EVOLUTIONARY PROCESSES & ORGANIZATIONAL ADAPTATION

a mendelian perspective
on strategic management

DANIEL A. LEVINTHAL

4
Selection

Ideas, business plans, off-sites, and design efforts do not themselves directly receive rewards from the market. Organizations do. While there is a large nuanced literature on the theory of firm, the term "theory of the firm" has taken on very different connotations within the economics literature (Coase, 1937; Williamson, 1975; Gibbons and Roberts, 2013) and behavioral traditions (Cyert and March, 1963). In the economics literature, the motivating question under this rubric is the question of the appropriate scope of the firm's boundaries. For Cyert and March (1963), the motivating agenda was an empirically grounded account of firm behavior. The ensuing "Carnegie School" has emphasized the role of search, problem-solving, and feedback learning processes (Gavetti et al., 2007). Within the evolutionary economics tradition (Nelson and Winter, 1982), the firm is the central carrier of the enduring basis of capability differences and the object over which the selection force of competitive dynamics operates.

Building on this evolutionary perspective, a basic overarching fact about organizations is that firms receive profits and losses, while individuals generally only receive rewards as mediated by an organization's accounting system and incentive structure. In that sense, a firm can be considered to be a credit assignment mechanism (Holland, 1975). Understanding the nature of these processes is a fundamental challenge for management scholars. In this chapter, we consider three basic challenges in this regard: the problem of diversity of selection criteria, the challenge of the timing of selection relative to developmental processes, and the issue of units of aggregation and selection.

One might object to this line of argument in the context of for-profit enterprises as the "for-profit" objective would nominally seem to obviate the need to consider multiple selection criteria. However, even putting aside the issue of divergent stakeholder interests, even in the context of a for-profit enterprise a superordinate goal of maximizing the net present value of the enterprise still leaves open the issue of what might constitute the most meaningful and reliable metrics associated with progress towards this end-goal. The property that strategic initiatives have implications across time and/or "space" (other initiatives within the organization), corresponds to Andrews's (1971) classic contrast between what is considered "strategic" and what is "tactical." Reflecting these properties of temporal and spatial linkages, in this chapter we link the consideration of selection criteria to the issue of the timing of developmental processes and the units of aggregation which form the bases of evaluation. Further, we recognize that the environment, or contexts, in which the organization operates is itself an object of selection, which in turn influences the feedback processes the organization experiences.

4.1 Challenge of Diversity of Selection Criteria

It is important to contrast the diversity of underlying elements—people, ideas, routines—and the diversity of selection criteria. While we tend to privilege the former sort of diversity as critical to processes of innovation and change, there is a relative neglect of the role of the diversity of selection criteria—the diversity of perspectives as to what constitutes useful endeavors for an organization, the alternative means by which an agreed upon goal might be best achieved and the associated implications of these potentially diverse perspectives for an organization's resource allocation processes. Underlying this difficulty of organizations sustaining a diversity of selection criteria is the tendency for resources to be allocated by a hierarchical authority structure within an organization.

Our discussions of innovation and change tend to highlight the role of variety. However, variety alone is clearly not sufficient for

innovation. To take Kanter's (1988) imagery of "letting a thousand flowers bloom," such diversity in blooming will not be of consequence if the organization only has one type of "lawnmower," or less metaphorically, one type of screening criteria.[1] While obviously a caricature, the point is that experiments must be complemented by sufficient variety in the feedback mechanisms and selection criteria that inform the internal selection process within an organization (Adner and Levinthal, 2008).[2] Innovation within organizations requires resources; therefore, sustaining diversity requires ongoing resource commitments to a diverse set of emergent efforts.

We explore this question of heterogeneity of selection criteria in a number of respects. First, feedback, whether through learning at the level of an individual actor or via differential selection among a set of entities, is dependent on context. Therefore, it is important to consider the heterogeneous settings to which an organization is, or latently is, exposed. Second, a critical role of organizations is to mediate between aggregate outcomes, profit and loss, and payoffs to individual actors and acts. As a result, it is important to consider in detail the organization as an artificial selection environment (Levinthal and Warglien, 1999) or credit assignment mechanism (Holland, 1975). A critical issue in this mediation of external outcomes or payoffs and the distribution of rewards and feedback to the set of actors and subunits within the organization is the degree to which diversity of selection criteria are present. Diversity can be mitigated by a high degree of centralization of resource allocation as it can be difficult for a single actor to be of "multiple minds" regarding alternatives. However, diversity can also be mitigated by a high degree of socialization and convergent thinking among a set of nominally independent actors (Van Maanen, 1973; Levine and Moreland, 1991).

[1] Ironically, Kanter's (1988) phrase, in turn, was presumably at least indirectly inspired by Mao's (1957) famous encouragement to Chinese dissidents to come forth and "let a hundred flowers blossom." Consistent with the argument here regarding the importance of not only initial variety but subsequent selection processes, many of those dissidents who did come forth subsequently faced punishment by the state.

[2] A different mechanism is the role of slack (March and Simon, 1958) or a single screening criterion that is noisily applied (Knudsen and Levinthal, 2007) or imperfectly implemented (Puranam, 2018).

Learning processes are feedback driven. As a result, the particular context in which one operates critically influences the feedback received. Christensen's (1997) work on the disk-drive industry is usefully interpreted in this light. One can take a bundle of performance characteristics regarding cost, processing capabilities, weight, and power consumption and get very different responses in terms of perceived value depending upon which customer constituency one asks. The desktop-user community responded with a shrug of their collective shoulders when offered drives that were smaller and lighter, while the emerging community of laptop producers responded with enthusiasm for such possibilities.

The fact that firms and the products they produce compete in heterogeneous demand environments is an issue that has been of long-standing interest to marketing researchers. However, marketing research tends to suffer the opposite problem of a neglect of "supply side" considerations. This tradition offers methods and techniques to identify heterogeneity of demand, techniques of conjoint modeling and the like; however, this work tends to operate with an implicit assumption of enormous plasticity in the range of what the firm is capable of producing. For instance, the marketing challenge is to understand the appropriate degree of bitterness of a beer, and perhaps what the desired images are to be associated with a product, but there is no question of brewing, of how one might actually produce the beer with the desired attributes.

From a learning and adaptation perspective, heterogeneity in demand context not only says something about identifying desired positioning, but also about what sort of capabilities might emerge as a consequence of the path-dependent development of those capabilities in particular contexts. These ideas are expanded upon both in Chapter 5 where the diversity in selection criteria among actors is argued to be critical in understanding the contrast between exploration and exploitation and in Chapter 6 in the discussion of punctuated change, where it is argued that critical junctures in the evolution of a technology are its shifts to a new application domain (Basalla, 1988; Levinthal, 1998) and similarly how shifts in an organization's artificial selection environment are critical catalysts to organizational change.

4.2 Challenge of Timing of Selection

A different set of considerations revolves around the timing of the selection process relative to the unfolding of any given initiative. Selection is inherently myopic (Levinthal and Posen, 2007). Evaluation of fitness, whether a conscious choice by an organizational actor or external to the firm in the form of product market competition or financial evaluation, is based on the current instantiation of a project or entire organization. Thus, the interplay between the developmental trajectory of an object of selection and the timing and intensity of selection is central to how selection processes will ultimately play out.[3] A salient managerial expression of this tension is the pragmatic playing out of "real options," as well as the development and funding journey of new ventures. These challenges may involve "false negatives" as often lamented, situations in which the near-term feedback is not positive but in fact the unfolding developmental process would lead to an attractive outcome. However, organizations also face the challenge of "false positives" in which early positive signals are not terribly indicative of longer-run prospects. The potential for false negatives should temper enthusiasm for a focus on "early wins"; indeed, Levinthal and March (1993) argue that feedback-driven learning processes can be an important source of myopia. Navigating between these dual challenges of false negatives and false positive is an important organizational challenge (Sah and Stiglitz, 1988; Christensen and Knudsen, 2010; Csaszar, 2013). Indeed, Guler (2018) finds that the effective capability to evaluative incipient ventures is associated with substantially higher returns for venture capital firms.

The link between actions and outcomes occurs on varying time scales. Some actions, such as dynamic pricing for a web retailer, will have immediate observable consequences. However, even such immediate tactical acts may have longer-run consequences on subsequent period sales and potentially even a brand's reputational

[3] One might argue that selection forces can operate regarding beliefs or indications of this developmental trajectory; however, the critical property remains that those attributes that inform these interim evaluations need to be visible in some fashion.

capital. Immediate evaluation (selection) of the shift in pricing can only reflect the near-term outcome of a shift in near-term sales. For many initiatives of interest from a strategic perspective, near-term outcomes are a mere shadow of their ultimate payoff. The difficult challenge in devising selection processes is developing temporally proximate indicators that are suggestive of these ultimate payoffs.

In this regard, a striking observation when one looks at the literature on organizational learning is the extent to which this work has examined learning issues in, essentially, selection-free environments. Formal models of organizational learning tend to have the structure of seeding a population of organizations with diverse learning strategies or organizational structures and then observing the variance in performance within the population after some large number of learning trials. However, these nominally process-oriented modeling efforts tend to ignore the path to these performance asymptotes. Imagine that learning does not take place in the benign petri dish of a simulation model, but in a competitive environment in which survival until the end of the period of observation cannot be taken for granted. What then are the implications for the desirability of alternative learning strategies?

First, once learning dynamics are placed in a context of selection pressures, the meaning of what is a high-performing learning strategy becomes non-trivial. Is a good strategy one that generates high expected performance conditional on survival? This is implicitly the criteria of the business press, which tends to extol the virtues of dramatic gambles that paid off well. Alternatively, is a good strategy one that leads to a higher probability of survival? A third option, and a criterion that is used in most models of organizational learning, is the average performance of alternative strategies assuming that *all* organizations survive.

A fundamental problem for selection processes is that selection is occurring over a "moving target" (Levinthal and Posen, 2007). Indeed, as is clear from evolutionary arguments, selection can only be "intelligent" if there is a high degree of stability over what is being selected. However, in the innovation context, it is inevitable and quite appropriate that selection processes are enacted even when

development processes are far from complete.[4] Firms need to make interim judgments as to whether to continue to commit resources to a technology or product development effort and often cannot afford to wait for its full fruition or failure. Similarly, capital markets, particularly markets for venture capital, need to make interim evaluations as to whether a given concern is worthy of further resources.

Development paths are subject to more discerning intermediate selection, that is selection prior to the full realization of their potential that is indicative of this ultimate potential, to the extent that the correlation in the performance of development efforts across time is relatively high. To the extent to the extent that early success is suggestive of ultimate success, then intermediate selection can operate effectively. Development approaches, however, are likely to vary in their degree of correlation across time. Levinthal and Posen (2007) contrast development efforts in which initial efforts focus on one facet of the overall development effort, which they term the technical development sub-problem, with efforts in which the full business system of technology, manufacturing, and marketing is jointly searched. Exploring sub-problems has the virtue that it leads to rapid early performance gains and therefore is more likely to survive early screening efforts. However, such a search strategy that initially attempts to optimize a particular subsystem will tend to lead to lower correlation in performance across time than an integrated development effort. Thus, while the focused strategy leads to higher survival from early screening efforts, that early filtering process is less aligned with the ultimate selection criterion of the performance of an integrated system. As a result, integrated search strategies are shown to lead to higher average performance *conditional* on survival, even though the average performance under this search strategy in the absence of selection is inferior. Work illustrating the power of modularity tends to ignore these selection and survival considerations.

[4] Indeed, one might argue in the context of many organizational initiatives that there is a priori no fixed end-state, but rather these initiatives continue until some future termination or redirection decision.

A further implication of this argument is that introducing survival concerns turns on its head the now established view of managing the dynamics of exploration and exploitation. The standard result from search models is that in early stages one should engage in exploration so as to learn more about the set of possible actions and then, after some knowledge has developed, to engage in more exploitative behavior. However, again, these analyses do not concern themselves with survival. Young, small, vulnerable firms often have an acute survival problem. They need to exploit whatever modicum of wisdom they have about the world if they are to survive. Exploration is arguably for richer, more established firms; indeed, this idea is suggested by the notion of slack search (March and Simon, 1958; Cyert and March, 1963). With slack search, more innovative, exploratory efforts are prompted by performance results in excess of aspiration levels.

A context in which the issue of intermediate selection tends to be both quite important and under-attended to is the use of real options to justify and guide projects and resource allocation. There has been much enthusiasm among academics and practitioners for the tool of real options as a solution to the problem of how firms should manage their uncertain futures, particularly with regard to technological uncertainty (McGrath, 1997; Amran and Kulatilaka, 1999; Trigeorgis and Reuer, 2017). However, as Adner and Levinthal (2004) argue, real options are not quite the panacea that its proponents tend to suggest. The real options argument as applied to problems of strategic management has the following basic structure: the world is uncertain, therefore the firm should make lots of modest-sized "bets" and, as future states are revealed, the firm should exercise those options that now appear attractive, having positioned itself to do so as a result of its earlier investments. One of the basic concerns that Adner and Levinthal (2004) pose is how is the firm to know in this metaphoric "stage 2" which investments are attractive to strike or not. Unlike financial options, for which observation of current pricing in financial markets suffices, real options on technology provide no such clarity.

Indeed, the typical early-stage innovative effort results in a partial failure, or put more optimistically, a partial success. A typical scenario

is as follows. Deadlines for technical hurdles are not quite met, but some substantial progress is made. Potential users have not reacted with unabashed enthusiasm for the product, but it appears some modification of the feature set may result in a product with considerable appeal. If this is the modal outcome, what is the implication for managerial action and subsequent resource commitments? In the same spirit, critical to the logic of real options that enhances the value of initial "bets" on risky technologies is that exit and the termination of initiatives is a real possibility. However, analogous to Popper's arguments regarding hypothesis testing that we can only prove hypotheses false, but can never prove them to be true: an innovative effort cannot in general demonstrate the impossibility of future success (Adner and Levinthal, 2004). Rather, one observes a failure of the current embodiment of the technology to meet certain technical standards or satisfy the needs of a particular set of consumers. Such failure does not rule out the possibility that future incarnations of the technology might meet such standards, perhaps by pursuing somewhat different approaches; or, alternatively that the firm might be able to identify a different user community that would respond more positively to the current technology. This issue of a potential never-ending journey of search manifests itself as well in the context of the discussion of lean start-ups in which entrepreneurs are encouraged to "pivot" their way to success (Blank, 2003; Ries, 2011), but seemingly little attention is paid to the possibility that an innovative effort might best be aborted in its entirety.

To preserve the analytical logic of real options, a firm would have to put tight boundaries around the scope of an innovative effort, boundaries concerning technical approaches, markets to which the product is to be sold into, and perhaps temporal boundaries (Adner and Levinthal, 2004). However, imposing such boundaries has enormous potential costs as they deprive the firm from taking advantage of the unanticipated discoveries of possibilities that is common in innovative efforts. Thus, real options may certainly be applicable to situations of well-defined risk, where there is uncertainty over known possible states of the world, but are deeply problematic in the face of more ambiguous environments.

4.2.1 Online versus Offline Evaluation

A central building block of the behavioral theory of the firm is the notion of bounded rationality (Simon, 1955).[5] In contrast to the optimizing agent of neoclassical economics, Simon offered the satisficing decision-maker. Furthermore, the set of alternative actions are not presumed to be laid out in their entirety ex ante, but must be discovered or searched. This facet of the behavioral theory of the firm (March and Simon, 1958; Cyert and March, 1963) is by now well established. However, another critical facet of bounded rationality has been largely ignored in this tradition, and that is how alternatives, once identified, are to be evaluated.

Two points are focal in Simon's argument regarding bounded rationality. One is that only a subset of the entire space of alternatives is considered in a given choice setting. Furthermore, decision-makers may be confronted with a sequential unfolding of these possible alternatives, even among the limited set considered. Second, he postulated that these alternatives are evaluated by a simple discrete value function that distinguishes between satisfactory and unsatisfactory outcomes. In this sense, Simon substituted for the usual objective function of economic theory an additional constraint of what constitutes a feasible solution to the choice problem. What is less salient, though considered in the original discussion, is how actors are to evaluate the proposed solutions or alternatives. How do we know whether the various feasibility constraints are satisfied or not? Simon notes that there may be uncertainty as to whether a particular alternative may yield a state of nature that is in the satisfactory set or not, but suggests that this indeterminacy may be resolved by identifying a new alternative that does not suffer this risk.

Yet, this discussion points to an important lacuna in this early work and subsequent development of this line of inquiry. While ideas of search are central in behavioral theories of the firm, the mechanisms by which these alternatives are evaluated are less

[5] This subsection draws from Gavetti and Levinthal (2000) and Levinthal (2002).

clearly developed. Typical models of adaptive search have the following characteristics. Some space of possible alternatives is sampled. The realization from this "draw" is then compared either with the current status quo action or in other cases with an aspiration level. When the space of alternatives constitutes attributes such as prices, the model does not seem to require any elaboration. However, consider other possible spaces of alternatives, such as the space of possible new production technologies for a factory or the space of possible spouses.

When presented with a new alternative from one of these sorts of "spaces," how is one to recognize a satisfactory solution when one is confronted with one? Quick inspection of a possible spouse or a production plan may reveal certain proposed alternatives to be unsatisfactory, and some basic constraints or criteria may be revealed to be violated. However, the satisfaction of other constraints may not be so self-evident. How will the workforce respond to the production process? How reliable will the process prove? Similarly, will this proposed spouse prove to be an enjoyable companion upon repeated dining experiences, and will they prove reasonably tolerant of your array of annoying habits? The evaluation of proposed alternatives is a relatively undeveloped facet of the behavioral theory of the firm.

To provide some structure with which to consider such issues, it is useful to distinguish between two sorts of evaluation mechanism: the distinction between "online" and "offline" evaluation (Gavetti and Levinthal, 2000). Online evaluation refers to those settings in which evaluation can only take place by actual trial of the proposed alternative, whereas offline evaluation indicates the ability to assess value in the absence of such a trial. As with many dichotomies, this one is both informative and misleading.

The distinction is clearly important. Some possibilities are evaluated by thinking, by imaging possible futures should that alternative (spouse, production process, etc.) be adopted. Sometimes this thinking is supported with various tools of analytical reasoning, such as spreadsheets and yellow-pads. However, the dichotomy is also quite misleading. There is an enormous gray area between these two poles and most evaluation processes occur somewhere in

this intermediate zone. New production processes need not require shutting down the firm's entire operations and substituting the proposed process. One plant may serve as a test, while the prior technology is exploited in the remaining plants. In cases of more incremental changes, only a single line or shift of the production process may be sufficient to serve as an experiential basis on which to evaluate the proposal. The pilot plant operates at a smaller scale than the ultimate substantiation of the alternative would imply, but again it allows a detailed examination of feasibility at lower cost and lower risk than full adoption.

In other cases, an "artificial" environment is created in order to evaluate a proposed alternative that does not introduce the risk associated with a full commitment to a specific initiative. A particular type of artificial environment, wind tunnels to test the performance of new aircraft, offers some additional insights regarding the boundaries of on- and offline evaluation. Wind tunnels allow engineers to test loft and drag in a variety of conditions for a prototype of a possible airframe. However, wind-tunnels have substitute modes of evaluation. One, of course, is to engage in the enormous financial commitment of the full development of a working aircraft and the run the human risk posed to a pilot by the testing of such a craft. The other route is cognitive: to build computer models that simulate the performance of proposed designs. As knowledge of the underlying material and aeronautical engineering improves, offline evaluation can substitute for more online forms of evaluation. But note that this is really a matter of degree. The computer simulation in some form creates its own kind of experience base. It is simply a lower-cost artificial world than ones that involve building physical artifacts. Among the interesting properties of recent advances in 3-D printing is that this technology radically reduces the cost of bridging the digital world of possible forms and the physical world of particular forms.

A different sort of experience is the experience of others (March et al., 1991; Miner and Hanuschild, 1995). This sort of experience has the virtue that trial does not require the disruption of one's own activities; and, furthermore, that the set of alternatives that can be explored at a given time are potentially quite vast. Its weakness lies, of course, in the inferential difficulties that such a process poses.

How much does one learn by watching a potential romantic partner with another what they would be like as a potential partner? Or, perhaps less daunting, how much does some other plant's experience tell one about their own firm's likely success with a new production technology? We are probably more comfortable with generalizing from vicarious learning in the latter case, but that may in part stem from the fact that more of us have experience with the former context than in being plant managers and are more keenly aware of the idiosyncratic features of such relationships than of production processes.

In some sense, the issue of on- or offline search becomes less a categorical distinction than factors that influence the cost, risk, and possibly accuracy of the evaluation process. Online search often entails a particular sort of cost, that of the opportunity cost of not making use of established options. It is this opportunity cost that underlies the tension in the oft-cited exploration/exploitation trade-off. In turn, the degree to which current operations are disrupted by the need to evaluate a proposed alternative influences how painful that trade-off is.

Neighborhood search, in this regard, has a distinct virtue as it provides an effective, though not necessarily optimal, balance for the need to explore in an online manner alternative bases of action, while at the same time, neighborhood search exploits current wisdom about the world by means of the local nature of the search process. The need for this balance between exploration and exploitation depends centrally on whether the evaluation process of proposed alternatives is on- or offline. Thus, the wisdom of a particular sampling strategy is intimately connected to the possible form of evaluation of those samples. Many of our discussions of search processes have suffered by not sufficiently disentangling these two features of search processes (Knudsen and Levinthal, 2007).

4.3 Units of Aggregation

Behavioral arguments stemming from the Carnegie School treat organizations as being feedback driven (Levitt and March, 1988): actions that are associated with positive outcomes tend to be

reinforced and more likely to be invoked in the future in similar circumstances, while those associated with negative outcomes are less likely to be invoked in the future in similar circumstances. But this correspondence between action and outcome is a function of the level of aggregation at which both actions and outcomes are experienced, a consideration that has not generally been highlighted or extensively developed in the literature. Boundaries and the scope of activities have important consequences for both processes of feedback and adaptive dynamics, as well as for the playing out of selection processes.

As a consequence of the fact that selection processes are mediated by the organization, the scope of the enterprise impacts how selection processes play out. Perhaps the most fundamental fact about business organizations is that they comprise an aggregate unit by which a vast set of underlying activities is allocated payoffs by an economy. As a consequence, a critical organizational property is how the firm mediates between the aggregate outcomes experienced at the organizational level and the underlying initiatives and activities within the firm. The fundamental engine of adaptive learning is feedback and the linkage between a focal action and observed outcomes (Levitt and March, 1988). This linkage is more strongly felt and more apparent to the extent that actions and outcomes are local to one another and are more proximate in "space" or time (Levinthal and March, 1993). As a consequence, structures imposing or facilitating the patterning of relationships among actors affect the development of capabilities in a variety of ways.

The logic of the diversified firm is that there are non-trivial interdependencies among elements of the enterprise. These interdependencies, in turn, raise a challenge for resource allocation. A long-standing argument in the business strategy literature has pointed to the power of de-averaging, both with respect to understanding the firm's cost drivers and the value of customers, as a mechanism by which to create a more effective basis for resource allocation than that offered by more aggregate units of analysis that might mask key elements of heterogeneity (Ghemawat, 2002). Similarly, work in corporate finance has argued that diversified

firms tend to engage in inefficient resource allocation due to incentive conflicts among operating units (Stein, 2003). However, these efforts at "de-averaging" make the implicit assumption that the individual divisions and projects that serve as the elemental units in the resource allocation process are independent of one another—a property that runs counter to the fundamental logic of a diversified firm.

In this regard, it is useful to return to Simon's classic work on the architecture of complexity (Simon, 1962). Natural and artificially designed systems tend to be nearly decomposable, meaning that interdependencies are not scattered and widely diffused but tend to be focused and localized. Treating subsystems as independent may be a useful simplification or fiction, but it is important that strategists recognize that it is a fiction. Treating the system as a whole is not cognitively possible or practical—indeed Simon suggests that whether or not systems are in fact nearly decomposable, we need to perceive them as such in order to understand them. But, again, while not fully engaging with the full array of interdependences, it is important to maintain some level of mindfulness of their existence and the recognition that any strict decompositions are a possibly convenient fiction of a more complex reality of some degree of interconnectedness even among subsystems that are treated as being disjoint.

Organizational structures and budgeting systems serve as category structures and decompositions. The key levers of corporate strategy are essentially ways of bracketing the interdependence among subunits and averaging and de-averaging these sets of activities. For instance, consider GE under Reginald Jones who, with the help of McKinsey, developed the substructures of strategic business units (Joseph and Ocasio, 2012). Jones felt that the existing "coarse-grained" structuring of GE's activities did not allow him to really understand the underlying businesses and to allocate capital intelligently among them. In response to this, he developed the more fine-grained structure of strategic business units to address that challenge. Once we recognize that any given structural arrangement will inevitably imperfectly capture the true structure

of interdependencies, it suggests a healthy skepticism regarding any extant structure.

These arguments suggest that we should not be focused so much on the identification of a Platonian ideal structure, but rather to recognize the imperfections of any given structure and maintain some openness to the possible value of the re-bracketing of activities. The pattern of relatively frequent reorganization is often viewed as a pathology of organizational life and an indication of management failure or fickleness. While there certainly may be instances of failure to conceive of an appropriate design or executive "fickleness" and perhaps the need for new executives to provide their own imprimatur on the organizational structure, there may be a functional role for such changes as well.[6]

De-averaging helps to surface potentially important new investment opportunities that may lie buried within a larger budgetary unit. Indeed, the encouragement to "split off" units addressing new, emerging technologies, is the primary normative suggestion of Christensen (1997) in his work on how established firms may effectively confront "disruptive" technological change. De-averaging can also facilitate efforts to identify inefficiencies within the organization and support efforts at achieving a superior cost structure. This is a central premise of efforts at benchmarking in which a specific business process is decontextualized from the broader business system and the performance of this specific, isolated process is examined and contrasted with similar processes elsewhere in the focal firm or in other enterprises.

The primary benefit of averaging and more coarse-grained structures is to link within a common budgetary unit activities that have a high degree of interdependence. These activities in a more fine-grained structure, which specified administrative and/or

[6] The notion of "fickleness" suggested here differs from the idea of "efficiency fickle" put forth by Nickerson and Zenger (2002). Their argument suggests that the ideal organizational form lies intermediate between any pure form, such as functional or product-oriented structures. The suggestion here is that there is no ideal form, and one form merely makes different compromises and tradeoffs than another. In this spirit, Ethiraj and Levinthal (2009) show that shifting over time among a set of simplified goals, where the articulated goal at any point in time is only a subset of the full set of payoff determinants, facilitates adaptive learning relative to having a fixed goal structure, even one that contains the full set of performance dimensions.

accounting decompositions, would create externalities with the impact of choices made within one unit effecting other units of the firm. Per the issue of the importance of treating these issues from a dynamic perspective, it may be that at a certain point it is useful to operate at a relatively coarse-grain level of aggregation in order to allow managers to realize latent scope economies and to learn to manage these interdependences. Having identified these linkages and begun to develop routines and procedures to regularize their coordination, the "scales" may shift and the firm may be better served by then de-averaging these business initiatives to enhance incentives and clarify the best use of the firm's capital.

In sum, the corporation at its core is fundamentally a decomposition of economic activity within the broader economic system. Within its own boundaries, the firm must make further decompositions, or de-averaging, facing tradeoffs regarding the need for coordination, the provision of powerful incentives, and the appropriate allocation of capital. Work in business strategy and corporate finance has, rightly, highlighted the returns to de-averaging with respect to the capital allocation process. At the same time, we must not lose sight of the role of the firm in coordinating economic activity and managing interdependencies among facets of the firm. Absent strong interdependencies, the market and the price system serve as an effective means of incentive provision, coordination, and resource allocation (Baldwin, 2007). As nearly, not fully, decomposable systems, any bracketing of activity will be incomplete or, put differently, not true to the true underlying structure of interrelationships. Thus, efforts at specifying appropriate budgetary units must inevitably trade-off incentive intensity and an allocation of capital that more or less strongly corresponds to the underlying structure of business opportunities with realizing the latent benefits of the coordination and linkages among those opportunities.

4.3.1 Organizations and Multi-Level Selection

In addition to being linked spatially, another fundamental fact is that organizations have a hierarchical structure. This hierarchy

opens up the possibility of selection at multiple levels. In particular, while selection forces operate on the organization as a whole, the organization itself operates to select the various elements within its boundaries. The term artificial selection environment within the organization has been introduced here to contrast this intra-firm selection force with the "natural" selection environment of market forces. Some projects are curtailed while others are scaled up. Some products are taken to global scope and other products are supplemented. This process of internal selection may well itself be driven by stable operating routines. For instance, Burgelman (1994) observed how fabrication facilities were allocated to different product lines within Intel based on the profit margins the firm experienced on these products. In contrast, the R & D resources were allocated based on the firm's conception of its strategy, which at the time was that it was primarily a "memory company" and that memory components were key technology drivers for the firm. As a result, the firm shifted its production capacity to logic chips, where the profit margins were high, and at the same time continued to devote the bulk of its R & D resources to memory products. Thus, the firm internalized the external selection pressures with regard to production capacity, but with respect to its R & D activity the internal selection environment was loosely coupled to the external environment.

The management literature is replete with tales that illustrate this tension between how initiatives are valued by the organization versus their valuation in some facet of the external environment. The prototypical saga is one of a highly motivated manager (perhaps in a functional role of technologist or marketer) who identifies a promising new initiative for the organization that is subsequently evaluated by their superior. This evaluation may be based on the superior's sense of the external market and the possible payoff to the initiative; alternatively, the evaluation may be premised on the superior's interpretation of the fit between the proposal and their understanding of the firm's strategy.

Consider the feedback processes and selection criteria implied by this process. One basis of evaluation is an untested belief about the market. Such beliefs will persist by the failure to test them (Weick, 1979; Denrell and March, 2001), unless proven to the contrary

via vicarious learning from others. The other selection criterion corresponds to fitness with respect to the organization's ongoing policies and, more broadly, its conception of itself (Prahalad and Bettis, 1986). Such a selection criterion is not likely to introduce novelty or act as a source of change for an organization.

Thus, underlying this difficulty of organizations sustaining a diversity of selection criteria is the tendency for resources to be allocated by a singular authority structure within an organization. While a large organization may have sufficient resources to make multiple "bets," those individuals who control resource allocation decisions are unlikely to be of multiple minds. Further, while there may well be considerable diversity of opinion within the organization, there is typically a dominant political coalition (March, 1962) and the perspective of this ruling group will likely drive the resource allocation decisions.

Contrast this characterization to a population of organizations. Even if individual organizations make a singular "bet" with regard to a given opportunity, there may be tremendous diversity across the population of organizations. While there may be some pressure to conform to the perspective of other, respected organizations, individual organizations may receive highly differentiated feedback from their environment and this distinct feedback may lead them to different views of the same business opportunity. Indeed, the motivation of entrepreneurs to leave their prior organization often stems as much from their inability to convince their prior firm to pursue an opportunity that they feel has tremendous promise as it is associated with an incentive to appropriate for themselves the prospective returns associated with the pursuit of that opportunity (Klepper and Thompson, 2010).

Conceptually, a single firm could engage in selective intervention, using Williamson's (1985) terminology, and replicate the virtues of a population of independent organizations. However, as Freeland and Zuckerman (2018: 157) point out, Williamson recognizes that "because higher-level executives always retain fiat rights, they face the constant temptation to use them (and the information gleaned from monitoring mechanisms) strategically and opportunistically, especially vis-à-vis lower-level employees." One way to think about

this challenge of a commitment to delegating decision rights is to consider it a problem in "sequential rationality" (Selton, 1975).[7] The corporate office will, in all likelihood, have a point of view about the appropriate direction for the firm and the relative promise of individual initiatives. When faced with a given funding decision, it cannot commit to "throw away" its belief structure or to not act on it.

In contrast, committing to some form of structural decentralization or budget autonomy can act to constrain the central actors from imposing their perspective. In the limit, the formation of a new, distinct organization, freed from any authority structure from the corporate office, is a structural solution to this challenge. Alternatively, it may be possible to design commitment devices that restrict the impulses that make selective intervention problematic from the perspective of sequential rationality. Within an organization, one mechanism is to abdicate budgetary authority. This is often seen on a small scale when a corporation allows a modest percentage of a subunit's operating budget to be used at the subunits' discretion. 3M, and more recently Google, have received attention for instituting such a role at the level of individual managers, who are free to spend a portion of their time pursuing initiatives that they perceive to be valuable (Iyer and Davenport, 2008). The limitation of such an approach is that successful initiatives may not be financially self-sufficient and, as a result, ultimately require supplemental funding. One is then back in the position of having to convince some central authority of the merits of the particular initiative, albeit this evaluation would then occur with the initiative having proceeded with some level of development and therefore would constitute less of an abstract pitch, a glimmer on a white-board, and there would be a greater substantiation of the idea. Thus, the inherent hierarchy of organizations

[7] Selten (1975) developed the concept of sequential rationality as a refinement of Nash Equilibria. The critical distinction is that Selten lays out an extensive form treatment of strategic interaction that considers a player's decision at each point in the game tree. Strategies that may be consistent with a Nash Equilibrium that treats actors' overall strategies may not be sequentially rational—when confronted with a choice situation in the playing out of the game, the actor may not find it in their interest to carry out the action associated with the Nash Equilibrium.

(Michels, 1915) constrains the variety that a single organization, independent of its size, can sustain.[8]

4.4 Selection and Shaping of External Landscapes

The imagery of a Mendelian designer of contexts has been put forth up to this point with the focus on the analysis of the firm's internal landscape—the roles, structures, goals, incentives, screening criteria—that managers might set forth within the organization. However, organizations and their Mendelian executives can also influence the context in which the firm operates. There are two basic mechanisms by which such influence can occur. One, which could be viewed as a "selection" effect, is the choice of the contexts in which the organization operates—what are the markets, network of relationships, and so on in which the organization engages. The second mechanism can be thought of as a "treatment" effect. How might the organization influence its environment, whether through direct mechanisms such as lobbying regulatory authorities or somewhat indirectly by acting to catalyze some processes of collective action.

This "selection" into particular markets and market niches influences the feedback that the firm receives and the incentives that it perceives. What markets should the firm serve? What activities should be performed within the firm and what sorts of external linkages should the firm make? These choices provide managerial discretion over the evolutionary path that the firm's capability set takes. A firm's capabilities and market position emerge, are refined, or decay as a result of, or an absence of, product market activity. Therefore, the particular submarkets a firm serves will engender a distinctive, though not necessarily unique, set of capabilities. These capabilities do not follow directly from current operations. However, the incentives the firm has to make various investments

[8] Puranam (2018) offers an interesting counter-point to the loss of variety as a by-product of hierarchy arguing that to the extent hierarchy entails control and information loss, hierarchy will facilitate less centrally directed search and choice processes as a result of this loss of operational control.

and the political forces internal to the firm that may influence such decisions are not independent of its current product market activities.

The selection of the market context in which to operate is a higher-level form of adaptation in contrast to the possibly adaptive responses to the feedback from a given context. This contrast is somewhat akin to the distinction Argyris and Schon (1974) make between first-order and second-order learning processes. The impact on capabilities of serving particular markets is analogous to a first-order learning process. While not automatic, first-order learning processes are a fairly direct outcome of existing structures. By establishing a new set of linkages, whether by choice of a new submarket to serve, a new set of customer relations, or a new internal organizational structure, management sets in motion a new direction for the development of the firm's capabilities and its competitive position more generally. Prescient managers look ahead and anticipate such feedback effects when making decisions about what industries or emerging subfields to enter and which clients may help further the firm's development. Thus, in making a choice about what markets to serve, a firm is making a bet on a co-evolutionary process. The firm is, or should be, not only concerned with its current capability to compete within that domain, but also with how participating in that particular industry or subfield will affect the firm's future capabilities.

Perhaps the most basic attribute of the markets and customers served that will impact the development of the firm's capabilities is their growth rate. Is the firm serving customers and market segments that are growing rapidly, thereby providing a basis for not only significant growth in sales but likely opportunities for cost savings and greater efficiency as well as incentives for innovation (Klepper, 1996)? In addition, leading-edge customers may expose the firm to advances in technology and product offerings (Von Hippel, 1988). The role suggested here of leading-edge customers is analogous to Porter's (1990) discussion of demand factors associated with industry performance across nations. Porter (1990) points to two critical attributes of home country demand. One is timing: does the home country tend to be early or late in its demand for a

particular class of new products or services? The other is the level of sophistication and the degree to which customers are demanding in their quality requirements. These factors influence the speed and direction with which organizations proceed along their evolutionary trajectories.

4.4.1 Shaping

Gavetti et al. (2017) make an important conceptual distinction between efforts at "search," finding a more or less favorable location on some existing competitive landscape, and "shaping"—taking actions that change some of the properties of the landscape itself. Three different forms of "shaping" are considered here. First, firms are a node in broader technological and economic systems and the value that adheres to any set of firm attributes needs to be understood from this larger systems level perspective. A basic way in which such considerations manifest themselves is the presence of complementarities. For instance, a faster microprocessor may not have much value if the overall speed of computation is constrained by the capacity of the serial bus (Ethiraj, 2007). Work on ecosystems highlights that such interdependencies may be present in aspects of value creation beyond the physical product itself (Adner, 2012). Awareness of these interdependencies informs the preferred locus of innovative efforts (Ethiraj, 2007) and the architectural challenges of linking and inducing the necessary complements so as to support a firm's offerings (Jacobides et al., 2006; Adner, 2012, 2017).

Shaping can also take an important cognitive dimension. An important line of work points to the rule of category definition and legitimation (Wry et al., 2011; Pontikes, 2018). Sociologists have long noted the role of processes of social construction (Berger and Luckmann, 1967) in how societies create meaning and in particular how value is ascribed, whether aesthetics of beauty (Sontag, 2002), genres of music (Phillips, 2011), or food (Rao et al., 2005). These processes of meaning creation do not lie outside the influence of individual organizations. This has been noted in particular in settings of entrepreneurship and market creation, or the "birth" of a

new category (Santos and Eisenhardt, 2009; Wry et al., 2011). Even established market categories may be redefined as Carroll and Swaminathan (2000) demonstrate with reference to craft brewers who defined a distinct niche within the broader brewing industry—a niche defined by the nature of their production process, its technical features and its localization. What is critical is not that the niche was defined merely by these technological or market-positioning choices, but that there was an effort to form a collective identity around the new form.

A firm's institutional environment not only influences what might constitute an appropriate set of actions on the firm's part, but the firm's actions may also influence the nature of its institutional context as well (Ahuja et al., 2018). This has been discussed in terms of standard setting (Ranganathan and Rosenkopf, 2014), contesting and supporting the legitimacy of alternative organizational forms (Hsu and Hannan, 2005), organizational practices (Myer and Rowan, 1977; Fligstein, 1985), product market categories (Pontikes, 2018), the design of a firm's ecosystem (Adner, 2012), as well as more traditional considerations pertaining to the role of advertising and brand identity on demand.

One of the challenges in considering the endogeneity of an organization's external environment is the question of what is the degree of plasticity of an organization's context? While an organization may influence its demand environment and institutional setting, there is considerable rigidity and path-dependence in these structures. An organization is not operating on a "blank canvass," but rather a setting of other organizations and possibly entrenched institutional structures. Further, effective influence efforts are often the result of some degree of collective action, identifying shared interest among a set of actors. For instance, legitimating wind power as an alternative form of energy among latent consumers, working with regulators to support its inclusion on the exiting power grid, and addressing technical and environmental challenges is in the collective interest of the broad set of actors interested in developing and pursuing this opportunity (Sine and Lee, 2009). Thus, in assessing what might constitute the "adjacent possible"[9] in

[9] Kauffman (2000: 142) characterizes the "adjacent possible" as "all those molecular species that are not members of the actual, but are one reaction step away from the actual."

the space of alternative external contexts, one needs to be mindful of the existing configuration of actors and institutional structures.

4.4.2 What Makes a Good Niche?

One generally considers the question of the environment as a constraint to which the organization must respond. The environment is treated as an "iron cage" that pressures the organization to conform. While there is some truth in that perspective, the environment is also, to some degree, an environment of choice. To the extent that there is such choice, that raises the question of what constitutes a better or worse choice. Three properties of the environment or niche are suggested in this regard: fit, feedback, and malleability.

Fit is the classic consideration of the strategy literature (Andrews, 1971). Strategists are encouraged to map their organization's current strengths and weaknesses into possible market contexts. The literature on diversification adopts a similar sensibility with its emphasis on leveraging current capabilities and resources into additional settings. These are essentially myopic considerations and fall comfortably under the label of "exploitation." These are not trivial considerations and, as with exploitive acts more generally, are central to an organization's near and medium-term survival and performance. However, per the exploration/exploitation tradeoff, choosing contexts to maximize contemporaneous fit is unlikely to enhance long-run survival prospects.

Thus, in addition to these standard considerations of what might constitute a favorable environmental context, there are dynamic considerations of fit. First, what might be more or less promising environmental contexts in the future? When considering this question, it is important to recognize that that promise is not a bet on some exogenously determined "Wheel of Fortune." Rather, a context is promising as a joint property of some of its intrinsic properties and the potential for the shaping of the environment by the organization, possibly in collaboration with other entities. In addition, there is the question of what settings might be more or less generative of future capabilities. Thus, a choice of niches in which to operate is, effectively, a co-evolutionary bet on what

opportunities the niche may offer and what trajectory of capabilities and resource the niche may engender.

Thus, organizations themselves are not only complex adaptive system, but they reside in a broader ecology of other organizations and institutions. Distinct niche spaces in this broader ecology provide not only potentially quite divergent selection pressures, but these divergent selection pressures in turn serve as diverse bases of feedback that may influence the adaptive journey undertaken by a particular organization. Further, these niche spaces in which the organizations operate are not inert, nor are the changes they exhibit exogenous to organizations' behaviors. While these influence processes are not determinative and may only be poorly understand by the relevant actors, these external contexts are, to some degree, malleable. The space of the "adjacent possible" niches is to some degree an organization's choice and to some degree of its making.

4.5 Summary

It is argued that a fundamental role of organizations is to mediate between the selection forces in their external context and the specific initiatives and activities within the organization itself. The critical properties of this "mediation" consist of the selection criteria enacted by the organization, the timing and intensity of the selection process that is imposed, and the unit of aggregation at which the selection process operates. Selection is typically not based on some unidimensional criteria such as "fitness" or, in the language of business enterprises, profitability. Even with a shared consensus within the organization of a superordinate goal such as profitability, projecting this objective to particular initiatives situated in a particular temporal and spatial context (i.e. location within the enterprise's broader set of initiatives and substructures) is problematic. Selection inevitably must be made on the basis of various imperfect indicators of broader objectives. The diversity of these selection criteria a facet of the diversity that has been under-appreciated in

the literature, which primarily focuses attention on the degree of diversity of underlying initiatives and activities. The enlightened Mendelian executive needs to create structures and processes that recognize the inherent limitations of a singular viewpoint and create structures and processes that moderate their own potentially heavy hand in the internal selection process. Further, the Mendelian executive selects the contexts in which the organization operates, contexts which provide their own distinct feedback and selection pressures, as well as possibly shaping and influencing those contexts.

Selection in Artificial Intelligence: Ex-ante Wisdom and Model-Based Learning

This chapter introduces the premise that an organization can be thought of as a "credit assignment" mechanism that mediates between outcomes experienced between the organization and its environment and the rewards and resources allocated to individuals and initiatives within the organization. Given this perspective, it is worthwhile to consider the body of work in artificial intelligence related to this challenge and, in particular, how different lines of work within artificial intelligence have addressed the problem of how actions and decisions generated by computer algorithms are generated and evaluated.

Since its onset in the early 1960s, there have been two basic approaches in the development of artificial intelligence. In one line of development, there is an assumption that there is some pre-existing knowledge and the challenge is to develop a program that will reliability enact that knowledge. Early efforts within this approach, with Newel and Simon (1972) as the initial contributors, strove to codify the expertise of domain experts through the use of protocol analysis. The decision rules of the experts were codified in a series of "if-then" rules and the broader effort adopted the term "knowledge engineering" (Feigenbaum, 1978).

Continued

Continued

With the development of neural networks (Rumelhart and McClelland, 1986), a very different approach emerged, but an approach that nonetheless still operated on the property that there was clear knowledge of the domain ex ante. This approach became known as "supervised learning." Under this approach, a neural net is trained on a training data set, for instance mammogram images for which clinical assessment had been carried out as to whether the image was associated with the presence of a tumor or not. The algorithm having been trained in this matter is then tested on a hold-out sample to assess its performance.

In important respects, the knowledge engineering approach development by Simon and his students and that of neural nets are vastly different. Knowledge engineering involves making explicit the possibly tacit knowledge of the domain expert and representing that knowledge in an explicit computer code. A neural net does not require or engender any explicit causal model or understanding; rather, it generates a complex set of relationships among characteristics of the stimulus, the mammogram in this example, and the outcome (presence or absence of a tumor in this case). However, there is a known "truth" by which the computer algorithm is judged and hence the label of "supervised learning." In that sense, supervised learning and protocol analysis share the presence of ex-ante expertise. The neural net trained to assess mammograms does not try to peer inside the decision process of medical experts, but it depends on medical experts having created a clinical assessment of each of the images used in the training and assessment process. While the basic structure of neural nets has been long established, their ascendency and prominence awaited complementary changes in the cost and availability of computer power and, even more important, the availability of vast libraries of digital data sets on which these algorithms could be trained and tested.

An alternative approach emerged in parallel from the earliest days of the development of artificial intelligence (Samuel, 1959; Minsky, 1961) that extended classic approaches of operant conditioning (Skinner, 1957) and removed or effectively endogenized the "trainer." Even in task

environments that have a finite state representation, such as board games like chess, the number of possible states can be sufficiently large so as to make the exhaustive search for an optimum impractical. Samuel (1959) developed what Minsky (1961: 19) termed an "expectation reinforcement" mechanism. More concretely, moves within the play of the game were reinforced both by an immediate reward that is associated with the move and the valuation placed on the position to which the move lead. However, a critical feature of this approach is that these "valuations" are themselves a learned property, based both on the ultimate outcomes of the game (win, lose, or draw in Samuel's checker-playing program) and, importantly, by the play within a given game (what was the valuation of the new position to which a particular position led). As Denrell et al. (2004) suggest, such an approach allows for a kind of "bootstrapping" of a cognitive model or representation of the problem domain.

Processes of credit assignment, or what some have termed "actor-critique" (Holland et al., 1986) models, have a "Mendelian" quality to them. Assessment is made not merely on the basis of direct trials, but also on the basis of the actor's valuation function. Further, this valuation function itself evolves through experience within the particular context as well as across contexts, with the efficacy of that process being a function of the quality of the categorization schema by which contexts are encoded, a further element in the learning system. The temporal differencing approach developed by Sutton and Barto (1998) has served as the basic "engine" by which such valuation functions are modeled as evolving. Temporal differencing is essentially a behavioral variant of the recursive logic of dynamic programming. In place of the optimal value associated with any state for the subsequent play of the game that is specified in the context of dynamic programing, the actor's current valuation function is substituted and this valuation function itself is reinforced throughout the play of a given game and across the play of multiple, distinct games.

Yet another strand within artificial intelligence is "unsupervised learning." Here there is no "outcome"; rather, the underlying structure of the input data itself is assessed. Variants of this approach have entered the management literature (Hannigan et al., 2019) as a means of distilling

Continued

Continued

from rich text data underlying constructs and relationships. A blending of unsupervised learning and reinforcement learning has developed under the label of "deep reinforcement learning" (Krizhevsky et al., 2012). When the state space of a problem domain becomes quite large and complex, reinforcement learning becomes increasingly inefficient if the process of reinforcement is based on the visitation of specific states. Reinforcement learning in a vast state space benefits from a parallel effort of generalization. An important early exemplar of this dual approach is Tesuaro and Sejnowksi's (1989) program to play backgammon which used unsupervised learning to represent the state space of the game of backgammon and a reinforcement learning process to develop insight about more or less favorable moves within the game given that state space representation. When one moves to even more challenging contexts such as autonomous vehicles, this blending of the representation challenge with the process of reinforcement learning becomes even more critical. Can certain types of slow-moving masses be treated as a category called humans, and, given this category representation, what are reasonable rates of forward movement of the vehicle that both brings it closer to its target destination while keeping the likelihood of intersecting with this slow-moving mass to a minimum. Further, these algorithms are often evaluated in both an "offline" and "online" manner, with an initial training period based on a simulated traffic environment and then the more refined algorithm tested on more naturally occurring settings.

The arch of efforts at machine learning has interesting parallels with our characterization of modes of choice in Chapter 2 and our general identification of a Mendelian "middle ground" between "rational"/divine conception of choice and Darwinian processes of blind variation and selection. The early efforts at knowledge engineering can be interpreted as attempting to capture the wisdom of experts, though in contrast to the usual conception of rational choice the experts were often viewed as having implicit heuristics and much of the effort at knowledge engineering was to bring those implicit heuristics to the surface. While from the perspective of the knowledge engineer and the effort at creating a computer

algorithm, the expertise exists a priori. However, the knowledge of the domain expert was a process of both assimilating pre-existing codified knowledge and drawing inferences from considerable direct experience. Credit assignment and actor-critique models are particularly interesting and suggestive from a Mendelian perspective. The valuation function guides the near-term process of reinforcement learning; however, this valuation function itself evolves over longer time intervals. The reinforcement for the valuation function is both a kind of internal validity (did actions driven by the current model and beliefs lead to outcomes that are viewed favorably by the current model?), and the externality validity of feedback from the external environment. Actions are guided by the "artificial selection" of the credit assignment mechanism, but the basis of credit assignment is treated as an informed speculation that itself may change over longer time scales. In that sense, our Mendelian actor serves as a guide to the process of learning and selection, but a guide with modesty about the wisdom of any given manifestation of these guidelines.