \cong f^*(\tau, \epsilon) \tag{5}$$

It follows that when adaptation is slower than environmental changes, an organism will need to spend more time in adjustment than in deployment. As well, it will be forced to develop strategies with poorer solution quality. That is coherent with empirical data:

- In adult viviparous lizards r_{dep} is around 0.7 and 0.8 for general locomotion, while it is reduced to nearly 0.25 when the lizards are actively searching for prey [4]. That is, when an agent has enough time to exploit its adjustment, it can afford high-fitness strategies (Figure 2c), while low-fitness strategies will be developed by an agent when the available deployment time is smaller (Figure 2b).




Figure 3. Quality of the solution for the optimal strategy $f^*(\tau, \epsilon)$ for different values of $\epsilon(t)$ and $\tau(t)$. The value of $f^*(\tau, \epsilon)$ is going to determine the relative amount of time spent in deployment.




Figure 4. Distribution of average duration of search and motion phases for various species such as fishes, birds, and lizards exhibiting saltatory search behavior, adapted from [8] (permission granted by the author). The parameter ρ (τ in the original article) represents the relation between the durations of the two phases. The first peak ($\rho \sim 0.1$ s) corresponds to foragers in regime S, which spend more time searching than moving. The second peak ($\rho \sim 25$ s) corresponds to foragers in regime M, which spend more time moving than searching.

- Several studies have pointed out behavioral changes of animals looking for prey as the search environment changes. When prey are more difficult to detect or when environments are visually more complex, the value of r_{dep} decreases [20, 24].

The percentage of the time spent in deployment varies greatly among different organisms. As seen in [17], r_{dep} ranges from 0.04 to 0.94 for different tasks and species. Also, according to experimental data [8] (Figure 4), r_{dep} follows a binomial distribution in foraging animals, meaning that most foragers either spend more time searching than moving or spend more time moving than searching; very few foragers spend similar amounts of time searching and moving. Such results are seen in the binomial distribution of Figure 3, where, if ε/τ is assumed to be log-uniformly distributed (e.g., if we assume that it is reasonable that activity in nature occurs with similar probability at all temporal scales), in most cases r_{dep} will be either small or large, and only in a small percentage of cases will it have medium values. The sigmoidal relation between r_{dep} and ε/τ makes it likely that distributions of r_{dep} will tend to be overrepresented in the extremes (for small and large values).

5 A Minimal Model Implementation of the Adjustment-Deployment Dilemma

We have presented a formal characterization of the adjustment-deployment dilemma and a formal optimal solution for different parametric configurations of the dilemma. We can now ask the following questions: Can evolutionary, developmental, or learning processes lead to an organism that can find (or approximate) this solution? If so, what is the simplest mechanism that can match an optimal solution to the adjustment-deployment dilemma?

In order to answer these questions we have used artificial evolution to evolve a behavioral selection mechanism whose simplicity and biological plausibility could be assumed for wide range of organisms. We have used continuous-time recurrent neural networks (CTRNNs) to implement a dynamical system capable of developing the optimal strategy for a wide range of possible situations (i.e., ε -values). CTRNNs have been among the most popular neural controllers for designing adaptive systems within a dynamical perspective [6]. They constitute a good choice for the proposed task because (1) they are the simplest nonlinear, continuous dynamical neural network model; (2) despite their simplicity, they are universal dynamics approximators in the sense that, for any finite interval of time, CTRNNs (provided that there is no constraint on the number of nodes) can approximate the trajectories of any smooth dynamical system [15].

The general form of a CTRNN with N neurons is

$$j_i = \frac{1}{\alpha_i} \left(y_i + \sum_{j=1}^N \omega_{ij} \cdot \sigma(y_j + \theta_j) + I_i \right) \tag{6}$$

where $i = 1, 2, \dots, N$, y is the state of each neuron, τ is its time constant ($\tau > 0$), w_{ij} is the strength of the connection from the j th to the i th neuron, θ is a bias term, g is a gain term, $\sigma(x) = 1/(1 + e^{-x})$ is the standard activation sigmoidal function, and I represents a constant external input. We allow each neuron to have external input information about three variables in the environment. Each neuron will have access to (1) the current quality of the solution being implemented (the value of the fitness $f(t)$), (2) how fast the current quality of the virtual fitness improves over time in the adjustment phase, and (3) how fast the effective fitness decays during the deployment phase. We define the external input for each neuron as the weighted sum of these three variables:

$$I_i = s_{i1} \cdot f(t) + s_{i2} \cdot \dot{f}_A(t) + s_{i3} \cdot \dot{f}_D(t) \tag{7}$$

where $\dot{f}_A(t)$ is equal to the last value of $\dot{f}(t)$ when the system was in the adjustment phase ($\gamma(t) = 0$), and $\dot{f}_D(t)$ is equal to the last value of $\dot{f}(t)$ when the system was in the deployment phase ($\gamma(t) = 1$); here $\dot{f}(t)$ is the first derivative of $f(t)$. Note that $\dot{f}_A(t)$ and $\dot{f}_D(t)$ cannot be mutually activated, so in adjustment $\dot{f}_D(t) = 0$ and in deployment $\dot{f}_A(t) = 0$.

One of the neurons (say $i = N$) is considered as the output of the system. Whenever an adjustment or a deployment period is finished (it has been active for a time that is a multiple of b_A or b_D , respectively), the value of $\gamma(t)$ will be actualized according to the state of the output:

$$\gamma(t) = \begin{cases} 0, & \sigma(y_N(t) + \theta_N) \leq 0.5 \\ 1, & \sigma(y_N(t) + \theta_N) > 0.5 \end{cases} \tag{8}$$

The CTRNN runs with an Euler step of $b = 0.01$ s, and we initially assigned arbitrary values to the adjustment and deployment minimal steps, $b_A = b_D = 10 \cdot b = 0.1$ s, which, according to [1], is approximately the mean duration of pauses and movement periods in intermittent animal behavior.

Once the neural networks are defined, we want to find the minimal configuration that successfully solves the adjustment-deployment dilemma. Also, we want the network to be able to adapt to different values of τ and ϵ .

In order to find the appropriate values for the parameters of the network to solve the problem, we used artificial evolution: a rank-based genetic algorithm with elitism and binary encoding [28]. We ran different genetic algorithms for populations of 60 neural networks for 12 generations, each genetic algorithm having a population of networks of a different size, from one to six neurons. Each neural network was evaluated against an environment with changing parameters during trials of $T = 200$ s. The environmental parameters were changed every 10 s, generating new random values of τ and ϵ from a function 2^x , where x was uniformly distributed in the interval $[-4, 4]$. The fitness function of the genetic algorithm was equal to the value of $\bar{p}(T)$ (the performance, i.e., the mean effective fitness, of the system in the adjustment-deployment dilemma; see Equation 3). The mutation probability was set to 0.01 for each binary digit of the chromosome.

5.1 Minimal Intermittent Adaptive Structure

Running the set of genetic algorithms, we found that even the one using a single neuron (see Figure 5) obtains optimum results for a wide variety of values of τ and ϵ (see Figure 6). From now on, we take the best result from the genetic algorithm with a single-neuron network as a




Figure 5. Behavior of the best single-neuron network from the genetic algorithm adapting to 20 randomly generated values of τ and ε in the adjustment-deployment dilemma.

representative example of a minimal intermittent adaptive structure. The neuron follows the equation

$$\dot{y} = \frac{1}{\alpha} (y + \omega \cdot \sigma(y + \theta) + (s_1 \cdot f(t) + s_2 \cdot \dot{f}_A(t) + s_3 \cdot \dot{f}_D(t))) \tag{9}$$

where $\alpha = 0.1067$, $\omega = -3.0645$, $\theta = 4.8387$, $s_1 = -10$, $s_2 = 9.3548$, and $s_3 = -4.8387$.

As shown by the value of ω , the neuron has a negative feedback self-connection. It behaves as a nonlinear oscillator when interacting with its environment. As we see (Figure 5), the neuron is able to adapt perfectly to different values of τ and ε , by modulating its oscillations through interaction with the environment (Figure 6). Furthermore, the neuron is also able to adapt to different values of h_A and h_D without the need of previous training (Figure 6). However, the dynamics of the neuron (concretely, its time constant α) limits how fast intermittency can happen, since the speed of the neural network is limited by the parameter α .

6 Discussion

The present model can be expanded and improved in different ways. Some of the underlying assumptions could be relaxed and the model complexified. For instance, many crucial temporal aspects of

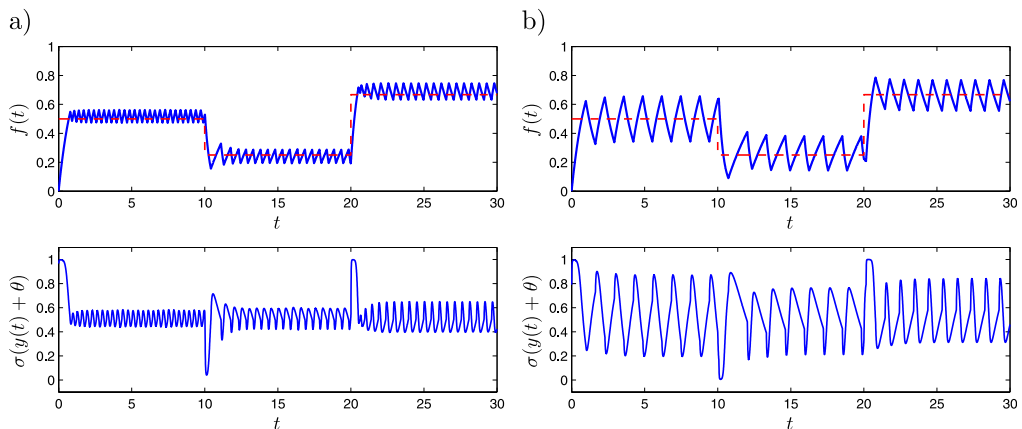


Figure 6. In the electronic version, fitness function $f(t)$, optimal fitness value $f^*(\tau, \varepsilon)$ (red dashed line), and system's response $\sigma(y(t) + \theta)$ for the single-neuron network with (a) $h_A = h_D = 10 \cdot h$ and (b) $h_A = h_D = 60 \cdot h$.

the adjustment-deployment dilemma were left aside in this study, and many of them might provide avenues for future research. The inclusion of forced perceptual delays, evaluation delays (organisms need to take some time to taste a food source, or to evaluate the outcome of its interaction), possible overlap between adjustment and deployment, constraints on deployment duration, and the like could be included in future developments. The measurement of fitness could also be enriched by including additional cost functions associated with deployment (energy expenditure), adjustment (risk of being detected/hunted), or intermittency itself. It is also important to acknowledge the lack of embodiment of the current model. This was crucial to achieve a model of wide generality, but its application will have to include a variety of spatial and embodiment constraints. In this line, future development should also include reference to and modeling of specific examples of animal behavior that face different versions of the adjustment-deployment dilemma in order to compare the model's predictions with experimental data and adjust the relevant parameters and dimensionality of the model.

Finally, for a more general model of intermittency, we should study cases with more than two possible behaviors, together with more complex dependences of the fitness function on the world, allowing the agent to perform behaviors that are able at the same time to adapt to the environment (adjustment) and to benefit from it (deployment) in different degrees.

Regarding the minimal mechanism capable of optimally solving the dilemma, its strongest limitation lies in the required input. Although the mechanism itself is simple (a unique and highly simplified neuron), it demands high-quality information about the problem (current fitness and indicators of current fitness change during deployment and adjustment) in order to perform the task. It is very unlikely that an organism has direct access to this information in any given task environment. Future work should reduce this assumption and try to find mechanisms that can solve the problem with poor or partial information. Alternatively, the possibility of more complex mechanisms could be explored, including specific cognitive mechanisms that could process sensory information and deliver the required input to an "adjustment-deployment neuron." It is more likely, however, that organisms do not modularize the problem and alternative solutions emerge out of brain-body-environment dynamics that were not considered in this article. In order to explore this possibility, full agent-environment models could be developed where the dilemma and its solution might emerge as higher-order phenomena from lower-level behavioral/adaptive capacities and the recurrent sensorimotor coupling with the environment. In any case the presented minimal mechanism provides a proof of concept to show that, despite the apparent mathematical complexity of the dilemma, relatively simple organisms could, at least in principle, be able to find optimal solutions to it.

7 Conclusion

We have shown how characteristic patterns of intermittency result from solutions to the *adjustment-deployment dilemma*: the dynamic interplay between the time spent adjusting a solution to the changing environment and the execution time taken by the deployment of the solution.

Despite its ubiquity in biological behavior, to our knowledge this is the first characterization, formalization, and modeling approach to the adjustment-deployment dilemma. We have formalized mathematically the structure of the dilemma and numerically computed its optimal solution for different configurations. The problem-structuring parameter was found to be the ratio between the rate of adjustment and the rate of fitness decay while deployment takes place. The optimal solution always results in a high intermittency between adjustment and deployment around a non-maximal fitness value. Furthermore, we have shown that this non-maximal fitness value is directly determined by the ratio between the exponential coefficient of the fitness increase during adjustment and the decay coefficient during deployment.

Our hypothesis is that at least part of the intermittent behavior displayed by living organism is a response to this dilemma, whose general solution can be captured by the motto "When the environment changes, the best behavior is the one that maximizes the number of interactions with the world, the optimal fitness level being determined by the dynamic ratio between adjustment speed and environmental change."

The distribution of optimal strategies over the range of parameter values takes a sigmoidal shape. It follows that most solutions will be distributed over the two extremes of the solution spaces: one where adjustment is very fast with long periods of deployment, and the another the opposite, where long periods of adjustment are followed by quick deployment. It turns out that the distribution of intermittency patterns found in animal behavior matches our model's optimal-solution distribution.

A simple model composed of a single neuron with negative feedback is able to display this optimal behavior, assuming the following inputs: an indicator of the success of its deployment, and the current adjustment and degradation rates. The implemented mechanism shows a high degree of robustness, being able to adapt to a wide range of possible configurations of the dilemma. These results suggest that optimal solutions to the adjustment-deployment dilemma could, in principle, be instantiated by very simple mechanisms, given the appropriate input, and should therefore be accessible even to unicellular systems.

Our model also brings forth the need to include the temporal dimension of agents and environment into current modeling frameworks. Adjustment speed, decay rates, deployment duration, patterns of intermittency, and so on crucially matter when it comes to real-world problem solving. Computational and representational approaches to cognition are prone to neglect such time-dependent phenomena and might often fail to account for natural behavior. Furthermore, they can fail to provide models that solve adaptive problems by means of temporally rich and structured agent-environment coordination patterns. The adjustment-deployment dilemma might constitute one such case where, given strong cognitivist assumptions, one would be tempted to build models that first compute a near-maximum-fitness solution and only then deliver an output command. We have shown, however, that the optimal solution to the adjustment-deployment dilemma exploits non-maximal solutions by means of fast intermittent behavior in a manner that, in addition, requires only very simple control mechanisms. It could be further conjectured that, under certain constraints, intermittency and, perhaps more generally, recurrent agent-environment suboptimal interactions provide robust and simple solutions to many adaptive problems. We have shown that the adjustment-deployment dilemma is one such case and that its solutions' distribution matches the patterns of intermittent behavior found in animals.

Acknowledgments

Miguel Aguilera, Manuel G. Bedia, and Francisco Serón were supported in part by the project TIN2011-24660 funded by the Spanish Ministerio de Ciencia e Innovación. Miguel Aguilera currently holds a FPU predoctoral fellowship from the Spanish Ministerio de Educación. During the development of this article Dr. Xabier E. Barandiaran held a postdoctoral position funded by FP7 project eSMCs IST-270212. He also acknowledges funding from the research project Autonomia y Niveles de Organización financed by the Spanish Government (FFI2011-25665) and IAS research group funding IT590-13 from the Basque Government (in which M.B. and M.A. are also collaborators).

References

1. Aguilera, M., Bedia, M. G., Barandiaran, X. E., & Serón, F. (2011). The adjustment-deployment dilemma in organisms' behaviour: Theoretical characterization and a model. In *Proceedings of the IEEE Symposium Series on Computational Intelligence 2011*, 116–123.
2. Alonzo, S. H., & Warner, R. R. (2000). Dynamic games and field experiments examining intra- and intersexual conflict: Explaining counterintuitive mating behavior in a Mediterranean wrasse, *Symphodus ocellatus*. *Behavioral Ecology*, 11(1), 56–70.
3. Anderson, J. P., Stephens, D. W., & Dunbar, S. R. (1997). Saltatory search: A theoretical analysis. *Behavioral Ecology*, 8(3), 307–317.
4. Avery, R. A., Mueller, C. F., Smith, J. A., & Bond, D. J. (1987). The movement patterns of lacertid lizards: Speed, gait and pauses in *Lacerta vivipara*. *Journal of Zoology*, 211(1), 47–63.
5. Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72(1–2), 173–215.

6. Beer, R. D. (2008). The dynamics of brain-body-environment systems: A status report. In *Handbook of cognitive science: An embodied approach*. Amsterdam: Elsevier.
7. Bellman, R. E. (1957). *Dynamic programming*. Princeton, NJ: Princeton University Press.
8. Bénichou, O., Coppey, M., Moreau, M., Suet, P. H., & Voiturier, R. (2005). A stochastic theory for the intermittent behaviour of foraging animals. *Physica A*, *356*(1), 151–156.
9. Bénichou, O., Loverdo, C., Moreau, M., & Voiturier, R. (2011). Intermittent search strategies. *Reviews of Modern Physics*, *83*(1), 81–129.
10. Cabrera, J. L., & Milton, J. G. (2002). On-off intermittency in a human balancing task. *Physical Review Letters*, *89*(15), 158702.
11. Campos, D., Méndez, V., & Bartumeus, F. (2012). Optimal intermittence in search strategies under speed-selective target detection. *Physical Review Letters*, *108*(2), 81–129.
12. Cannon, C. H., & Leighton, M. (1994). Comparative locomotor ecology of gibbons and macaques: Selection of canopy elements for crossing gaps. *American Journal of Physical Anthropology*, *93*(4), 505–524.
13. Clark, A. (1997). The dynamical challenge. *Cognitive Science*, *21*(4), 461–481.
14. Evans, B. I., & O'Brien, W. J. (1988). A reevaluation of the search cycle of planktivorous Arctic grayling, *Thymallus arcticus*. *Canadian Journal of Fisheries and Aquatic Sciences*, *45*(1), 187–192.
15. Funahashi, K.-i., & Nakamura, Y. (1993). Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks*, *6*(6), 801–806.
16. Hall, S. R., Duffy, M. A., & Cceres, C. E. (2005). Selective predation and productivity jointly drive complex behavior in host-parasite systems. *The American Naturalist*, *165*(1), 70–81.
17. Kramer, D., & McLaughlin, R. (2001). The behavioural ecology of intermittent locomotion. *American Zoologist*, *41*(2), 137–153.
18. Lock, A., & Collett, T. (1979). A toad's devious approach to its prey: A study of some complex uses of depth vision. *Journal of Comparative Physiology A—Neuroethology, Sensory, Neural, and Behavioral Physiology*, *131*(2), 179–189.
19. Mitchell, M. (1998). *An introduction to genetic algorithms*. Cambridge, MA: MIT Press.
20. O'Brien, W. J., Evans, B. I., & Browman, H. I. (1989). Flexible search tactics and efficient foraging in saltatory searching animals. *Oecologia*, *80*(1), 100–110.
21. Rayner, J. M. V., Viscardi, P. W., Ward, S., & Speakman, J. R. (2001). Aerodynamics and energetics of intermittent flight in birds. *American Zoologist*, *41*(2), 188–204.
22. Schall, J. D., & Thompson, K. G. (1999). Neural selection and control of visually guided eye movements. *Annual Review of Neuroscience*, *22*, 241–259.
23. Seth, A. K. (2001). Modeling group foraging: Individual suboptimality, interference, and a kind of matching. *Adaptive Behavior*, *9*(2), 67–89.
24. Sonerud, G. A. (1992). Search tactics of a pause-travel predator: Adaptive adjustments of perching times and move distances by hawk owls (*Surnia ulula*). *Behavioral Ecology and Sociobiology*, *30*(3), 207–217.
25. Tye, A. (1989). A model of search behaviour for the northern wheatear *Oenanthe oenanthe* (Aves, Turdidae) and other pause-travel predators. *Ethology*, *83*(1), 1–18.
26. van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy*, *92*(7), 345–381.
27. Vergassola, M., Villermaux, E., & Shraiman, B. I. (2007). “Infotaxis” as a strategy for searching without gradients. *Nature*, *445*(7126), 406–409.
28. Videler, J. J., & Weihs, D. (1982). Energetic advantages of burst-and-coast swimming of fish at high speeds. *Journal of Experimental Biology*, *97*, 169–178.

Appendix I: Formal Solution of the Adjustment-Deployment Dilemma

Once the problem is defined, we proceed to compute the values that offer a maximum value of $\bar{p}(T)$. We have

$$\begin{cases} \dot{f}(t) = \frac{1}{\tau}(1 - f(t)) - \gamma(t) \cdot \left(\frac{1}{\tau} + f(t) \cdot \left(\frac{1}{\epsilon} - \frac{1}{\tau}\right)\right) \\ \dot{p}(t) = \gamma(t) \cdot f(t) \end{cases}$$

where $\gamma(t) = \{0, 1\}$ and we want to find the set $\{\gamma_k\}$ that maximizes $p(t)$, where each γ_k corresponds to a period of adjustment or deployment, with duration b_A or b_D , respectively. For simplicity, we will take $b_A = b_D = b$, but the result is the same if the adjustment and deployment periods are different. Now we can discretize the system with step b (with $b_A \neq b_D$, the discretization step would be the least common multiple between them):

$$\begin{cases} f_{k+1} - f_k = -b\left(\frac{1}{\tau}(1 - f_k) - \gamma_k \cdot \left(\frac{1}{\tau} + f_k \cdot \left(\frac{1}{\epsilon} - \frac{1}{\tau}\right)\right)\right) \\ p_{k+1} - p_k = b(\gamma_k \cdot f_k) \end{cases}$$

where b is a temporal step, $k = 0, 1, 2, \dots, N$, so $f(0) = f_0, p(T) = p_N$, given $T = \{t_1, t_2, \dots, t_N\}$. For the sampled version, the problem can be reformulated by the following (knowing that b is constant): “Find the set of decisions $\{\gamma_k(t_k)\}$ that maximizes $\sum_{k=0}^N \gamma_k f_k$.” That is, the $\{\gamma_k(t_k)\}$ values must be computed providing that

$$p_N = \max_{\gamma_0, \gamma_1, \dots, \gamma_N} \sum_{k=0}^N \gamma_k f_k$$

which, since it starts at a_0 , will be denoted by $p_N^{\text{MAX}}(a_0)$.

For solving the problem we apply the Bellman algorithm: “An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision” [7]. The Bellman algorithm computes the complete sequence $(\gamma_0, \gamma_1, \dots, \gamma_N)$ in a recursive way and backward:

$$p_N^{\text{MAX}}(f_0) = \max_{\gamma_0} \left[\gamma_0 f_0 + \max_{\gamma_1, \dots, \gamma_N} \sum_{k=1}^N \gamma_k f_k \right]$$

where

$$p_{N-1}^{\text{MAX}}(f_1) = \max_{\gamma_1, \dots, \gamma_N} \sum_{k=1}^N \gamma_k f_k.$$

Iterating, we obtain the sequence

$$p_N^{\text{MAX}}(f_0) = \max_{\gamma_0} \left[\gamma_0 f_0 + \max_{\gamma_1} \left[\gamma_1 f_1 + \max_{\gamma_2} \left[\gamma_2 f_2 + \dots + \dots + \max_{\gamma_N} [\gamma_N f_N] \right] \dots \right] \right]$$

For solving the system, we must proceed from the last decision to the first. Since the last does not affect the future, the maximization is local. In our case, it is

$$\gamma_N(t_N) = \begin{cases} 1 & \text{if } f_N \in (f_M, 0), \dot{f}_N < 0 \\ 0 & \text{if } f_N \in (0, f_M), \dot{f}_N > 0 \end{cases}$$

and therefore,

$$p_0^{\text{MAX}}(f_N) = \begin{cases} f_N & \text{if } f_N \in (f_M, 0), \dot{f}_N < 0 \\ 0 & \text{if } f_N \in (0, f_M), \dot{f}_N > 0 \end{cases}$$

that is, depending on whether the system is in an adjustment or a deployment phase. Once we know what is the optimal decision for $\gamma_N(t_N)$, the previous instant $\gamma_{N-1}(t_{N-1})$ is computed, applying the following equation:

$$p_1^{\text{MAX}}(f_{N-1}) = \max_{\gamma_{N-1}} [\gamma_{N-1} f_{N-1} + p_0^{\text{MAX}}(f_N)]$$

We know that

$$f_N = f_{N-1} - \frac{b}{\tau} \left((1 - f_{N-1}) - \gamma_{N-1} \cdot \left(1 + f_{N-1} \cdot \left(\frac{\tau}{\varepsilon} - 1 \right) \right) \right)$$

Therefore,

$$p_1^{\text{MAX}}(f_{N-1}) = \max_{\gamma_{N-1}} \left[\gamma_{N-1} f_{N-1} + p_0^{\text{MAX}} \left(f_{N-1} - \frac{b}{\tau} \left((1 - f_{N-1}) - \gamma_{N-1} \cdot \left(1 + f_{N-1} \cdot \left(\frac{\tau}{\varepsilon} - 1 \right) \right) \right) \right) \right]$$

Given $\gamma_{N-1} = \{0, 1\}$, we only have to compute which one of the two cases is larger:

$$f_{N-1} + p_0^{\text{MAX}} \left[\left(1 - \frac{b}{\varepsilon} \right) \cdot f_{N-1} \right] \geq p_0^{\text{MAX}} \left[\left(1 - \frac{b}{\tau} \right) \cdot f_{N-1} + \frac{b}{\tau} \right]$$

The equilibrium condition is met for a critical f_{N-1} value, denoted as f_{N-1}^* , that allows us to rewrite the equation in the following way:

$$p_1^{\text{MAX}}(f_{N-1}) = \begin{cases} f_{N-1} + p_0^{\text{MAX}} \left[\left(1 - \frac{b}{\varepsilon} \right) \cdot f_{N-1} \right] & \text{if } f_{N-1} > f_{N-1}^* \\ p_0^{\text{MAX}} \left[\left(1 - \frac{b}{\tau} \right) \cdot f_{N-1} + \frac{b}{\tau} \right] & \text{if } f_{N-1} < f_{N-1}^* \end{cases}$$

The procedure can be repeated for $k = 2, \dots, N$, obtaining the values of $\{f_0^*, f_1^*, \dots, f_{N-1}^*, f_N^*\}$ by iteratively solving the equation

$$p_{N-k}^{\text{MAX}}(f_{N-k}) = \max_{\gamma_{N-k}} \left[\gamma_{N-k} f_{N-k} + p_{k-1}^{\text{MAX}} \left(f_{N-k} - \frac{b}{\tau} \left((1 - f_{N-k}) - \gamma_{N-k} \cdot \left(1 + f_{N-k} \cdot \left(\frac{\tau}{\varepsilon} - 1 \right) \right) \right) \right), \right. \\ \left. k = 1, \dots, N. \right]$$

Appendix 2: Approximation of the Optimal Adjustment in the Adjustment-Deployment Dilemma

Given that the best solution to the adjustment-deployment dilemma is the one that intermittently alternates adjustment and deployment to maintain the fitness level around a given value of $f^*(\tau, \epsilon)$, we can assume some conditions that will allow us to accurately compute an approximation of the value of $f^*(\tau, \epsilon)$. If we imagine a system performing an optimal behavior for the adjustment-deployment dilemma (as represented in Figure 2), the assumptions that allow us to compute an approximation of the optimal fitness level are the following:

- The time spent in the transitions at the start and the end of the performance (when the system is not around $f^*(\tau, \epsilon)$) is negligible compared to the time when the system is intermittently alternating adjustment and deployment.
- The minimum durations of adjustment and deployment (b_A and b_D) are small enough to consider that the changes in $f(t)$ during the cycles of either b_A and b_D are small enough to be negligible.

These assumptions allow us to reduce the problem to what happens when the system is alternating adjustment and deployment in a state where $f(t)$ is always very close to the optimal level $f^*(\tau, \epsilon)$. In this situation, it is going to be much easier to compute how effective the system is for given parameters. Thus, we can simplify Equation 3 in the following way:

$$\bar{p}(T) = \frac{1}{T} \int_0^T \gamma(t) \cdot f(t) dt \cong f^*(\tau, \epsilon) \cdot \frac{1}{T} \int_0^T \gamma(t) dt = f^*(\tau, \epsilon) \cdot r_{\text{dep}} = p^*$$

where r_{dep} represents the relative time spent in adjustment. As well, if we just consider the time when the system is alternating adjustment and deployment, we can approximate r_{dep} as the value that equilibrates the effects of adjustment and deployment reaching a steady value of $f(t)$:

$$r_{\text{dep}} \cdot \left. \frac{\partial f(t)}{\partial t} \right|_{\gamma(t)=1} + (1 - r_{\text{dep}}) \cdot \left. \frac{\partial f(t)}{\partial t} \right|_{\gamma(t)=0} = 0$$

which, assuming that we can approximate the value of $f(t)$ by $f^*(\tau, \epsilon)$, gives us

$$r_{\text{dep}} \cdot \left(-\frac{1}{\epsilon} \cdot f^*(\tau, \epsilon) \right) + (1 - r_{\text{dep}}) \cdot \left(\frac{1}{\tau} \cdot (1 - f^*(\tau, \epsilon)) \right) = 0$$

$$r_{\text{dep}} = \frac{1}{1 + \frac{\tau}{\epsilon} \cdot \frac{f^*(\tau, \epsilon)}{1 - f^*(\tau, \epsilon)}}$$

$$\bar{p}(T) \cong p^* = \frac{f^*(\tau, \epsilon)}{1 + \frac{\tau}{\epsilon} \cdot \frac{f^*(\tau, \epsilon)}{1 - f^*(\tau, \epsilon)}} = \frac{f^*(\tau, \epsilon) - f^*(\tau, \epsilon)^2}{1 + f^*(\tau, \epsilon) \cdot \left(\frac{\tau}{\epsilon} - 1 \right)}$$

Now we can compute the value of $f^*(\tau, \epsilon)$ that gives us a maximum for $\bar{p}(T)$ just by computing when its first derivative is equal to zero:

$$\frac{\partial p^*}{\partial f^*(\tau, \epsilon)} = \frac{(1 - 2 \cdot f^*(\tau, \epsilon)) \cdot (1 + f^*(\tau, \epsilon) \cdot \left(\frac{\tau}{\epsilon} - 1 \right)) - (f^*(\tau, \epsilon) - f^*(\tau, \epsilon)^2) \cdot \left(\frac{\tau}{\epsilon} - 1 \right)}{(1 + f^*(\tau, \epsilon) \cdot \left(\frac{\tau}{\epsilon} - 1 \right))^2} = 0$$

which can be simplified to

$$\left(\frac{\tau}{\varepsilon} - 1\right) \cdot f^*(\tau, \varepsilon)^2 + 2 \cdot f^*(\tau, \varepsilon) - 1 = 0$$

which has as a positive solution:

$$f^*(\tau, \varepsilon) = \frac{1}{1 + \sqrt{\frac{\tau}{\varepsilon}}}$$

As well, we can compute now the value of r_{dep} :

$$r_{\text{dep}} = \frac{1}{1 + \frac{\tau}{\varepsilon} \cdot \frac{f^*(\tau, \varepsilon)}{1 - f^*(\tau, \varepsilon)}} = \frac{1}{1 + \sqrt{\frac{\tau}{\varepsilon}}} = f^*(\tau, \varepsilon)$$

This is an approximate solution, valid for small values of b_A and b_D , and small transition times before and after the intermittent phase; but it is still a good approximation if these conditions are not met.

