



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Capability Building in Sluggish Organizations

Kfir Eliaz, Ran Spiegler

To cite this article:

Kfir Eliaz, Ran Spiegler (2023) Capability Building in Sluggish Organizations. Management Science 69(3):1703-1713. <https://doi.org/10.1287/mnsc.2022.4445>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.



For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Capability Building in Sluggish Organizations

Kfir Eliaz,^{a,b} Ran Spiegler^{a,c,*}

^aTel Aviv University, 6997801 Tel Aviv, Israel; ^bUniversity of Utah, Salt Lake City, Utah 84112; ^cUniversity College London, London WC1E 6BT, United Kingdom

*Corresponding author

Contact: kfire@tauex.tau.ac.il,  <https://orcid.org/0000-0001-8988-0726> (KE); rani@tauex.tau.ac.il,  <https://orcid.org/0000-0001-9662-0236> (RS)

Received: September 14, 2021

Revised: January 12, 2022; March 28, 2022

Accepted: April 10, 2022

Published Online in Articles in Advance:
July 7, 2022

<https://doi.org/10.1287/mnsc.2022.4445>

Copyright: © 2022 INFORMS

Abstract. In order to thrive, organizations need to build and maintain an ability to meet unexpected external challenges. Yet many organizations are sluggish: their capabilities can only undergo incremental changes over time. What are the stochastic processes governing “routinely occurring” challenges that best prepare a sluggish organization for unexpected challenges? We address this question with a stylized principal-agent model. The “agent” represents a sluggish organization that can only change its capability by one unit at a time, and the “principal” represents the organization’s head or its competitive environment. The principal commits ex ante to a Markov process over challenge levels. We characterize the process that maximizes long-run capability for both myopic and arbitrarily patient agents. We show how stochastic, time-varying challenges dramatically improve a sluggish organization’s preparedness for sudden challenges.

History: Accepted by Joshua Gans, business strategy.

Funding: This work was supported by the Foerder Institute for Economic Research.

Keywords: training programs • organizational preparedness • sluggish adaptation • periodization • dynamic capabilities

1. Introduction

To thrive in the long run, organizations need to cope with unexpected challenges. In the private sector, a firm that can quickly adapt to new circumstances will gain a competitive advantage. The strategy literature refers to this trait as *dynamic capabilities*, which can be broadly described as “the firm’s ability to integrate, build, and reconfigure internal and external competences to address rapidly changing environments” (Teece et al. 1997, p. 516). In the public sector, organizations such as the military or emergency-response agencies must have the expediency to respond to unexpected crises such as wars, natural disasters, or epidemics. Their effectiveness is tested precisely when rare, unanticipated challenges arise.

However, organizations tend to be “sluggish” in their response to external challenges. As Hannan and Freeman (1984, p. 149) write, “Organizations are subject to strong inertial forces ... they seldom succeed in making radical changes in strategy and structure in the face of environmental threats.”¹ How, then, do organizations overcome their innate sluggishness and manage to build and maintain preparedness for random challenges?

We consider two perspectives into this question. First, an organization’s dynamic capabilities are shaped by its *natural environment*. As Eisenhardt and Martin (2000, p. 1110) note, “The pattern of effective dynamic capabilities depends upon market dynamism.” A

market environment that involves volatile changes in competition, technology, or regulation may be more conducive for building and maintaining dynamic capabilities.

Second, organizations can actively *simulate* random challenges via systematic *training programs*. This is particularly relevant for military and emergency-response organizations. Unable to quickly adapt to sudden real challenges, such an organization’s level of preparedness will gradually deteriorate unless its training regimen summons simulated ones. Lakoff (2007, p. 254) notes that, in the context of planning for emergencies, “since the probability and severity of such events cannot be calculated, the only way to avert catastrophes is to have plans to address them already in place and to have exercised for their eventuality—in other words, to maintain an ongoing capability to respond appropriately.”

The first perspective raises a natural question: What kind of stochastic environment is best for nurturing dynamic capabilities in sluggish organizations? The second perspective rephrases the question: What is the optimal way to “train” a sluggish organization to meet unexpected challenges?

This paper addresses both versions of this question with an economic theory approach.² We construct a stylized, dynamic *principal-agent* model in which the agent represents an organization that can only sluggishly respond to exogenous challenges. The agent

trades off the cost of failing to meet a challenge and the cost of maintaining a capability. The principal can be interpreted literally as a sluggish organization's head, who designs a dynamic "training program" that simulates challenges in order to develop and sustain the organization's capabilities. Alternatively, the principal can be viewed as a *fictitious* entity representing the organization's "regular" competitive environment, which generates time-varying challenges according to a stable stochastic pattern. This environment is similar to a crucible that forges the organization's dynamic capabilities, which would be tested against "irregular" challenges that arise unexpectedly and independently of the regular process.

More specifically, we represent capability and challenge levels by integers and measure them on the same scale. Time is discrete, and the agent adjusts its capability at every period *after* learning the current challenge level. Incremental sluggish adjustment means that the agent can change its capability (in either direction) only by *one unit* at any time period. Even when external stimuli change dramatically from one period to the next, capability adjusts slowly. For illustration, a firm may be unable to take immediate advantage of a large competitor's sudden demise. Conversely, when the firm faces prolonged low demand, its production capabilities will not disappear overnight but gradually decay ("use it or lose it").

The agent's adjustment process balances two opposing forces. On the one hand, maintaining capability is costly. For example, effective response to a technical challenge requires constant availability of skilled staff or computing resources. This maintenance cost exerts a downward force on capability. On the other hand, when capability falls below the challenge, this is recorded as a cost that exerts an upward pull on capability. Under the "fictitious principal" interpretation, this performance gap cost can be viewed as an opportunity cost of forgoing a source of revenues or as a reputational loss when a firm fails to rise up to a technical challenge. Under the literal "organizational training" interpretation, the cost may be part of a mechanism that incentivizes the organization's preparedness (e.g., performance-based bonuses and promotion prospects for the organization's members). However, we take it as given and focus on the dynamic training regime itself.

We assume that the principal commits *ex ante* to a Markov process that governs the evolution of the challenge level over time. The agent knows the process and monitors it throughout its evolution. We hold fixed the average challenge that the Markov process induces, such that this is a parameter of the principal's problem. Under the fictitious principal interpretation, this parameter is a characteristic of the organization's environment (similar in spirit to the distinction between "moderate"

and "high-velocity" markets made by Eisenhardt and Martin (2000)). Under the training program interpretation, the parameter may represent an allotted monthly amount of time for training. The principal's objective is to *maximize the agent's long-run capability*—defined as the *lowest* value it gets under the long-run distribution induced by the principal's Markov process and the agent's response.

To complete the model, we must specify the agent's planning horizon. We consider two extreme cases. The first case involves a *myopic* agent that balances the two cost components only for the current period. This effectively means that the agent's adjustment process is *mechanistic*: when capability is below (respectively, above) the current challenge, it goes up (respectively, down). That is, capability always changes incrementally in the direction of the current challenge level. The second case involves a *forward-looking* agent that minimizes the long-run average cost. Unlike the first case, here the agent's behavior is not mechanistic: it involves dynamic optimization that takes into account the agent's knowledge of the regular stochastic evolution of future challenges and the constraints on its own ability to adjust.³

Although the two cases require different proof methods, they share important commonalities. First, in a benchmark model with no sluggishness (in which the agent can adopt *any* capability at any period), the maximal capability that the principal can implement coincides with the average challenge level. The principal can attain this level with a constant challenge, which elicits the same response from the agent whether it is myopic or forward looking. Thus, the model is trivial in the flexible-adjustment benchmark.

Second, the principal's optimal Markov process under sluggish adjustment exhibits similar features in the two cases. It has two states: a "rest" state with a zero challenge level and a "high intensity" state. Some transitions between the two states are stochastic. For instance, in the myopic case (and for some parameter values in the forward-looking case), a high-intensity period is followed by another one with positive probability. However, the role of stochastic transitions is different in the two cases. In the myopic case, it ensures that the agent's long-run capability is insensitive to initial conditions. In the forward-looking case, it manages the agent's dynamic incentives. Randomization keeps the agent "on its toes," deterring it from lowering its capability during idle periods. As we discuss in Section 4, this effect has a subtle relation to models of auditing and inspection.

In both cases, the principal's optimal plan sustains a long-run capability that is considerably higher than what she could achieve in the flexible-adjustment benchmark. In the myopic case, long-run capability is nearly *twice* as large. In the forward-looking case, the

factor of increase can be *arbitrarily large* when the agent’s “maintenance cost” is small. Thus, our main theoretical insight is that in the presence of sluggish adjustment, a high-variance regular process that involves zero- and high-intensity phases enhances long-run capability. Furthermore, sluggish adjustment leads to an increase in the organization’s long-run capability. At the optimum, there will be periods in which the agent holds idle capabilities, which may appear wasteful to an outside observer. Nevertheless, this idleness is a feature of sluggish organizations’ optimal preparedness.⁴

2. The Model

We formulate our model as a principal-agent problem, in which the agent (it) is an organization or an organizational unit. For expositional focus, throughout Sections 2–4 we refer to the principal as a “trainer” (she) and adhere to the *literal* interpretation of the trainer as an organizational leader who wishes it to attain and maintain a high level of preparedness for unexpected external challenges. We discuss alternative interpretations in Section 6.

The trainer commits *ex ante* to a pair (P, f) , where P is a discrete-time, finite-state Markov process over some finite set of states S , and $f: S \rightarrow \mathbb{N}_+$ is an output function that assigns a challenge level to every state $s \in S$. The set of states S is *endogenous*: the trainer can choose a set of any finite size. We denote by s_t and d_t the state and challenge levels at period t , respectively. In keeping with the literal interpretation of the trainer, we sometimes refer to (P, f) as a *training program* and to d as *training intensity*.

We impose the following constraints on (P, f) . First, P is irreducible. This ensures that it has a unique invariant distribution λ_P and therefore enables us to talk about long-run average quantities unambiguously. Second,

$$\sum_{s \in S} \lambda_P(s) f(s) \leq \mu + \varepsilon, \quad (1)$$

where $\mu \geq 1$ is an integer and $\varepsilon \in (0, 1)$ can be arbitrarily close to 0. That is, the long-run average challenge level administered by the training program cannot exceed μ by more than a negligible amount. The approximate formulation of the constraint is due to μ getting integer values.

After the trainer chooses (P, f) at period 0, the agent chooses a nonnegative integer $m_t \in \{m_{t-1} - 1, m_{t-1}, m_{t-1} + 1\}$ at every $t = 1, 2, 3, \dots$. The agent’s choice at period t takes place *after* the realization of s_t . We refer to m_t as the agent’s *capability* at time t . Let $m_0 \in \mathbb{N}_+$ be the agent’s initial capability. The restricted choice set for m_t reflects sluggish adaptation. Note that the Markov process P does not condition on the agent’s

history of capability realizations. In particular, it is insensitive to the initial condition m_0 . We discuss this assumption in Section 6.2.

Define

$$C_t = cm_t + \max(0, d_t - m_t),$$

where $c \in (0, 1)$. This is the total cost that the agent incurs at period t . It consists of two terms. First, cm_t is the “maintenance cost” of the capability level. Second, the gap between m_t and d_t (when the latter is higher) represents a “performance gap” cost that arises when the agent’s capability is lower than the challenge it faces. Under the training program interpretation of the model, the performance gap cost captures a disutility that members of the organization experience when failing to meet the program’s challenges (foregone performance-based bonuses, thwarted promotion prospects, reputational damage, etc.). This disutility can be part of a broader incentive scheme. However, unlike most of the economic literature on organization design (e.g., Bolton and Dewatripont 2005), we take these incentives as given and focus on the problem of designing a dynamic, stochastic challenge.

The agent faces a trade-off whenever its current capability is not enough to meet the current challenge: increasing capability requires higher maintenance costs but lowers the performance gap. Our piecewise linear cost specification implies that moving up to the next capability rung reduces net costs by $1 - c$ in the current period, regardless of the agent’s current capability (as long as it is below the current challenge). Of course, a forward-looking agent still has to take into account that increasing m today will delay its ability to scale it back down in response to low future challenge levels.

We consider two alternative specifications of the agent’s intertemporal aggregation.

- *Myopic/mechanistic adjustment*: At every period $t \geq 1$, the agent chooses m_t to minimize C_t . That is, the agent is myopic: it does not take into account future costs. Because $c \in (0, 1)$, this immediately implies the following strategy for the agent:

$$m_t = \begin{cases} m_{t-1} + 1 & \text{if } d_t > m_{t-1}, \\ m_{t-1} & \text{if } d_t = m_{t-1}, \\ \max\{0, m_{t-1} - 1\} & \text{if } d_t < m_{t-1}. \end{cases} \quad (2)$$

That is, capability always moves in the direction of the current challenge level. This adjustment rule is *mechanistic*: it does not require the agent to know the trainer’s Markov process or to monitor the evolution of its state.

- *Forward-looking adjustment*: The agent knows the trainer’s choice of (P, f) . At every period t , it observes the realized state s_t before choosing m_t . The agent’s

objective is to minimize

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T C_t. \quad (3)$$

This is the long-run average cost that the agent incurs. The limsup criterion reflects the assumption that the agent is not only forward looking but also arbitrarily patient.

Under both rules of adjustment, the agent faces an extended Markov problem, in which the state at period t is (m_{t-1}, s_t) . Therefore, the agent has an optimal response that is also Markovian with respect to this extended state space. In the myopic case, this strategy is explicitly given by (2). We assume that the agent plays a *Markovian best response* in the forward-looking case as well. This ensures that the extended Markov process induced by the two parties' strategies has a unique invariant distribution over (d_t, m_t) . Consequently, all the limit quantities we will invoke going forward are well defined. In particular, let m^* be the lowest value that m takes beyond a sufficiently large t . This quantity is well defined, independent of the initial condition m_0 . We refer to m^* as *the lowest long-run capability* that is induced by the extended Markov process.

The *trainer's objective* is to maximize m^* subject to the feasibility constraint (1). A higher m^* means that the system has greater preparedness; that is, it can *consistently* meet bigger *actual* challenges that may arise unexpectedly, outside the regular process defined by (P, f) .

We make some additional comments on components of our model.

Comment on the Feasibility Constraint. Under the literal training program interpretation, we can regard (1) as a hard “budget constraint” that limits the resources (hours, ammunition) that the trainer can devote to training. We can also view μ as a parameter that the trainer controls at a cost. Our analysis characterizes the trainer's gross payoff as a function of μ , and a more complete analysis would trade off this payoff against the cost of increasing μ .

Comment on the Cost Function. The key assumption embodied by the cost function C_t is that if $d_t \geq m_{t-1}$, the maintenance cost saved when the agent lowers m by one unit is more than offset by the increase in the performance gap cost. In the case of a myopic agent, this feature leads to the mechanistic adjustment rule (2)—namely, m_t always chases d_t . In the case of a forward-looking agent, our analysis in Section 4 will also make use of the piece-wise linearity of the performance gap cost component. We conjecture that our results will remain intact if we replace the term $\max(0, d_t - m_t)$ by $g(\max(0, d_t - m_t))$, where g is an increasing, convex function satisfying $g(0) = 0$ and $g(1) > c$.

2.1. Benchmark: Completely Flexible Adjustment

Suppose the agent could choose *any* $m_t \in \mathbb{N}_+$ at every period, regardless of m_{t-1} . In particular, it could always choose m_t to minimize C_t . Recall that the agent chooses m_t *after* observing d_t . Therefore, it would set $m_t = d_t$ at every t . Under this flexible-adjustment rule, the long-run average of m_t coincides with the long-run average of d_t , which by assumption cannot exceed μ (more than negligibly). Therefore, the trainer cannot do better than play a constant strategy $d_t = \mu$ at every period, such that $m_t = \mu$ at every t as well. When the agent is sluggish, this deterministic process attains the same long-run capability of μ . The reason is that the agent will eventually reach this capability level and stay there indefinitely. The question is whether the trainer can outperform this benchmark with a nondegenerate Markov process.

3. Myopic/Mechanistic Adjustment

In this section we analyze the trainer's problem when the agent behaves according to the myopic/mechanistic adjustment model. The proofs of all formal results are relegated to Section 5.

Proposition 1. *Assume the agent follows the strategy given by (2). Then, we have the following.*

(i) *For any trainer strategy, the lowest long-run capability is at most $2\mu - 1$.*

(ii) *This upper bound can be implemented by the following (P, f) . The Markov process P has two states, H and L, and a transition matrix given by*

$$\Pr(s_t \rightarrow s_{t+1}) \begin{array}{cc} & \begin{array}{c} L \quad H \end{array} \\ \begin{array}{c} L \\ H \end{array} & \begin{array}{cc} 0 & 1 \\ \beta & 1 - \beta \end{array} \end{array}$$

where β is arbitrarily close to 1. The output function is $f(H) = 2\mu$ and $f(L) = 0$. In the $\beta \rightarrow 1$ limit, the invariant capability distribution assigns probability $\frac{1}{2}$ to $m = 2\mu$ and $m = 2\mu - 1$.

Thus, a slightly perturbed cyclic training program can dramatically increase the long-run capability of a sluggish agent, relative to the flexible-adjustment benchmark. When μ is large—corresponding to a very sluggish agent, given that we normalized the adjustment increment to 1—the increase is by a factor of nearly 2.

The training regime approximately consists of alternating periods of “high intensity” ($d = 2\mu$) and “rest” ($d = 0$). After a period of high-intensity training, there is a small chance $1 - \beta$ that the high-intensity episode will be repeated. This stochastic perturbation ensures that the set of capability values $\{2\mu, 2\mu - 1\}$ is absorbing: the agent will reach it in finite time with probability 1, regardless of m_0 . The *only* role of randomness is thus to ensure that the agent's long-run behavior is insensitive to initial conditions. Note that the long-run

average intensity under the trainer’s strategy is $2\mu/(1 + \beta)$. Therefore, for every $\varepsilon > 0$, we can select β to be sufficiently close to 1 such that average intensity will not exceed $\mu + \varepsilon$.

The intuition for the result is that changes in m depend only on the *sign* of $d - m$, whereas the trainer’s budget constraint is expressed in terms of the *average* of d . The contrast between the *cardinal* constraint and the *ordinal* adjustment rule—which itself is a consequence of the agent’s sluggishness—is the key to our result. The most economical way to get the agent’s capability to go up at period t is to set $d_t = m_{t-1} + 1$; the most economical way to bring it down is to set $d_t = 0$. In the long run, because the agent’s capability moves around in increments of one unit, m goes up and down with equal frequencies. This explains the approximate factor 2 by which the trainer can increase long-run capability, relative to the flexible-benchmark μ .

4. Forward-Looking Adjustment

In this section we characterize the solution to the trainer’s problem when the agent is forward looking. For expositional convenience, we assume μ/c is an *integer*.

Proposition 2. *Assume the agent evaluates cost streams by (3). Then, we have the following.*

- (i) *The lowest long-run capability is at most $\mu/c - 1$.*
- (ii) *This upper bound can be implemented by the following (P, f) . The Markov process P has two states, H and L , and a transition matrix given by*

$$\Pr(s_t \rightarrow s_{t+1}) \begin{array}{cc} & \begin{array}{c} L \quad H \end{array} \\ \begin{array}{c} L \\ H \end{array} & \begin{array}{cc} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{array} \end{array}$$

where $\alpha = 1$ if $c \geq \frac{1}{2}$, $\beta = 1$ if $c < \frac{1}{2}$, and $\alpha/(\alpha + \beta)$ is arbitrarily close to c from above. The output function is $f(H) = \mu/c$ and $f(L) = 0$. In the $\alpha/(\alpha + \beta) \rightarrow c$ limit, the invariant capability distribution assigns probability c to $m = \mu/c$ and probability $1 - c$ to $m = \mu/c - 1$.

When $c < \frac{1}{2}$, the upper bound on the agent’s lowest long-run capability is higher than in the myopic case. Moreover, it gets arbitrarily high when $c \rightarrow 0$. As c gets closer to 1, the highest minimal long-run capability approaches the flexible-agent benchmark μ .⁵ Note that the long-run average intensity under the trainer’s strategy is

$$\frac{\alpha}{\alpha + \beta} \cdot \frac{\mu}{c}$$

For every $\varepsilon > 0$, we can set $\alpha/(\alpha + \beta)$ to be sufficiently close to c , such that the average intensity does not exceed $\mu + \varepsilon$.

The Markov process that attains the upper bound is similar to the one in Section 3. However, the reasoning behind the result is different. Because the mechanistic agent of Section 3 responds only to the *current*

realization of d , the only role of randomization in that case is to ensure insensitivity to initial conditions. By contrast, a forward-looking, patient agent responds to the trainer’s *entire* continuation strategy. Randomization serves as an incentive to keep the agent “on its toes” and deter it from lowering its level of preparedness during periods of rest. In particular, when $c < \frac{1}{2}$, a rest period is followed by another one with probability approximately equal to $(1 - 2c)/(1 - c)$. Hence, the trainer’s optimal program allows for a streak of $d = 0$ realizations. When this happens, the agent does not lower its capability below $\mu/c - 1$ because it takes into account the future loss $d - m$ in the event that d switches from 0 to μ/c .

The trainer designs the transition probabilities such that the agent’s intertemporal trade-offs lead it to be nearly indifferent between lowering its capability and remaining at $m = \mu/c - 1$. By contrast, the mechanistic agent cannot be made indifferent when faced with a streak of $d = 0$ realizations: it repeatedly lowers its capability. This difference enables the trainer to achieve higher long-run capability when the agent is forward looking, as long as $c < \frac{1}{2}$.

To further elucidate why randomization is necessary, consider the following example, which shows that the minimal long-run capability attained by the optimal stochastic strategy *cannot* be sustained by a particular deterministic strategy with the same long-run distribution over d . Suppose $\mu = 4$ and c is slightly below $\frac{4}{11}$. Then, the optimal training strategy of Proposition 2 induces an invariant distribution that assigns probability $\frac{4}{11}$ to $d = 11$ and probability $\frac{7}{11}$ to $d = 0$. This strategy sustains a minimal long-run capability level of $m = 10$.

Now consider a deterministic strategy that induces the same long-run frequencies of d . The strategy follows an 11-period cycle consisting of four consecutive periods of $d = 11$ and seven consecutive periods of $d = 0$. If the agent plays $m = 11$ when $d = 11$ and $m = 10$ when $d = 0$ —as it does against the strategy presented in Proposition 2—the minimal long-run capability is $m = 10$. Moreover, this strategy is optimal for the agent among all strategies that induce this minimal long-run capability. However, given the predictable evolution of d under the cyclic deterministic strategy, a forward-looking agent can do better. Suppose that it plays the following sequence of m against the cyclic sequence of d :

d	11	11	11	11	0	0	0	0	0	0	0
m	11	11	11	10	9	8	7	7	8	9	10

Compared with the benchmark strategy of playing $m = 11$ (10) against $d = 11$ (0), the agent saves approximately

$$c \cdot (1 + 1 + 2 + 3 + 3 + 2 + 1) - 1 \approx \frac{41}{11}$$

per cycle. It follows that the agent's best reply to the cyclic deterministic strategy leads to a minimal long-run capability below $m = 10$.

This example highlights a key role of the stochasticity of the trainer's optimal strategy in the forward-looking case. The fact that there is always a chance that the agent will face a big challenge following a rest period incentivizes the agent not to lower its capability. By contrast, the predictable nature of the cyclic deterministic strategy allows the agent to gradually lower its capability and gain it back by the time the big challenge arrives. In particular, it is profitable for the agent to lower its capability already in the final period of the high-intensity phase of the cycle, even though this involves a costly performance gap during that period, because this is more than offset by the cumulative maintenance cost saving over the cycle.

4.1. A Random-Audit Analogy

The literal training program interpretation of our model invites an analogy to models of auditing or inspections. Think of d_t as the audit's intensity, such that the budget constraint (1) represents limited resources for auditing. The organization adapts its capability to the auditing regime because of underlying incentives, which are captured by the cost function C .

The idea that optimal inspection may involve random audits is familiar in game theory and economics: when auditing is costly, making it unpredictable deters the agent from shirking (e.g., Lazear (2006), Eeckhout et al. (2010), Varas et al. (2020), and Solan and Zhao (2021)).

Despite this analogy, there is a crucial difference between conventional models of auditing and the present model, where the agent's move at period t is taken *after* d_t is realized. That is, the agent can condition its action on the principal's "auditing" choice. In a standard auditing model, this would entirely rob audits of their potency. Indeed, in the flexible-agent benchmark described in Section 2.1, randomizing over d is useless for the trainer. What makes randomization valuable in our model is the element of sluggish adjustment. Successive periods of shirking can magnify the agent's failure at an audit, and sustained effort may be required to rebuild the ability to pass it. To use a metaphor, a restaurant chef may learn on Monday that a famous food critic will come for dinner on Friday. The chef cannot realistically raise the quality of the restaurant staff and recipes in this tight time frame but would have had been able to if given more notice.

What our results demonstrate is that when faced with a sluggish agent, an "auditor" can use randomization as if she were facing a simultaneously moving, flexible

agent. However, as Propositions 1 and 2 demonstrate, this is *not* an exact equivalence: the details of the optimal random "auditing" strategy depend on the agent's patience; in addition, the optimal strategy does not involve independent and identically distributed randomization.

4.2. Basic Ideas Behind the Proof of Proposition 2

In part (i) we actually prove something stronger than the stated result: to attain a *strictly positive* minimal long-run capability, the *average* long-run capability cannot exceed $\mu/c - 1 + c$. The Markov process we construct in part (ii) approximates this upper bound. This means that among all trainer strategies that attain the minimal long-run capability of $\mu/c - 1$, this process cannot be outperformed in terms of average capability.

The proof of part (i) proceeds in several steps. First, note that by playing a constant $d = \mu$, the trainer can attain a long-run capability of μ . Therefore, the trainer can attain a minimal long-run capability that is at least as large as μ . Hence, the invariant distribution over (m, d) —induced by an optimal trainer strategy and an agent's best reply—satisfies $\Pr(m > 0) = 1$.

Second, we establish a lower bound on the long-run frequency of positive training intensity: under the invariant distribution induced by the two parties' strategies, $\Pr(d > 0) \geq c$. To prove this, we consider the following possible deviation by the agent: pick a history in which m is at its lowest long-run value (which is positive, as we saw), move one notch below the original plan, and afterward, proceed as if the deviation never took place. The piecewise linearity of the cost function enables a simple calculation of the net long-run profit from this deviation: it saves c per period but raises the performance gap cost by one unit whenever $d \geq m$ under the original strategy. When the agent is forward looking, this deviation is unprofitable only if $\Pr(d \geq m) \geq c$. Because $\Pr(m > 0) = 1$, we have that $\Pr(d > 0) \geq \Pr(d \geq m) \geq c$.

The third and final step of the proof shows that the long-run average capability cannot exceed $\mu/c - 1 + c$. If this were not true, then the average long-run cost would exceed $\mu - c(1 - c)$. But then, using the previous step, we obtain that the following deviation is profitable for the agent: descend all the way to $m = 0$ and play $m = 1 (d > 0)$ thereafter. The upper bound on the lowest long-run capability then immediately follows.

The proof of part (ii) begins by noting that the agent has a best reply to the trainer's strategy that induces two (and therefore adjacent) long-run values of m (this is a consequence of the fact that P has two states and that $\alpha\beta = 0$). We then show that by the piecewise

linearity of the agent's cost function and the condition on α, β, c , the two long-run capability values are μ/c and $\mu/c - 1$. The induced long-run average capability is then $\mu/c - 1 + c$.

5. Proofs

This section provides proofs of the formal results in Sections 3 and 4.

5.1. Proof of Proposition 1

5.1.1. Proof of Part (i). Consider an arbitrary strategy for the trainer. Let $(m_{t-1}, d_t)_{t=1,2,\dots}$ be a possible sample path that results from the extended process. The long-run frequency of every (m, d) in the sample path, denoted as $\lambda(m, d)$, coincides with the probability of this pair according to the invariant distribution induced by the two parties' strategies. Let X be the set of recurrent pairs (m, d) in the sample path. Partition X into three classes:

$$\begin{aligned} X^+ &= \{(m, d) \in X \mid d > m\}, \\ X^- &= \{(m, d) \in X \mid d < m\}, \\ X^0 &= \{(m, d) \in X \mid d = m\}. \end{aligned}$$

The proof now proceeds by a series of steps. Recall that we use the notation $d(s)$ as a substitute for $f(s)$.

Step 1. The distribution λ satisfies

$$\sum_{(m,d) \in X^+} \lambda(m, d)(m + 1) = \sum_{(m,d) \in X^-} \lambda(m, d)m. \quad (4)$$

Consider some period t along the sample path such that $(m_t, d_{t+1}) \in X^+$. By definition, this pair is recurrent. Therefore, m_t must be visited again in some later period. Let $t' + 1$ be the *earliest* such period (although $m_{t'+1} = m_t$, we do not require $d_{t'+2} = d_{t+1}$). Because $(m_t, d_{t+1}) \in X^+$, $m_s > m_t$ for every $s = t + 1, \dots, t'$. Therefore, by the definition of t' , it must be the case that $m_{t'} = m_t + 1$ and $(m_{t'}, d_{t'+1}) \in X^-$. In other words, because the trajectory of m is upward at t , it must be downward at t' by the definition of this period.

We have thus defined a *one-to-one* mapping from periods t for which $(m_t, d_{t+1}) \in X^+$ to periods t' for which $(m_{t'}, d_{t'+1}) \in X^-$, such that $m_{t'} = m_t + 1$. It follows that

$$\begin{aligned} &\lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T \mathbf{1}[(m_t, d_{t+1}) \in X^+] \cdot (m_t + 1)}{T} \\ &= \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T \mathbf{1}[(m_t, d_{t+1}) \in X^-] \cdot m_t}{T}; \end{aligned}$$

we can rewrite this equation as (4), because

$$\lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T \mathbf{1}[(m_t, d_{t+1}) = (m, d)]}{T} = \lambda(m, d). \quad \square$$

Step 2. The average long run m is at most 2μ (approximately).

The long-run average of m induced by the trainer's strategy can be written as

$$\begin{aligned} \mathbb{E}(m) &= \sum_{(m,d) \in X^+} \lambda(m, d)m + \sum_{(m,d) \in X^-} \lambda(m, d)m \\ &\quad + \sum_{(m,d) \in X^0} \lambda(m, d)m. \end{aligned} \quad (5)$$

By the feasibility constraint,

$$\sum_{(m,d) \in X^+} \lambda(m, d)d + \sum_{(m,d) \in X^-} \lambda(m, d)d + \sum_{(m,d) \in X^0} \lambda(m, d)d \lesssim \mu.$$

By definition, $d \geq m + 1$ for every $(m, d) \in X^+$, $d \geq 0$ for every $(m, d) \in X^-$, and $d = m$ for every $(m, d) \in X^0$. Therefore,

$$\begin{aligned} &\sum_{(m,d) \in X^+} \lambda(m, d)(m + 1) + \sum_{(m,d) \in X^-} \lambda(m, d) \cdot 0 \\ &\quad + \sum_{(m,d) \in X^0} \lambda(m, d)m \lesssim \mu. \end{aligned}$$

This means that

$$\begin{aligned} &\sum_{(m,d) \in X^+} \lambda(m, d)m \leq \sum_{(m,d) \in X^+} \lambda(m, d)(m + 1) \lesssim \mu \\ &\quad - \sum_{(m,d) \in X^0} \lambda(m, d)m. \end{aligned}$$

By (4), it follows that

$$\sum_{(m,d) \in X^-} \lambda(m, d)m \lesssim \mu - \sum_{(m,d) \in X^0} \lambda(m, d)m$$

as well. Plugging the last two inequalities in (5), we obtain

$$E(m) \lesssim 2\mu - \sum_{(m,d) \in X^0} \lambda(m, d)m \leq 2\mu. \quad \square$$

Step 3. The minimal long run m is at most $2\mu - 1$.

Suppose the long-run distribution over d is degenerate at some d^* . Therefore, $d^* \lesssim \mu$. The agent's myopic best reply implies that, eventually, its capability coincides with d^* . It follows that to reach a minimal long-run capability above μ , the long-run distribution over d must assign positive probability to at least two values. This means there are infinitely many periods t in which $d_t \neq m_{t-1}$. By myopic best replying, this precludes the possibility that the long-run distribution over m is degenerate. Because the long-run average of m cannot exceed 2μ by more than an infinitesimal amount, there must be infinitely many periods t in which $m_t \leq 2\mu - 1$. \square

5.1.2. Proof of Part (ii). Consider the trainer's strategy described in part (ii) of the statement of the result. As long as $\beta \in (0, 1)$, the Markov process over m that is induced by the strategy, and the agent's best reply (given by Step 1) has a unique invariant distribution, with $m = 2\mu$ and $m = 2\mu - 1$ being the only recurrent

capability values. The reason is that if $m_t > 2\mu$, $m_{t+1} = m_t - 1$ with certainty; if $m_t < 2\mu - 1$, there is a positive probability that there will be a streak of realizations $d = 2\mu$ such that m will keep adjusting upward until it reaches $m = 2\mu$; and finally, if $d_t = 0$, then $d_{t+1} = 2\mu$ for sure, which means that once m hits 2μ and later goes down to $2\mu - 1$, it will return to 2μ immediately in the next period. As the exogenous upper bound on average intensity gets arbitrarily close to μ , β can be made arbitrarily close to 1. In the $\beta \rightarrow 1$ limit, the invariant distribution over m assigns probability $\frac{1}{2}$ to each of the values $m = 2\mu$ and $m = 2\mu - 1$. \square

5.2. Proof of Proposition 2

5.2.1. Proof of Part (i). Let p be the unique invariant distribution over (d_t, m_t) that results from the trainer’s strategy and the agent’s best-replying strategy. (Note the different time subscripts of d and m , compared with the proof of Proposition 1; our different notation highlights this difference.) We abuse notation and write $p(d)$, $p(m)$, and $p(d | m)$ to represent marginal and conditional distributions induced by p . As in the myopic agent case, we first derive an upper bound on the expected capability according to p , which we use to derive the upper bound on the minimal long-run capability. Then, we show how to implement this upper bound.

In Section 2, we saw that the trainer can implement a minimal long-run capability of at least μ (by playing $d = \mu$ at every period). Therefore, we take it for granted that the minimal value of m in the support of p is at least $\mu \geq 1$.

Step 1. The invariant distribution p satisfies $p(d > 0) \geq c$.

Consider the following deviation by the agent. Pick some period t history for which $m_{t-1} \geq 1$ is at the lowest value according to p . Therefore, $m_t = m \in \{m_{t-1}, m_{t-1} + 1\}$. At this history, the agent deviates to $m'_t = m - 1$. The agent subsequently behaves according to its original strategy as if the deviation did not occur.

This deviating strategy induces an invariant distribution p' such that for every (d, m) in the support of p , $p'(d, m - 1) = p(d, m)$. Therefore, the deviation saves c at every period but raises costs by one unit per period whenever $d \geq m$ under the original strategy. In order for this deviation to be unprofitable for an arbitrarily patient agent, it must be the case that $p(d \geq m) \geq c$. Because $m > 0$ with probability 1, $p(d > 0) \geq p(d \geq m)$; hence $p(d > 0) \geq c$. \square

Step 2. The expectation of m according to p is at most $\mu/c - 1 + c$.

Assume the contrary. Then, the agent’s average long-run cost exceeds

$$c \cdot \left[\frac{\mu}{c} - 1 + c \right] = \mu - c(1 - c).$$

Now consider a deviation to the following strategy. Descend from m_0 to $m = 0$, and then implement the

following rule: $m_t = 0$ whenever $d_t = 0$ and $m_t = 1$ whenever $d_t > 0$. When the agent is arbitrarily patient, the average long-run cost from this strategy is approximately

$$p(d = 0) \cdot 0 + p(d > 0) \cdot \left[c + \sum_{d>0} p(d | d > 0)d - 1 \right] \\ \approx p(d > 0)(c - 1) + \mu.$$

Because $c < 1$, Step 1 implies that

$$p(d > 0)(c - 1) + \mu < \mu - c(1 - c)$$

such that the deviation is profitable—a contradiction. \square

Step 3. The minimal long-run capability is at most $\mu/c - 1$.

Because μ/c is an integer, $\mu/c - 1 + c$ is not an integer. Hence, in order for the average long-run cost to be weakly below $\mu/c - 1 + c$, the minimal long-run capability cannot exceed $\mu/c - 1$.⁶ \square

5.2.2. Proof of Part (ii). Consider the strategy described in the statement of part (ii). Our objective is to show that given this strategy, there is a best reply for the agent such that for every sufficiently high t , $m_t = \mu/c$ whenever $s_t = H$ and $m_t = \mu/c - 1$ whenever $s_t = L$.

Because the agent faces a Markovian decision problem with an extended state space (s, m) , there exists a best reply that is Markovian with respect to this state space. To derive such a best reply, we proceed in four steps.

Step 1. There is no best reply in which the invariant distribution assigns probability one to a single m .

Proof. Assume the contrary. If $m < \mu/c$, then it is profitable for the agent to deviate to a strategy that plays $m + 1$ whenever $s = H$ and m whenever $s = L$. Likewise, if $m > 0$, it is profitable for the agent to deviate to a strategy that plays m whenever $s = H$ and $m - 1$ whenever $s = L$. \square

Step 2. The set of recurrent values of m (according to the unique invariant distribution induced by the two parties’ strategies) is a set of consecutive numbers $\underline{m}, \underline{m} + 1, \dots, \bar{m}$, where $\bar{m} \leq \mu/c$.

Proof. The agent’s sluggishness implies that if the agent visits two nonadjacent capabilities m and m' , then it must also visit every m'' between them. Therefore, if m and m' are recurrent, so is m'' . Suppose $\bar{m} > \mu/c$. Then, there is a profitable deviation for the agent that instructs to remain at $\bar{m} - 1$ whenever the original strategy instructs to switch to \bar{m} . \square

Step 3. There is a best reply that induces an invariant distribution that assigns positive probability to exactly two values of m .

Proof. Consider the invariant distribution over (d, m) induced by the trainer’s strategy and the agent’s best reply. By Step 1, $\bar{m} - \underline{m} \geq 1$. If $\bar{m} - \underline{m} = 1$, we are done. Therefore, assume $\bar{m} - \underline{m} > 1$. There are two cases to consider.

First, let $\alpha = 1$ (this fits the case of $c \geq 1/2$). This means that whenever $s = L$, the state switches immediately to $s = H$ in the next period. Consider the top two values of m in the invariant distribution—namely, \bar{m} and $\bar{m} - 1$. By Step 2, $\bar{m} \leq \mu/c$. Moreover, when $s = L$ (at which d attains its lowest value according to the trainer’s strategy), the agent strictly prefers $\bar{m} - 1$ to \bar{m} . Consider some t for which $m_t = \bar{m}$ (there are infinitely such periods because \bar{m} is recurrent). If $s_{t+1} = L$, the agent necessarily switches to $m_{t+1} = \bar{m} - 1$. If, on the other hand, $s_{t+1} = H$, we need to consider two possibilities.

- Suppose that when $s_{t+1} = H$, it is not optimal for the agent to play $m_{t+1} = \bar{m}$. That is, the agent switches from $m_t = \bar{m}$ to $m_{t+1} = \bar{m} - 1$ for any realization of s_{t+1} . But this also means that if $m_{t'} = \bar{m} - 1$ at some period t' and $s_{t'+1} = H$, it cannot be optimal for the agent to switch to $m_{t'+1} = \bar{m}$. The reason is that, by revealed preference, the agent prefers being at $\bar{m} - 1$ to being at \bar{m} when the state is H . And because we already saw that the agent prefers being at $\bar{m} - 1$ to being at \bar{m} when the state is L , this means that the agent will *never* switch from $\bar{m} - 1$ to \bar{m} , contradicting the definition of \bar{m} as a recurrent state.

- Suppose that when $s_{t+1} = H$, it is optimal for the agent to play $m_{t+1} = \bar{m}$. This reveals a weak preference for \bar{m} over $\bar{m} - 1$ when the state is H . Therefore, there is a best reply for the agent that prescribes $m_{t+1} = \bar{m}$ whenever the extended state (s_{t+1}, m_t) is $(H, \bar{m} - 1)$ or (H, \bar{m}) . We already saw that when the extended state is (L, \bar{m}) , the agent switches to $\bar{m} - 1$. Because $\alpha = 1$, this means that we have constructed a best reply for the agent such that once it reaches \bar{m} , it will only visit \bar{m} and $\bar{m} - 1$ from that period on, contradicting the assumption that there are additional recurrent values of m .

Thus, we have ruled out the possibility that $\bar{m} - \underline{m} > 1$ when $\alpha = 1$. Now suppose $\beta = 1$ (this fits the case of $c \leq 1/2$). An analogous argument establishes that there is a best reply for the agent that induces an invariant distribution with only two recurrent capability values, \underline{m} and $\underline{m} + 1$.

It follows that we can restrict attention to strategies of the agent that induce an invariant distribution, which assigns positive probability to precisely two consecutive capability values, m and $m - 1$, where $0 < m \leq \mu/c$. □

Step 4. There is a best reply for the agent that induces an invariant distribution on the capability values μ/c and $\mu/c - 1$.

Proof. Given Step 3, it is clearly optimal for the agent to be at m when $s = H$ and at $m - 1$ when $s = L$. In addition, when $m > \mu/c$ ($m < \mu/c - 1$), the agent clearly wants to move downward (upward).

The invariant distribution of the trainer’s two-state Markov process assigns probability $\alpha/(\alpha + \beta)$ to state H and $\beta/(\alpha + \beta)$ to state L . Therefore, because the agent is arbitrarily patient, its long-run expected payoff is approximately

$$-\frac{\alpha}{\alpha + \beta} \cdot \left(cm + \frac{\mu}{c} - m \right) - \frac{\beta}{\alpha + \beta} \cdot c(m - 1).$$

It is now easy to see that given that $\alpha/(\alpha + \beta) > c$, this expression increases with m , such that the optimal value of m is μ/c . The expected value of m according to this strategy is

$$\frac{\alpha}{\alpha + \beta} \cdot \frac{\mu}{c} + \frac{\beta}{\alpha + \beta} \cdot \left(\frac{\mu}{c} - 1 \right),$$

which is arbitrarily close to the upper bound. □

6. Discussion

In this section we discuss the interpretation and two features of our model.

6.1. The Competitive-Environment Interpretation

In Sections 2–4, we adopted the literal interpretation of the principal as an actual trainer. Accordingly, we interpreted the trainer’s Markov process as a training program. However, recall that in the introduction we proposed an alternative interpretation of the principal as a fictitious entity that represents the organization’s *competitive environment*. From this point of view, the Markov process is an autonomous process that generates regular challenges. Our results then suggest that when organizational adjustment is sluggish, the processes that maximize the organization’s dynamic capabilities are those that have high variance, fluctuating between two states of zero and high intensity. Organizations that adapt to such environments are also better prepared for unexpected, irregular challenges—certainly in comparison with static environments in which challenges remain at a constant level. It should be emphasized that this distinction arises when the organization is sluggish; as we saw in Section 2.1, it vanishes in the flexible-adjustment case.

6.2. Conditioning on the Agent’s Past Capability

In our model, the trainer does not condition the choice of d on past realizations of m . There are several reasons for this modeling decision. To begin with, under one interpretation (discussed in the previous subsection), (P, f) is an autonomous process over “states of nature,” and allowing this process to condition on m would be nonsensical. However, even under the

literal interpretation of (P, f) as a training program, there are good reasons to rule out conditioning on m .

First, in the myopic/mechanistic case, monitoring m is irrelevant because the agent's adjustment rule is not forward looking and hence does not respond to threats to change the evolution of d if m fails to meet some target. Therefore, in what follows, we focus on the case of a forward-looking agent. For expositional ease, we will let $\varepsilon = 0$ when discussing the trainer's budget constraint (1).

Second, the trainer's gain (in terms of her objective function) from conditioning on m can only be modest. Recall that the max-min capability in the case of forward-looking adjustment is $\mu/c - 1$. By playing $m_t = 0$ for every sufficiently large t , the agent can guarantee a long-run cost of $\mathbb{E}(d) \leq \mu$ because of the trainer's budget constraint (where $\mathbb{E}(d)$ is the long-run average d). Therefore, the highest minimal capability that the trainer can hope to sustain with a more complex policy is μ/c . This means a maximal gain of one capability unit. This gain may be outweighed by the implicit cost of a more complex training program that monitors m .

Finally, training programs that condition on m and attain a minimal capability of μ/c are not credible, in the following sense. In order to incentivize the agent not to deviate to a capability below μ/c , the trainer needs to threaten the agent that such a deviation would trigger a "punishment" phase in which $d > \mu/c$ with some probability. Impose the restriction that Constraint (1) holds after every history. The following is an example of a punishment phase that satisfies these properties: At every period, the trainer plays $d > \mu/c$ with probability $p < c$ and $d = 0$ with probability $1 - p$, such that $pd = \mu$. By the same methods as in the proof of part (ii) of Proposition 2, it can be shown that the agent's best reply to the punishment phase is $m_t = \mathbf{1}(d_t > 0)$. Clearly, near-zero long-run capability is a bad outcome for our trainer. Therefore, if the trainer is interested in meeting her objective after any history—including those that result from trembles by the agent—she would not want to use a strategy that relies on such punishments.

Although we demonstrated this argument for a particular punishment strategy, the argument holds for any policy that sustains a long-run capability of $m = \mu/c$ on the equilibrium path (and satisfies (1) after every history). It follows that designing a Markov training policy that does not condition on the history of m entails no loss of generality if the trainer wants to maximize the minimal level of m both on and off the equilibrium path. This also means that the trainer's optimal policy in Section 4 is robust to relaxing the assumption that the trainer commits to her policy ex ante, whereas policies that condition on m and implement a minimal capability above $\mu/c - 1$ fail this criterion.

6.3. The max min Criterion

In our model, the trainer's objective is to maximize the agent's minimal long-run capability. Alternatively, we could use the long-run average m as a criterion. However, this criterion is less attractive in our context because it does not reflect the ideas of "preparedness" and "dynamic capabilities." In particular, the average criterion allows 0 to be a recurrent value for m , which means that the agent will sometimes be completely unprepared for any surprise challenge.

A by-product of our analysis in Section 3 is that in the myopic case, 2μ is an upper bound on the average long-run capability that the trainer can attain. It can be shown that this upper bound can be approximated arbitrarily well, but this must come at the price of arbitrarily long recurrent stretches of $m = 0$ (which are compensated for by periods in which m reaches arbitrarily high values). Obviously, such paths imply that the agent's minimal long-run capability is 0. By comparison, the process we constructed in Section 3 induces an average long-run capability of approximately $2\mu - \frac{1}{2}$ and a minimal long-run capability of $2\mu - 1$.

A similar diagnosis pertains to the forward-looking case (let $\varepsilon = 0$, for the sake of argument). An upper bound on the average long-run capability is μ/c . The reason is that if average m exceeds this value, it implies that the agent's average long-run cost is above μ . However, the agent can ensure an average cost of μ by always playing $m = 0$; hence a long-run capability in excess of μ/c is inconsistent with the agent's best replying. We believe that, as in the myopic case, this upper bound can be approximated arbitrarily well, at the same price of long stretches of $m = 0$. By comparison, the process we constructed in Section 4 induces an average long run m of approximately $\mu/c - 1 + c$ and a minimal long-run m of $\mu/c - 1$. It follows that many combinations of the minimal and average criteria would lead to the same result.

6.4. A "Bodybuilding" Interpretation

Taking the training program interpretation even more literally, our model can be read through a very different lens. Instead of an organization, the agent in our model can be viewed as a *physiological* system, such as a muscle or a cognitive function. The capability m thus stands for things such as muscle mass. The trainer engages in physical or cognitive training. The variable d represents a physical or cognitive challenge, and the system adjusts its capability m in a way that trades off the *energy* cost of maintaining capability against a cost of failing to meet the challenge. For illustration, when m_t represents muscle mass, and cm_t captures the caloric cost of maintaining it, whereas $d_t - m_t$ may represent physical damage (inflammation, torn tissue) as a result of excessive stress that occurs when training intensity exceeds the muscle's capability. The

system's energy-saving motive creates an agency problem, because the trainer does not share this motive. Note that under this interpretation, the agent is the biological system itself, *not* the person of which it is part of.

In light of this physiological interpretation, our optimal training plan may be viewed as a stochastic variant on “periodization” training techniques familiar from *exercise physiology*. Numerous studies have documented the success of periodization in terms of increased muscle mass and athletic performance (Issurin 2010, Kiely 2012, Bompa and Buzzichelli 2018, Kiely et al. 2019). Although the literature offers biological explanations for the superiority of cyclical training (e.g., Issurin 2019), our results provide a complementary perspective by deriving the effectiveness of stochastic periodization as a logical conclusion of sluggish adaptation (resulting from rational cost-benefit calculus) to random physical stimuli. To our knowledge, this perspective is new: we are not aware of life science studies that examined the hypothesis that building and maintaining long-run physiological or neurological capabilities involves optimizing mechanisms. This theoretical conclusion does not require knowledge of details of the adjustment mechanism of the system in question (although it does make use of a number of simplifying assumptions). Therefore, it might be relevant for various biological systems that exhibit sluggish adjustment.

Acknowledgments

This is a substantial revision of papers formerly titled “Anabolic Persuasion” and “Training a Sluggish System.” The authors thank Martin Cripps, Michael Crystal, Tuval Danenberg, Israel Halperin, Nathan Hancart, Ron Peled, Ariel Rubinstein, Mickey Scheinowitz, the editor, and referees of this journal and seminar audiences for helpful comments.

Endnotes

¹ Other articles articulate similar ideas; see Hollnagel et al. (2006) for a collection of papers.

² For very different economics-based approaches to the subject of dynamic capabilities, see Sutton (2012) and Gans (2017).

³ Characterizing the optimum for an intermediately patient agent is a challenging technical problem that remains an open question.

⁴ See Iliev and Welch (2013) for a complementary rationale for “optimal idleness.”

⁵ By requiring μ/c to be an integer, we effectively rule out the case that c is *arbitrarily* close to 1. In that case, the trainer would be unable to outperform the flexible-agent benchmark of μ .

⁶ The proof of this step utilizes the convenient assumption that μ/c is an integer. An alternative proof that does not rely on this assumption is analogous to Step 3 in the proof of Proposition 1(i).

References

- Bolton P, Dewatripont M (2005) *Contract Theory* (MIT Press, Cambridge, MA).
- Bompa T, Buzzichelli CA (2018) *Periodization: Theory and Methodology of Training*, 6th ed. (Human Kinetics, Champaign, IL).
- Eeckhout J, Persico N, Todd P (2010) A theory of optimal random crackdowns. *Amer. Econom. Rev.* 100(3):1104–1135.
- Eisenhardt KM, Martin JA (2000) Dynamic capabilities: What are they? *Strategic Management J.* 21(10–11):1105–1121.
- Gans JS (2017) Negotiating for the market. Furman J, Gawer A, Silverman BS, Stern S, eds. *Entrepreneurship, Innovation, and Platforms*, Advances in Strategic Management, Vol. 37 (Emerald Group Publishing, Bingley, UK), 3–36.
- Hannan M, Freeman J (1984) Structural inertia and organizational change. *Amer. Sociol. Rev.* 49(2):149–164.
- Hollnagel E, Woods D, Leveson N (2006) *Resilience Engineering: Concepts and Precepts* (CRC Press, Boca Raton, FL).
- Iliev P, Welch I (2013) A model of operational slack: The short-run, medium-run, and long-run consequences of limited attention capacity. *J. Law Econom. Organ.* 29(1):2–34.
- Kiely J (2012) Periodization paradigms in the 21st century: Evidence-led or tradition-driven? *Internat. J. Sports Physiol. Performance* 7(3):242–250.
- Kiely J, Pickering C, Halperin I (2019) Comment on “Biological background of block periodized endurance training: A review.” *Sports Medicine* 49(9):1475–1477.
- Issurin VB (2010) New horizons for the methodology and physiology of training periodization. *Sports Medicine* 40(3):189–206.
- Issurin VB (2019) Biological background of block periodized endurance training: A review. *Sports Medicine* 49(1):31–39.
- Lakoff A (2007) Preparing for the next emergency. *Public Culture* 19(2):247–271.
- Lazear E (2006) Speeding, terrorism, and teaching to the test. *Quart. J. Econom.* 121(3):1029–1061.
- Solan E, Zhao C (2021) Dynamic monitoring under resource constraints. *Games Econom. Behav.* 129(September):476–491.
- Sutton J (2012) *Competing in Capabilities: The Globalization Process* (Oxford University Press, Oxford, UK).
- Teece DJ, Pisano G, Shuen A (1997) Dynamic capabilities and strategic management. *Strategic Management J.* 18(7):509–533.
- Varas F, Marinovic I, Skrzypacz A (2020) Random inspections and periodic reviews: Optimal dynamic monitoring. *Rev. Econom. Stud.* 87(6):2893–2937.