Founding Conditions, Learning, and Organizational Life Chances: Age Dependence Revisited

Gaël Le Mens
Universitat Pompeu Fabra and *Barcelona Graduate School of Economics*
Michael T. Hannan
Stanford University
László Pólos
University of Durham

© 2011 by Johnson Graduate School, Cornell University.
0001-8392/11/5601-0095/$3.00.

•

We thank Glenn Carroll, Jerker Denrell, Balázs Kovács, Dan Levinthal, Ray Reagans, participants at the Organization Ecology Conferences, 2009 and 2010, at the Resource Partitioning Conference, 2010, at the Annual Meeting of the Academy of Management, 2010, at the Meeting of the Theoretical Organization Modeling Society, 2010, and at a seminar at the Massachusetts Institute of Technology (April 2010) for discussions and comments. Gaël Le Mens was supported by the Barcelona School of Management and a Juan de la Cierva Fellowship from the Spanish Ministry of Science and Education. Michael Hannan was supported by the Stanford Graduate School of Business, the Graduate School of Business Trust, and the Durham Business School. László Pólos was supported by the Durham Business School.

Empirical evidence about the relation between organizational age and failure is mixed, and theoretical explanations are conflicting. We show that a simple model of organizational evolution can explain the main patterns of age dependence and reconcile the apparently conflicting theoretical predictions. In our framework, the predicted pattern of age dependence depends crucially on the quality of organizational performance immediately after founding and its subsequent evolution, which in turn depends on the intensity of competition. In developing our theory, we clarify issues of levels of analysis as well as the relations between organizational fitness, endowment, organizational capital, and the hazard of failure. We show that once organizational learning is considered, founding conditions affect the fate of organizations in ways more complex than previously acknowledged. We illustrate how the predictions of our theory can be tested empirically and evaluate the effect of aging on the mortality hazards of American microbreweries and brewpubs by estimating the parameters of a random walk with time-varying drift. We also make some conjectures about expected patterns in other empirical settings.•

Scholars have debated whether young, middle-aged, or old organizations are more likely to fail. Despite considerable study, no clear answer has emerged from prior work (see Carroll and Hannan, 2000, for a review). Empirical evidence is mixed. Some studies have found that young organizations have higher failure hazards (Carroll and Delacroix, 1982; Carroll, 1983; Freeman, Carroll, and Hannan, 1983), others that somewhat older organizations have higher failure hazards (Carroll and Huo, 1988; Brüderl and Schüssler, 1990; Fichman and Levinthal, 1991), and still others that old organizations have higher failure hazards (Barron, West, and Hannan, 1994; Ranger-Moore, 1997). Moreover, existing theoretical explanations conflict, because they support only one of the three patterns of age dependence of the failure hazard, leaving the others unaccounted for (Hannan, 1998; Hannan, Pólos and Carroll, 2007).

There are several reasons for the conflicting theoretical explanations. First, there is an issue of levels of analysis: some explanations invoke organization-level processes whereas others invoke population-level processes. Second, most earlier work on age dependence suffered from a lack of careful attention to the distinction between stocks and flows of resources, which have distinct effects on organizational survival.

In this paper, we propose a simple model of organizational evolution that can both explain the distinct patterns of age-dependence in failure hazards and reconcile the existing theoretical explanations. Our approach builds on Levinthal's (1991) random walk, which makes the distinction clear. In the random walk model, the stock of resources (referred to as the stock of organizational capital), not the flows that arise from current performance, determines the risk of failure. Organizations with little capital can be destroyed by a random environmental shock that would do little damage to an organization with substantial capital. Yet performance and environmental shocks matter—they shape the dynamics of the stock of organizational capital. Specifically, the random walk model assumes that each producer is endowed with an initial stock of organizational capital. In each subsequent period, the stock of organizational capital increases or decreases according to the quality of organizational performance and environmental shocks. A producer fails when its stock of organizational capital reaches a critically low level, usually referred to in the random walks literature as an absorbing barrier. In this model, poor performance depletes the stock of organizational capital and moves a producer closer to the absorbing barrier, causing the hazard of mortality to rise. The situation is reversed when performance is good.

Although Levinthal's (1991) analysis focused on cases in which the systematic component of organizational performance remained constant, we explore what happens when performance improves with age, as when organizations learn from experience. This simple modification of the random walk has interesting implications. If initial performance is poor but improves with age, then the stock of capital initially decreases but later rises. This implies that the mortality hazard rises initially with aging and then falls—there is a liability of adolescence. But if initial performance is good and improves with age, then the stock of organizational capital rises with age and the mortality hazard falls with age—there is a liability of newness. This reasoning suggests that a simple model can produce several of the documented patterns of age dependence of the failure hazard.

In this paper, we develop this insight into a full-fledged model that also reconciles the leading theoretical explanations. We take particular care to incorporate environmental effects to limit the Panglossian implications of the initial (working) assumption that organizational performance continuously and unboundedly increases due to the positive effects of learning. We do so by integrating the ideas underlying the random walk with contemporary niche theory. This integration has two important features. First, by tying learning to constructs such as the appeal of a producer's offering to audience members and organizational fitness, we see that the effects of learning on the flow of resources are bounded by an upper limit. This upper limit is a function of the competitive conditions prevailing in the organizational population. Second, the most important implications of this theoretical integration concern the effect of the conditions at founding on organizational evolution, which are likely to be more subtle than previously acknowledged. For example, it has been proposed that intense competition at the time of founding (or entry) degrades organizational life chances over the life course (Carroll and Hannan, 1989). The new argument suggests a more complicated picture: the intensity of competition prevailing at founding affects not only the level of the failure hazard, but also its pattern of variation over time. Weak initial competition allows the quality of early organizational performance to be high, and a liability of newness will be the rule. But intense initial competition degrades early organizational performance, and a liability of adolescence or even a liability of aging will arise. In other words, the integrated model identifies the conditions that give rise to the three patterns of age dependence (liabilities of newness, of adolescence, and of old age) that have preoccupied research on the subject.

It is clear that subtle differences in arguments and some lack of clarity about units of analysis have hampered efforts at coming up with a unified treatment of this subject. This is exactly the kind of situation in which much can be gained from using a formal language to disambiguate the argument and to control attempts at integrating lines of argument (cf. Hannan, Pólos, and Carroll, 2007). Fortunately, much prior work on age and organizational mortality has at least a semi-formal character, often relying on parametric specifications (Freeman, Carroll, and Hannan, 1983; Brüderl and Schüssler, 1990; Levinthal, 1991; Barron, West, and Hannan, 1994) or a formal logic (Hannan, 1998; Pólos and Hannan, 2002). So we can build on these foundations for our integrated model of age dependence.

EXISTING PERSPECTIVES ON AGE DEPENDENCE

Several theories have been proposed to account for the distinct patterns of age dependence in the hazard of failure (for a comprehensive review, see Carroll and Hannan, 2000). Two main lines of argument make conflicting predictions. While several attempts at unifying conflicting theoretical predictions and empirical findings have been proposed, we contend that this previous work has overlooked an important distinction between the two basic lines of arguments.

Two Perspectives on Organizational Aging

According to the first perspective, the hazard of failure decreases with age. This kind of pattern is usually called a liability of newness, following Stinchcombe (1965). This line of argument hinges on ideas about capabilities and position. Stinchcombe proposed that younger organizations lack some crucial capabilities and positional advantages, which makes them vulnerable. More precisely, they (1) lack the technical and social skills needed for smooth functioning, (2) must invent roles, relationships between roles, and rewards and sanctions, (3) face uncertainty pertaining to social relations among strangers, and (4) normally lack strong social ties with external constituencies, which makes it harder for them to mobilize resources and ward off attacks. Here, for simplicity, we concentrate on the effects of capabilities. As organizations age, they learn from experience and improve their capabilities, especially when the environment remains stable (Nelson and Winter, 1982; see also Sørensen and Stuart, 2000). Improvement in capabilities enhances performance and lowers the hazard of failure. This line of argument implies that hazards decline monotonically with age, because capabilities improve over an organization’s lifetime.

The second perspective makes the opposite prediction: the hazard increases with age, at least initially. This line of argument assumes that newly founded organizations possess an initial endowment that lasts for some fixed time and buffers them from failure while it lasts (Stinchcombe, 1965; Brüderl and Schüssler, 1990; Carroll and Hannan, 2000; Fichman and

Levinthal, 1991). Scholars have defined endowment as the initial resources such as "seed capital, credit and commitment from others, and political support" (Hannan, Pólos, and Carroll, 2007: 154) that provide some immunity to a newly founded organization. And they have referred to the initial period during which organizations benefit from the immunity conferred by their initial endowment as the endowment period. Proponents of this view posit a monotonic positive relationship between the size of the endowment and the strength of immunity. But an important (often implicit) assumption of this perspective is that endowments get spent down, whatever their initial sizes. The immunity conferred by the endowment therefore declines over time, causing failure hazards to rise with age within periods of endowments.

Limitations of These Perspectives

The two arguments, as summarized above, focus on different constructs. The first focuses on capabilities, and the second focuses on initial endowments. Resolving the theoretical tension between the two arguments demands attention to the relationship between endowments and capabilities, to distinguishing stocks and flows.

An endowment is a stock of resources, whereas capabilities drive the flow of resources. This distinction between stocks and flows underlies a fundamental difference between the two arguments. Taken literally, the capability argument assumes implicitly that the past evolution of capabilities does not matter for the current risk of failure. All that matters is current capability, which determines current performance. But if the hazard depends on the stock of resources, as suggested by the second argument, then the past evolution of resource flows does matter indirectly. An organization with a history of substandard capabilities and weak performance generally has a meager stock of resources, even if its current capabilities have become superior. In contrast, an organization with a history of superior capabilities and good performance generally has an abundant stock of resources, even if its current capabilities have become inferior.

Similarly, the endowment argument assumes implicitly that learning and capabilities do not matter during the endowment period. This assumption can be justified by the conceptual linkage of endowment to the idea of initial immunity or a honeymoon period. We see this claim as questionable. The endowment at founding is simply the initial stock of resources. Just as an organization can improve its capabilities and strengthen its social position, it can also replenish its stock of resources. The theory we develop in this paper builds on this foundation to make novel predictions about the dynamics of the hazard of organizational failure. Fichman and Levinthal (1991: 447), in a discussion of social relationships, also suggested that the initial "stock of assets" can be replenished by "adaptation processes and the development of relationship-specific capital." But in their discussion of organizational failure and of the "burn rate" of initial capital, they emphasized selection processes instead of learning processes to explain the liability of adolescence. By contrast, we will show that learning processes are sufficient to explain the various empirically observed patterns of age dependence in organizational failure.

We adopt the perspective that the risk of failure depends more systematically on stocks of resources than on current flows. This view is consistent with the finding that failure in a given year rarely stems only from poor performance in that year or the year before but instead often results from sustained substandard performance (Hambrick and D'Aveni, 1988). The theoretical challenge therefore consists in integrating ideas about the evolution of capabilities, insofar as they affect the flow of resources, with the view that the stock of resources shapes the hazard of failure.

STOCKS, FLOWS, AND ORGANIZATIONAL EVOLUTION: A RANDOM WALK MODEL

We can illustrate the importance of the distinction between stocks and flows of resources for analyzing the dynamics of organizational failure by building on Levinthal's (1991) random walk. The crucial feature of this model is that it does not assume that initial endowments get depleted no matter what, almost independently of organizational performance. Rather, it defines a new state variable, *organizational capital*, as a conceptual extension of endowment to the whole lifetime of the organization. As such, organizational capital works just like the initial endowment, but it can be replenished if the organization performs well enough. Alternatively, it gets depleted if the organization performs poorly.

For the purpose of developing our theory, we consider organizational capital as just the size of a buffer that separates an organization from failure. But more substantially, it is a combination of financial resources, such as ample cash reserves, as well as non-financial assets, such as goodwill, positional advantages, a good reputation and ties with powerful social actors, that allow an organization to survive the deleterious effects of environmental shocks or temporary periods of poor performance. In that context, an organization fails when its buffer becomes fully depleted.

The focus on organizational capital provides a simple, yet powerful way to analyze the evolution of the hazard of failure. All that is needed is knowledge of the initial stock of organizational capital and of how the stock changes with age. Although organizational capital and size are closely related (Levinthal, 1991), as evidenced by the robust finding that large organizations are generally less prone to failures than small ones (see Carroll and Hannan, 2000, for a review), organizational capital is a multi-dimensional construct that includes dimensions as diverse as financial resources and goodwill from stakeholders. To develop our analysis of the dynamics of organizational failure, we will focus narrowly on the resources that an organization can extract from the environment by providing an appealing offering to relevant audience members, including offers of products to consumers and of jobs to potential employees. That having been said, a similar kind of analysis could be applied by focusing on other aspects of organizational capital, such as the relationships with stakeholders, which can also be characterized in terms of a stock of assets, including trust, goodwill, or favorable prior beliefs (see Fichman and Levinthal, 1991). With that restriction in mind, we can view changes in the stock of organizational capital as given by the net flow of resources. Positive flows cause organizational capital to accumulate, while negative flows cause organizational capital to be depleted.

In developing theory about this process, we address solely organizational failure, excluding other types of organizational exits, such as voluntary acquisitions, mergers, or strategic bankruptcies (Delaney, 1999). Also, while firms can vary in terms of the size of the stock of organizational capital that leads to certain failure, as when founders or top managers have distinct psychological thresholds about the minimum stock of organizational capital acceptable for keeping the organization alive (see Gimeno et al., 1997), we do not explicitly consider such firm-level heterogeneity in our theory. But we do control for unobserved firm-level heterogeneity in our empirical estimations.

To clarify this perspective and its implications, we now formulate this argument in formal terms. At founding, a producer is endowed with some stock of organizational capital. In each period, its organizational capital increases or decreases according to the quality of organizational performance. More precisely, the stock of organizational capital of producer *x* at the end of period *t*, $\kappa_x(t)$, is defined recursively as follows:

$$\kappa_x(t) = \kappa_x(t-1) + \mu_x(t) + \varepsilon_x(t),\ \varepsilon_x(t) \sim N(0,1). \qquad (1)$$

Capital at the end of a period equals capital at the start of the period plus the flow of resources, which in turn is the sum of the systematic component of organizational performance, $\mu_x(t)$, and of a random component, $\varepsilon_x(t)$, that is assumed to follow a normal distribution with a mean of zero and a standard deviation of one. A producer fails when its stock of organizational capital reaches zero.

In this model, we see organizational capability as influencing flows of resources, setting aside for the moment the role of the environment. In terms of the random walk, organizational capability therefore corresponds to $\mu_x(t)$, the systematic component of organizational performance (in technical terms, this is the drift parameter of the random walk). When $\mu_x(t) < 0$, the producer generally experiences a net outflow of resources, and the stock of organizational capital generally declines. This decline causes the failure hazard to increase, because the producer moves closer to the absorbing barrier, which makes it more likely that, due to the randomness of organizational performance, the stock of organizational capital becomes completely depleted in the next period. The situation is reversed when $\mu_x(t) > 0$.

Levinthal's (1991) analysis constrained $\mu_x(t)$ to remain constant over age. Here, however, we explore what happens when $\mu_x(t)$ increases with age. This is equivalent to assuming that capabilities improve with age, consistent with one of the main existing lines of arguments about age dependence. Suppose for the moment that

$$\mu_x(t) = \mu_x(\tau_x) + (t - \tau_x)\lambda_x,\ \lambda_x > 0 \text{ and } t \geq \tau_x, \qquad (2)$$

where τ_x denotes the producer x's date of founding. The parameter λ_x characterizes the (constant) learning rate, and $\mu_x(\tau_x)$ characterizes the quality of initial performance.[1] This simple modification of the random walk yields novel predictions: (1) If $\mu_x(\tau_x) < 0$, then the stock of capital initially decreases and then rises, so the hazard rises initially with age and then falls; and (2) If $\mu_x(\tau_x) > 0$, then the stock of organizational capital rises with age so the hazard falls with age. These relations have important implications for the dynamics of organizational failure. But before detailing these, it is necessary to discuss issues related to levels of analysis.

The Dynamics of Hazard of Failure at the Organizational Level

Levinthal's analysis centered on the effect of selection pressures at the *population* level. He showed how various combinations of the two crucial parameters, the initial stock of organizational capital and the organizational performance (i.e., the drift of the random walk), affect the *average* hazard of failure of the organizations in a population of surviving organizations, which we call the hazard rate. Even if organizational performance remains constant over age for each organization, the evolution of the hazard rate in the *population* could display a liability of newness or a liability of adolescence. In other words, even if aging per se does not affect individual organizations (i.e., there is no systematic relation between organizational performance and age), the population can still display patterns in which the average hazard varies with the age of the (surviving) members. This is because selection pressures tend to eliminate early on organizations with small stocks of capital. Depending on the details, selection can give rise to the various patterns. These patterns of age dependence in the population's hazard rate ought to be interpreted as depicting spurious age dependence in the sense that they do not depend on an aging process at the organizational level.

Here we investigate the dynamics of the hazard of failure at the *organizational* level. We consider a producer's failure hazard as a time-varying state variable specific to each organization. The hazard, denoted by $\omega_x(t)$, equals the instantaneous probability of failure of producer x at age t (conditional on survival until that age). Levinthal (1991) did not explicitly discuss the dynamics of $\omega_x(t)$.

The specification of the dynamics of organizational capital as a random walk implies that the hazard of organizational failure is the instantaneous probability that the organization receives a negative shock that depletes its capital completely, in technical terms, hits the absorbing barrier in the next (brief) time period:

$$\omega_x(t) = \frac{1}{\sqrt{2\pi}} \int_{u=-\infty}^{-\kappa_x(t)} e^{-\frac{1}{2}(u-\mu_x(t))^2}\, du. \qquad (3)$$

In terms of organizational capabilities, the hazard of failure tends to increase when capabilities are inferior and tends to decline when capabilities are superior.

1
We use this specification with a constant learning rate parameter for its simplicity, and we readily grant that a constant rate of learning is likely to be an oversimplification. In the next section, we develop a theoretical argument that implies that the learning rate tends to decline with age and in the Empirical Implications section, we estimate a model that allows for this.

This rendering of the random walk suggests that the two leading theories of age dependence can be seen as two complementary aspects of a single perspective. Endowment (organizational capital) buffers organizations from failure. But the degree of development of capabilities regulates the dynamics of this buffer and, in turn, the dynamics of the failure hazard.

More generally, the crucial construct that determines change in the failure hazard is the flow of resources. In what follows, we build a more refined conceptualization of this construct. To this point we have assumed implicitly that resource flows depend mainly on capabilities, which ignores the influence of environmental conditions. A useful model must attend at least to the competitive conditions prevailing in the organizational population. Improved capabilities do not necessarily imply that the resource flow ever becomes positive if competition is intense or if competitors also improve their capabilities at the same rate as the focal producer. To integrate such considerations, we build on recent theorizing in organizational ecology and derive more refined predictions.

ORGANIZATIONAL PERFORMANCE AND EVOLUTION: THEORETICAL INTEGRATION

In developing the new theory, we use the niche theory delineated in Hannan, Pólos, and Carroll (2007: chap. 7) as a starting point. We consider a generic producer that operates in an unspecified category and tries to capture resources controlled by members of the audience for that category. The relevant audience consists of actual and potential customers, actual and potential members of the organization, holders of capital, and, more generally, any individual, organization, or governmental agency that controls resources useful to the organization and takes an interest in the category. For simplicity, we assume throughout that each producer (1) bears only one category label and (2) operates at a single (unspecified) social position. Following standard sociological arguments, we assume that the audience members at a social position have similar tastes.

Building on the previous section, we assume that organizational capital is the main determinant of the hazard of failure. Looking at variations of organizational capital amounts to looking at the history of the net inflow of resources. We now express this dependence explicitly. Let $\delta_x(t)$ denote the net inflow of resources from the environment to producer *x* at time *t*.

Definition 1 (Organizational capital). A producer's organizational capital at time *t* is the sum of its initial endowment and the integral, over its lifetime until *t*, of its net inflow of resources:

$$\kappa_x(t) = \kappa_x(\tau_x) + \int_{\tau_x}^{t} \delta_x(s)ds,$$

where $\kappa_x(\tau_x)$ is the "initial endowment."

Whereas the initial endowment generally depletes over time, organizational capital can be replenished. Although the nature and composition of the current stock of organizational capital and of the initial endowment might differ, we do not see this difference as a concern. We employ these constructs primarily as they matter for immunity from failure. Whether the current stock of organizational capital consists of resources that have been acquired at the time of founding or subsequently does not matter for our theory. In this paper, we do not discuss how to measure organizational capital. However, we use it as a building block of a theory that makes new empirical predictions. Because these predictions can be tested without measuring organizational capital, this potential measurement issue does not make the theory unfalsifiable.

To implement our assumption that organizational capital buffers a producer from failure, we follow the spirit of Hannan, Pólos, and Carroll's (2007: 156) Postulate 7.4: "During a period of endowment, a larger endowment normally yields a higher expected level of immunity." In expressing this postulate and those that follow, we use a nonmonotonic "normally" quantifier $\mathfrak{N}$ (Pólos and Hannan, 2004). The use of this quantifier expresses that the formula should be considered as a generic rule, a rule with possible patterned exceptions. Such generic rules can be overridden by more specific information. The possibility of using specificity considerations to control potentially conflicting arguments makes a crucial difference in a companion paper that builds on the present paper but introduces drifting tastes (Le Mens, Hannan, and Pólos, 2011). Because we want to preserve the possibility of such integration, we use the nonmonotonic quantifier even though arguments do not get overridden here. In the context of this paper, it would not be misleading to treat $\mathfrak{N}$ as an ordinary (first-order) universal quantifier.

Postulate 1 (Organizational capital and the failure hazard). A producer's failure hazard, ω_x, normally decreases with its stock of organizational capital, κ_x:

$$\mathfrak{N}\, x \,\forall t_1, t_2\, [(\kappa_x(t_1) < \kappa_x(t_2)) \rightarrow \omega_x(t_1) > \omega_x(t_2)].$$

This postulate is consistent with the relation between organizational capital and the hazard implied by equation (3) in the modified random walk. We introduce this nonparametric representation because we do not want to tie the theory to a particular assumption about the distribution of the shocks.

Together, Def. 1 and Post. 1 imply that knowledge of the initial endowment and of the history of the past flow of resources suffices to make predictions about the stock of organizational capital and, in turn, about the evolution of the hazard of failure. We now turn to modeling the flow of resources.

Resource Flows and Fitness Thresholds

To answer the question of how organizations acquire resources, Hannan, Pólos, and Carroll (2007) introduced the concept of fitness as the ability to capture resources from relevant audience members. Although offerings often simultaneously target several social positions, we make the simplifying assumption that the offering of the focal producer targets the audience at one social position. This assumption keeps the formalism tractable. The basic idea is that a producer's offering in a category has some actual appeal to the audience at the focal social position that depends on the fit of an offering to the taste of that audience and on how the producer engages the audience, as explained below. Let $\alpha_x(t)$ denote the function with values in [0,1] that tells the appeal of the offering of producer *x* in the market for the category to the audience at the unspecified social position at time *t*. By definition, a producer's fitness is proportional to the appeal of its offering relative to those of its competitors.

Definition 2 (Relative fitness). An organization's fitness, relative to the other producers in the category, is its share of the total appeal among the offerings in the category at the unique position it targets. (Hannan, Pólos, and Carroll, 2007: D9.1 specialized to one social position):

$$\phi_x(t) = \frac{\alpha_x(t)}{\alpha_x(t) + \mathsf{A}_x(t)},$$

where $\mathsf{A}_x(t)$ denotes the total appeal of all of *x*'s competitors: $\alpha_x(t) + \mathsf{A}_x(t) = \sum_{x' \in \mathbf{o}(l)} \alpha_{x'}(t)$ and $\mathbf{o}(l)$ denotes the set of producers in the category.

An offering has high fitness when the audience finds it relatively attractive. Such attractiveness ought to garner a high inflow of resources. Although resources can flow from other sources, such as successful lawsuits or settlements or investments in the financial markets, these other sources are less systematic and are akin to rainfall gains in a model of agricultural production, in our view. For us, resources garnered from the target audience constitute the most important systematic part of the inflow, and we decided to develop our theory around those. Specifically, we model the resources that flow to the producer as the product of an increasing function of fitness and the amount of resources controlled by audience members at the social position targeted by the focal producer, which we denote by $R(t)$. The inflow of resources to the producer at *t* therefore equals $h(\phi_x(t))R(t)$, where $h(.)$ is a non-negative, increasing function.

While we have limited our discussion so far to inflows of resources, outflows also matter. It generally costs more to create offerings that audiences will find very appealing. For one thing, a producer must take costly actions to engage the target audience so as to convert an intrinsically appealing offer into an actually appealing one. More generally, producers have to commit to certain cost structures to repeatedly generate offers of a given appeal. These considerations suggest that we define a cost-structure parameter $\mathfrak{c}_x(t)$. Producers that seek to occupy the high-quality/high-cost position in the market lock themselves into a high threshold requirement; those that follow a lost-cost/low-quality strategy face lower threshold requirements. The cost structure parameter is useful in formulating a postulate about the relation between resource flow and fitness.

Postulate 2 (Fitness and the net inflow of resources). A producer's net inflow of resources normally increases with the product of its fitness and the quantity of resources controlled by audience members at the social position it targets; and it normally decreases with its cost-structure parameter:

$$\mathfrak{N}\, x\, \forall t\, [(\tau_x \leq t) \rightarrow (\delta_x(t) = h(\phi_x(t))R(t) - \mathfrak{c}_x(t))].$$

A feature of our approach needs emphasis. We are interested in the effect of age on the hazard, and not just in the covariation of the two variables. Age, regarded as a producer-level state variable, exactly covaries with the passage of time. From a theoretical standpoint, this implies that analyzing the process of aging requires that an investigator abstract away from the influence of the passage of time unrelated to aging directly. At first sight, it might seem impossible to separate the influence of these two variables. Fortunately, a solution exists thanks to a difference in terms of levels of analysis. The aging process affects a specific producer, whereas the passage of time has a systematic influence on all the producers in the population. In other words, we define the effect of aging as the producer-specific effect of the passage of time. Capturing the effect of aging requires holding constant or controlling the systematic (category-wise) effect of the passage of time.[2]

In our model, the systematic component for the population of producers comes from variations in $R(t)$, the resources that the audience at the focal position has devoted to the activities of the category at that time. Variations in $R(t)$ affect the resource flows of all producers independently of their ages. We do not find it useful to tie our analysis to any particular pattern of time variation in resource availability. We therefore assume that this quantity remains constant, and we omit it from the formal specifications of the results. Because this assumption is made for analytical convenience and does not express an intuition about the world, we refer to it as an auxiliary assumption.

Auxiliary Assumption 1. The total amount of resources available to the focal category remains constant over time.

$$\mathfrak{A} \forall t_1, t_2\, [R(t_1) = R = R(t_2)],$$

where the "assumedly" quantification operates implicitly over the focal market and social position.[3]

We also introduce another assumption for tractability reasons: the cost-structure parameter does not change over time. This is a simplification that allows us to focus on the core intuitions but also potentially reduces the scope of application of the theory developed here. The challenge is that highly fit

2
Here we confront the problem of separating the effects of age, cohort, and period (historical time). We control for historical time by conditioning of $R(t)$, and our analysis can be regarded as holding for an age cohort of organizations (see Ryder, 1965). With these two dimensions controlled, we can obtain a meaningful characterization of age effects. In the empirical analyses presented later in the paper, we also control for heterogeneity in individual frailty, following standard practice in demographic analysis (Vaupel, Manton, and Stallard, 1979).

3
The use of a different nonmonotonic quantifier for the auxiliary assumptions signals their different status in the theory. Otherwise $\mathfrak{A}$ functions just as $\mathfrak{N}$.

organizations will gain more of the fixed pool of resources and grow in scale. Such growth can yield scale economies of the classic form, if the market allows them. If growth in scale causes per-unit costs of production to fall, a producer can modify its offering by some combination of lower prices and higher quality. If a producer chooses the lower-price option, then its cost-structure cannot meaningfully be said to be constant. But if it chooses the higher-quality option, then it might keep a constant cost structure.

By invoking the assumption of a constant cost structure, we limit the scope of application of the theory (in its current form). It applies to situations in which producers cannot realize scale economies, perhaps because the relevant audience associates quality with small scale. And, even when such economies are possible, it applies to organizations whose strategies emphasize competition on quality.

Auxiliary Assumption 2. A producer's cost-structure parameter does not vary over time.

$$\mathfrak{A}\, x\, \exists \mathfrak{c}_x \forall t\, [(\tau_x \leq t) \rightarrow (\mathfrak{c}_x(t) = \mathfrak{c}_x)].$$

Given this pair of auxiliary assumptions, it follows that a meaningful threshold can be defined such that fitness above the threshold ensures a positive inflow of resources. The threshold depends on the two parameters we hold constant: the quantity of resources devoted to the category by the audience and the producer's cost structure. We now build an argument based on the existence of such a threshold.

Fitness, Organizational Capital, and Mortality

We can now use Post. 2, Aux. 1, and Aux. 2 to derive some key results about the existence of a fitness threshold $\mathfrak{f}_x$. This threshold is a quantity such that, if fitness surpasses the threshold, then the net flow of resources is positive. In this case, the producer's organizational capital accumulates and the hazard declines with age. If fitness lies below the threshold, then the pattern is reversed: the hazard rises with age.

Proposition 1 (Fitness, resource flows, organizational capital, and the evolution of the hazard of failure):
Case A. A producer's net inflow of resources is presumably positive if its fitness exceeds $\mathfrak{f}_x$; and it is negative if its fitness falls below $\mathfrak{f}_x$:

$$\begin{aligned} \mathfrak{P}\, x\, \exists \mathfrak{f}_x\, \forall t\, [(\phi_x(t) > \mathfrak{f}_x \rightarrow \delta_x(t) > 0) \\ \wedge (\phi_x(t) = \mathfrak{f}_x \rightarrow \delta_x(t) = 0) \\ \wedge (\phi_x(t) < \mathfrak{f}_x \rightarrow \delta_x(t) < 0)], \end{aligned}$$

where the 'presumably' nonmonotonic quantifier $\mathfrak{P}$ expresses the consequence of an argument (a rule chain) that builds partly on generic sentences.

Case B. A producer's stock of organizational capital presumably increases with age if its fitness exceeds $\mathfrak{f}_x$; and it decreases if its fitness falls below $\mathfrak{f}_x$:

$$\begin{aligned}\mathfrak{P}\, x\, \exists \mathfrak{f}_x\, \forall t_1, t_2, u[(t_1 \leq u < t_2) \rightarrow ((\phi_x(u) > \mathfrak{f}_x) \rightarrow \kappa_x(t_1) < \kappa_x(t_2))\\ \wedge((\phi_x(u) = \mathfrak{f}_x) \rightarrow \kappa_x(t_1) = \kappa_x(t_2))\\ \wedge((\phi_x(u) < \mathfrak{f}_x) \rightarrow \kappa_x(t_1) > \kappa_x(t_2))].\end{aligned}$$

Case C. A producer's failure hazard presumably decreases with age if its fitness exceeds $\mathfrak{f}_x$; and it increases if its fitness falls below $\mathfrak{f}_x$:

$$\begin{aligned}\mathfrak{P}\, x\, \exists \mathfrak{f}_x\, \forall t_1, t_2, u[(t_1 \leq u < t_2) \rightarrow ((\phi_x(u) > \mathfrak{f}_x) \rightarrow \omega_x(t_1) > \omega_x(t_2))\\ \wedge((\phi_x(u) = \mathfrak{f}_x) \rightarrow \omega_x(t_1) = \omega_x(t_2))\\ \wedge((\phi_x(u) < \mathfrak{f}_x) \rightarrow \omega_x(t_1) < \omega_x(t_2))].\end{aligned}$$

The proof of this proposition and of all other propositions and theorems can be found in Appendix A.

Incorporating Learning: Improvements in Capabilities

We now incorporate Stinchcombean thinking about increasing capabilities into the new model. As suggested in the discussion of the extended random walk above, the essence of this argument holds that performance tends to improve with age as the organization gains more experience and its members learn to operate together and within its institutional environment. We now integrate such a learning argument into our framework. Figure 1 provides a schema of the constructs used in our model that is useful in following the flow of the argument.

Our model allows two distinct ways to evaluate performance. We can consider either the actual appeal of a producer's offering or a producer's fitness. Actual appeal is an absolute measure of performance that depends on the relation between the producer and its audience. Fitness measures performance in relative terms. It is not enough for a producer to perform well in some absolute sense to have a high fitness—it must perform better than the competition. Of course, this also means that a producer need not perform well in an absolute sense to have high fitness if the competitors perform at a much lower level.

Because we read Stinchcombe's argument about increasing capabilities as referring to an absolute measure of performance, we represent it formally as an argument that implies that the actual appeal of the producer's offering increases with age. But we see the link between age and actual appeal as indirect.

According to the niche theory on which we build, actual appeal depends on intrinsic appeal and engagement. An offering has intrinsic appeal to an audience if its various

Figure 1. The constructs used in our model of age dependence.

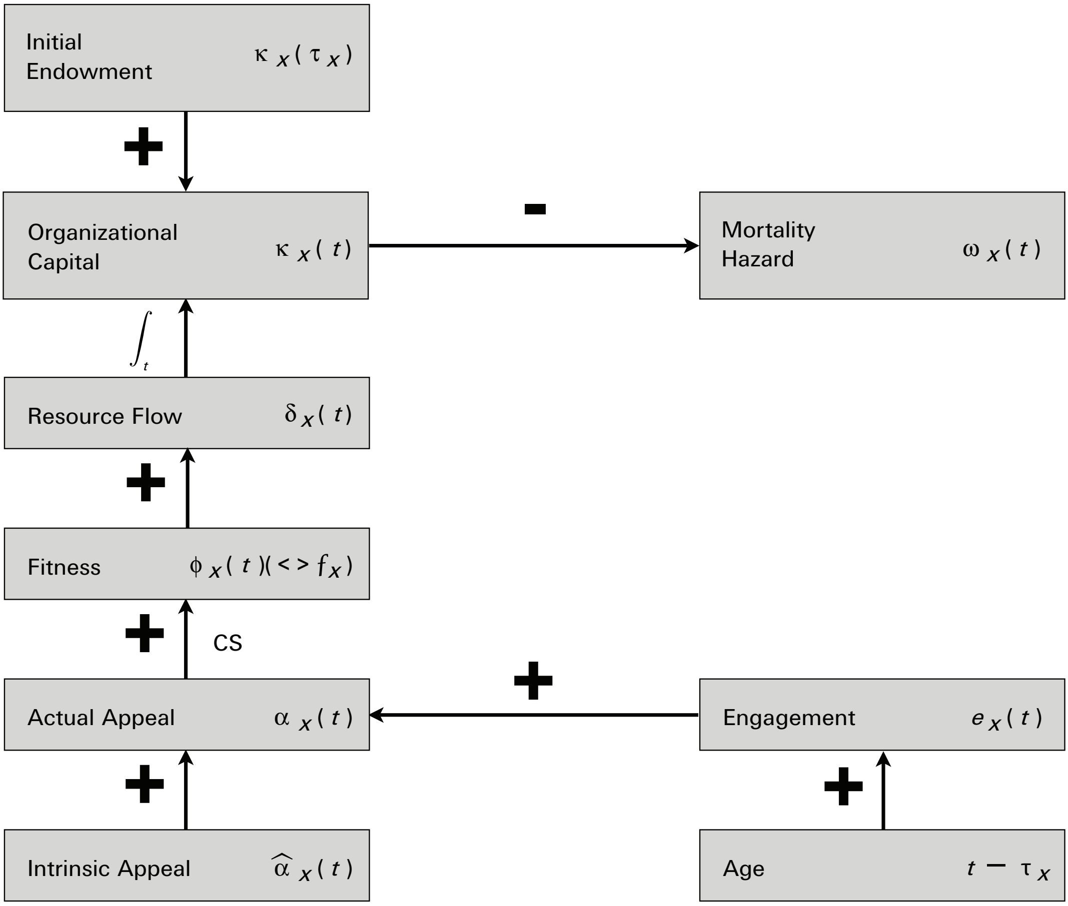

dimensions fit the tastes of the audience, but even those offerings that do fit generally do not have actual appeal unless they are made available (in an appropriate way), unless the producer engages the audience. Hannan, Pólos, and Carroll (2007: 179) sketched the notion of engagement as multifaceted:

> This term [engagement] refers to a diverse set of actions, including (1) learning about the idiosyncrasies of the local sub-audience and its aesthetics; (2) designing or redesigning features of the offering to make it attractive to that audience; and (3) trying to establish a favorable identity in the relevant sub-audience. In many cases of interest, key engagement activities include developing and displaying credible signals of *authenticity*. . . .

Some of the actions described are forward looking, involving efforts to learn about tastes so as to create offerings that will be intrinsically appealing at some later date. Other aspects of engagement involve presenting an existing offering in an appropriate way. And it is this second narrower meaning that animates Hannan, Pólos, and Carroll's theory of the niche, especially their treatment of the conversion of intrinsic appeal into actual appeal.

Continued development of the theory requires that the broad-brush notion of engagement be decomposed into its static and dynamic elements. This issue is especially germane when changing tastes are at issue (Le Mens, Hannan, and Pólos, 2011). We do not offer a systematic reanalysis of the concept here. Rather, we restrict our usage of the term engagement to the narrow sense of actions taken by a producer to bring its current offering to the attention of an audience in a manner that fits the aesthetics of the audience. In this narrow sense, engagement is tied to an offering. Let $e_x(t)$ be a non-negative real-valued function that records the level/quality of the engagement of producer x with respect to its offering at time t at the target social position. To integrate Stinchcombe's argument about increasing capabilities with these considerations, we introduce a postulate that claims that the quality/quantity of (offering-related) engagement normally increases with age.

Postulate 3 (Age and engagement). A producer's level/quality of engagement (with respect to its current offering) normally rises with age:

$$\mathfrak{N}\, x\, \forall t_1, t_2\, [(\tau_x \leq t_1 < t_2) \rightarrow e_x(t_1) < e_x(t_2)].$$

We see no need here to try to model variations over time in intrinsic appeal (but see Le Mens, Hannan, and Pólos, 2011). The capability story works with increasing engagement as long as intrinsic appeal is positive for the producer according to the construction introduced by Hannan, Pólos, and Carroll (2007: Post 8.3B), which they introduced as a way to fill in gaps in arguments when the available information is partial, e.g., when the analyst knows about engagement but knows less about intrinsic appeal (namely, only that it is positive).

We rely on a simpler relationship that fits the constrained setting we analyze. Let $\hat{\alpha}_x(t)$ denote the function with values in [0,1] that tells the intrinsic appeal of the offering of producer x in the market for the category to the audience at time t.

Auxiliary Assumption 3. The actual appeal of a producer's offering normally increases with its engagement (as long as its intrinsic appeal is nonzero):

$$\mathfrak{A}\, t_1, t_2, x\, [(\hat{\alpha}_x(t_1) > 0) \wedge (\hat{\alpha}_x(t_2) > 0) \wedge \\ (e_x(t_1) < e_x(t_2)) \rightarrow \alpha_x(t_1) < \alpha_x(t_2)].$$

Under these assumptions, it follows that actual appeal increases with age.

Proposition 2 (Age and actual appeal). The actual appeal of a producer's offering presumably rises with the producer's age:

$$\mathfrak{P} x\, \forall t_1, t_2\, [(\tau_x \leq t_1 < t_2) \rightarrow \alpha_x(t_1) < \alpha_x(t_2)].$$

Actual appeal presumably increases with age, and it is also bounded—it lies in the interval [0,1]. These two characteristics and the rules of calculus imply that the actual appeal converges toward a limiting value. We denote the limiting actual appeal of producer x by $\vec{\alpha}_x$.

This proposition has implications for the evolution of fitness, the key driver of accumulation or depletion of organizational capital. The relevant environmental condition is the strength of competitive interactions within the category. If competition becomes more intense, then improving capabilities do not necessarily translate into increased (relative) fitness. More generally, the argument should hold for an arbitrary producer in the category. But if capabilities improve in the same way for all producers, no one gains—fitness will remain constant over age for all members of a cohort. So we restrict the story to hold in environments in which the strength of competition remains stable over time. We do this by invoking the following predicate.

Definition 3 (Stable competitive pressure). A competitive environment for a producer x imposes stable pressure if the sum of the actual appeals of its competitors (in the focal producer's category) remains constant:

$$\mathrm{CS}(x) \leftrightarrow \exists \mathsf{A}_x\, \forall t\, [(\tau_x \le t) \rightarrow \mathsf{A}_x(t) = \mathsf{A}_x > 0].$$

With this condition, we can relate initial fitness to the failure hazard in a setting in which actual appeal improves with age.

Proposition 3 (Age and fitness). Suppose the environment is characterized by stable competitive pressure:

Part A. A producer's fitness presumably increases with age:

$$\mathfrak{P}\, x\, \forall t_1, t_2\, [\mathrm{CS}(x) \wedge (\tau_x \le t_1 < t_2) \rightarrow \phi_x(t_1) < \phi_x(t_2)];$$

Part B. A producer's fitness presumably becomes close to $\vec{\alpha}_x / (\vec{\alpha}_x + \mathsf{A}_x) \equiv \vec{\phi}_x$ (which we call long-run fitness):

$$\mathfrak{P}\, x\, [\mathrm{CS}(x) \rightarrow \lim_{(t-\tau_x)\rightarrow\infty} \phi_x(t) = \frac{\vec{\alpha}_x}{\vec{\alpha}_x + \mathsf{A}_x} \equiv \vec{\phi}_x].$$

This part of the argument imposes a bound on the effects of learning (improvements in capabilities) on fitness and mortality.

We now have the conceptual machinery needed to derive predictions about the evolution of the hazard with age. Our analysis shows that two conditions are crucial. The first is whether fitness at founding lies above or below the threshold $\mathfrak{f}_x$. The second pertains to the competitive environment. Three cases can occur, as illustrated by figure 2, which depicts the qualitative shapes of the time paths of the key variables of the theory. If initial fitness exceeds the threshold (Case A), then a pattern of capabilities increasing with age implies that fitness will remain above the threshold. This yields a pattern of increasing organizational capital and falling hazards of

Figure 2. Graphical representation of the predictions of Theorem 1.

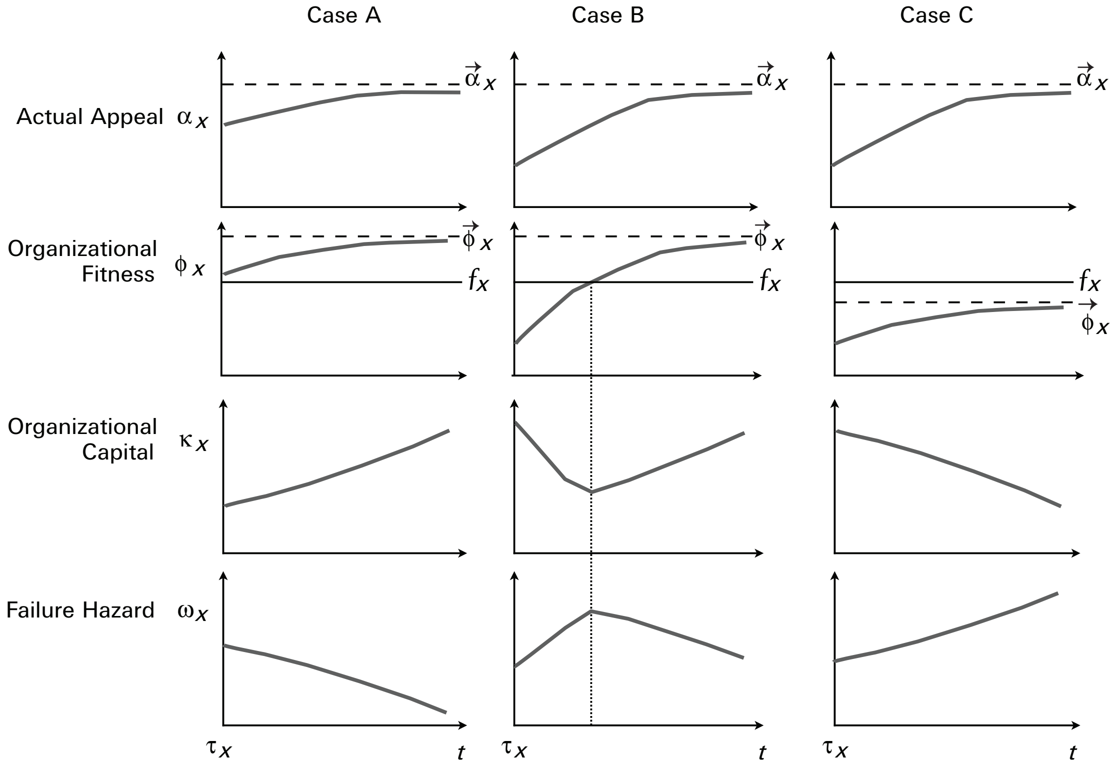

failure with aging. The second and third cases consider what happens when initial fitness falls below the threshold. A liability of adolescence emerges only if fitness at founding falls below the threshold, but the producer can reach a fitness high enough that it starts accumulating capital (Case B). In this case, the stock of capital gets depleted over an early period before it begins to accumulate. The hazard rises with age initially and then declines. If the current competition is so intense that a newly founded firm cannot reach the fitness threshold even if it can generate enough engagement, then organizational capital gets depleted monotonically over time, and the hazard increases with age (Case C).

Theorem 1 (Age dependence in organizational failure).
Case A. If a producer's initial fitness exceeds the threshold, then organizational capital presumably grows with age and the failure hazard presumably decreases with age—there is a liability of newness:

$$\mathfrak{P}\, x\, \forall t_1, t_2\, [\mathrm{CS}(x) \wedge (\phi_x(\tau_x) > \mathfrak{f}_x) \wedge (\tau_x \leq t_1 < t_2) \rightarrow \begin{cases} \kappa_x(t_1) < \kappa_x(t_2) \\ \omega_x(t_1) > \omega_x(t_2) \end{cases}].$$

Case B. If a producer's initial fitness falls below the threshold but its long-run fitness exceeds it, then organizational capital presumably first diminishes and then grows with age and the failure hazard presumably first increases and then decreases—there is a liability of adolescence:

$$\mathfrak{P}\, x\, \exists q\, \forall t_1, t_2, t_3, t_4\, [\mathrm{CS}(x) \wedge (\phi_x(\tau_x) < \mathfrak{f}_x) \wedge (\vec{\phi}_x > \mathfrak{f}_x) \\ \wedge (\tau_x \leq t_1 < t_2 \leq q \leq t_3 < t_4) \\ \rightarrow \begin{cases} (\kappa_x(t_1) > \kappa_x(t_2)) \wedge (\kappa_x(t_3) < \kappa_x(t_4)) \\ (\omega_x(t_1) < \omega_x(t_2)) \wedge (\omega_x(t_3) > \omega_x(t_4)) \end{cases}].$$

Case C. If both a producer's initial fitness and its long-run fitness fall below the threshold, then organizational capital presumably decreases with age and the failure hazard presumably increases with age—there is a liability of aging:

$$\mathfrak{P}\, x\, \forall t_1, t_2\, [\mathrm{CS}(x) \wedge (\phi_x(\tau_x) < \mathfrak{f}_x) \wedge (\vec{\phi}_x \leq \mathfrak{f}_x) \\ \wedge (\tau_x \leq t_1 < t_2) \rightarrow \begin{cases} \kappa_x(t_1) > \kappa_x(t_2) \\ \omega_x(t_1) < \omega_x(t_2) \end{cases}].$$

Theorem 1 keeps intact the spirit of the extended random walk. But the focus on fitness allows for a more refined perspective that considers the joint impact of early performance and the competitive conditions prevailing in the category. In proposing a modified random walk (equation 2), we had assumed that the learning rate was positive and constant. That is, we had assumed that the drift increases linearly with age, which leads to a prediction of liability adolescence when initial performance is poor. The development of this section leads us revisit this assumption. Even if a producer learns from experience, the marginal effect of learning on resource accumulation will tend to decline because of a ceiling effect due to (1) the fact that the actual appeal is bounded and (2) the competitive pressure. When early performance is poor relative to the threshold and competition is so intense that the focal producer cannot get to a point at which the net flow of resources is positive, i.e., cannot "break even," then the expected pattern is a liability of aging rather than of adolescence. In other words, initial conditions and the ecological conditions prevailing in the category constrain the applicability of the increasing capabilities argument. The stronger the competitive environment, the less likely it is that failure chances will fall with aging. To our knowledge, previous analyses of age and failure have not recognized that the predicted pattern of age dependence depends on these joint effects. We believe that this insight stems from our effort to integrate the ideas underlying the extended model into the comprehensive theoretical framework delineated in Hannan, Pólos, and Carroll (2007).

EMPIRICAL IMPLICATIONS

Our model makes novel predictions. In this section, we discuss issues related to empirical estimation and present an illustration to demonstrate how the key parameters of the model can be estimated from ecological data.

Levels of Analysis and Model Specification

The crucial issue in estimating the dynamics of failure as predicted by our theory involves choosing a level of analysis. Our theory concerns the dynamics of the organizational-level hazard of failure. But standard estimation techniques likely confound the consequences of organizational-level processes with the consequences of selection, as discussed above. This kind of confounding would not be a problem if organizational capital were observable. But even if some aspects of capital such as organizational size can be observed or reliably approximated, we believe it is generally not possible to get a measure of organizational capital good enough to control for selection effects. We can address this methodological problem by building on the random walk model discussed earlier in the paper. A largely overlooked implication of Levinthal's (1991) analysis is that estimating the parameters of a random walk can give some insight about the dynamics of the organizational-level hazard of failure. The two parameters of the standard random walk with drift are the initial stock of capital and the drift term. As we explained above, if the drift is positive, then the organizational capital generally increases with age, and the hazard of failure declines. And if the drift is negative, then the organizational capital generally decreases with age, and the hazard of failure increases.

The dynamics can differ at the organization and population levels. In fact, this is what happened in Levinthal's (1991) estimations for two organizational populations. He estimated a Makeham model, which specifies the hazard rate as $h(t)=\alpha+\beta\exp(\gamma(t-\tau_x))$, and obtained a positive estimate for β and a negative estimate for γ (for both populations), which suggests a liability of newness. But he also estimated the random walk with constant drift and obtained negative estimates of the drift (for both populations), which suggests that organizational capital declines with age: a liability of aging at the organizational level. These two sets of empirical estimations are not inconsistent, because they operate at different levels of analysis, and one takes selection into account while the other does not. In summary, empirical analysis of the dynamics of the organizational-level hazard of failure requires estimating models that properly control for unobserved heterogeneity due to selection. Random walk models seem appropriate for this task. We now turn to an illustration of how these considerations can be integrated into an estimation framework.

An Empirical Illustration

We build on the random walk, extending it to a specification that can incorporate the distinct patterns of Theorem 1. We continue to assume that the evolution of organizational capital follows the random walk given in equation (1), with an absorbing barrier at zero.[4] The drift of this random walk in our framework corresponds to the rate of accumulation of organizational capital. Whereas we assumed in the random walk considered earlier that the drift increases linearly with time (and is unbounded, see equation 2), the developments of the previous section suggest that the drift is generally bounded.

4
As before, we assume that the variance of the random shocks equals one. This is because, based on survival data only, one cannot estimate both the initial stock of organizational capital, $\kappa_x(\tau_x)$ and the variance of the shocks (see Levinthal, 1991).

We incorporate the idea of monotonically varying and bounded drift by formulating the representation of the drift as follows:

$$\mu_x(t) = \mu_x(\tau_x) + \beta \arctan(\gamma(t - \tau_x)), \quad t \geq \tau_x,$$

where $\mu_x(\tau_x)$ is the initial (time-of-founding) drift, γ characterizes the rate of change of the drift with age, and β characterizes the limiting drift.[5]

The random walk given by the pair of equations (1) and (4) implies a particular distribution of failure times. Knowledge of this distribution would allow maximum-likelihood estimation of the parameters $\mu_x(\tau_x)$, β, and γ, as well as the initial stock of organizational capital, $\kappa_x(\tau_x)$, from an empirical distribution of lengths of lifetimes (with information on censoring). To the best of our knowledge, however, an analytical formulation of this distribution for this specification does not currently exist. Nonetheless, it is possible to compute the likelihood function by using computer simulations. Therefore, we used maximum simulated likelihood (MSL) estimation to fit our model to the data (see Appendix B for further details).

We analyzed the dynamics of failure for microbreweries and brewpubs in the United States during 1961–1997. Carroll and Swaminathan (2000) provided a detailed description and discussion of these data.[6] Prior estimation of hazard rate models with a piecewise-constant specification showed no clear evidence of age dependence in failure for microbreweries and some evidence of negative age dependence for brewpubs (Carroll and Swaminathan, 2000: tables 3 and 4).

Our MSL estimates of the various random walk models are shown in table 1. The table reports estimates of the parameters of three specifications: (1) the random walk with constant drift, as in equation 1; estimates given in the second column of table 1; the random walk with linearly varying drift, as in equation 2; estimates given in the third column of table 1; (3) the random walk with monotonically varying but bounded drift as implied by our full theory (see equation 4); estimates given in the fourth column of table 1.

In interpreting the results of the empirical estimations, it is important to note that we did not impose any constraints on the signs of the coefficients (e.g., we did not constrain the drift to be positive, nor did we constrain it to be increasing). Estimations of the three models lead to distinct predictions. Estimations of the model with constant-drift specification (column 2) yield a negative drift term (λ). This means that the stock of organizational capital depletes with aging, a pattern that implies positive age dependence. In contrast, estimates of the model with linear drift (column 3) suggest that, although organizational capital initially depletes over time, it begins to be replenished after about eight years for microbreweries and six years for brewpubs, a pattern that implies a liability of adolescence.

The pattern of age dependence predicted by estimates of the model with bounded drift (column 4) also suggests a liability of adolescence, which is consistent with Theorem 1.B. More

5
For positive values, the arctangent function initially increases linearly with a slope of one and goes up to a ceiling of 0.5π. Equation 4 therefore implies that the initial rate of increase of the drift is given by $\beta * \gamma$ and that the limiting drift is equal to $\mu_x(\tau_x) + 0.5\pi\beta$.

6
We thank Glenn Carroll and Anand Swaminathan for providing these data.

Table 1

Maximum Simulated Likelihood (MSL) Estimates of Random Walk Models of Organizational Failure of American Microbreweries and Brewpubs*

	Model					
	(1)	(2)	(3)	(4)	(5)	(6)
Drift	None	Constant	Linear	Bounded	Linear	Bounded
Heterogeneity	No	No	No	No	Yes	Yes
			Microbreweries (N = 553)			
Initial capital $\kappa_x(\tau_x)$	1.77•	1.93•	2.44•	5.70•	2.73•	5.74•
	(0.03)	(0.03)	(0.07)	(0.02)	(0.03)	(0.05)
Initial drift $\mu_x(\tau_x)$		–0.03	–0.24	–2.96•	–0.37•	–3.00•
		(0.23)	(0.21)	(0.01)	(0.01)	(0.02)
Linear drift λ			0.03		0.03	
			(0.20)		(0.30)	
Age-rel. drift γ				2.06•		2.10•
				(0.09)		(0.21)
Limit. drift increase β				2.00•		1.99•
				(0.01)		(0.01)
Heterogeneity κ					0.26†	0.52
					(0.30)	(1.24)
Ln L	–464.3	–461.9	–447.3	–435.2	–440.7	–434.0
BIC	934.9	936.4	913.5	895.7	906.6	899.5
			Brewpubs (N = 929)			
Initial capital $\kappa_x(\tau_x)$	1.80•	1.94•	2.39•	5.79•	2.66•	5.78•
	(0.02)	(0.04)	(0.05)	(0.02)	(0.04)	(0.09)
Initial drift $\mu_x(\tau_x)$		–0.04	–0.24	–2.96•	–0.46•	–3.00•
		(0.34)	(0.16)	(0.01)	(0.02)	(0.08)
Linear drift λ			0.04		0.04	
			(0.15)		(0.09)	
Age-rel. drift γ				1.70•		1.68•
				(0.09)		(0.31)
Limit. drift increase β				2.05•		2.03•
				(0.01)		(0.08)
Heterogeneity κ					0.21†	0.47
					(0.15)	(1.07)
Ln L	–726.4	–719.6	–701.6	–675.9	–685.9	–672.3
BIC	1459.6	1452.8	1423.7	1379.1	1399.2	1378.7

• $p < .01$, one-sided t-tests.
* Standard errors are in parentheses.
† $p < .01$, one-sided t-test, with null hypothesis $k = 1$.

precisely, the drift of the random walk is initially negative but ultimately becomes positive. The limiting drift of the random walk is 0.18 (σ = 0.02) for microbreweries and 0.26 (σ = 0.03) for brewpubs.[7] Parameter estimates of the model with bounded drift suggest a shorter length of time until organizational capital begins to be replenished, compared with what is implied by the model with linearly varying drift. More precisely, parameter estimates imply that organizational capital starts to accumulate (and organizational mortality starts to decline) after about 5.3 years for microbreweries and about 4.6 years for brewpubs.[8]

Likelihood-ratio tests and analyses of Bayesian Information Criterion (BIC) scores support these conclusions. Likelihood-ratio tests show that the specifications with age-varying drift

7
Both limiting values are significantly positive. The standard errors on the limiting drift were computed using the variance-covariance matrix of the parameter estimates.

8
When $\mu_x(\tau_x) < 0$, the time at which organizational capital starts to accumulate is given by $(1/\gamma)\tan(-\mu_x(\tau_x)/\beta)$.

improve significantly over the model with constant drift for both populations.[9] A comparison of the BIC scores of the models with linear drift (column 3) and bounded drift (column 4) suggests that this latter model best explains the ecological data for both microbreweries and brewpubs. Besides, parameter estimates imply that the rate of resource accumulation is increasing over time. Overall, these results are consistent with the theory developed in the previous sections.

The foregoing analysis assumed that all producers are alike in the sense that they follow a similar developmental trajectory and that differences in failure risk are explained just by organizational age. But another source of heterogeneity is possible: some organizations might be systematically more immune to failure than others. To check the robustness of our results, we also estimated models that control for this source of potential unobserved heterogeneity. Using a standard frailty approach (Vaupel, Manton, and Stallard, 1979; Tuma and Hannan, 1984; Levinthal, 1991), we assumed that the hazard is multiplied by a random quantity η that corresponds to a time-invariant organization-specific factor. In our framework, this approach is akin to allowing for some unobserved heterogeneity in terms of the initial stock of organizational capital or in terms of the location of the absorbing barrier. We assumed that η follows a gamma distribution with parameters k and $1/k$.[10] This specification allows for a simple formulation of the probability density of the failure times using the density with no control for heterogeneity and the gamma distribution (see Appendix B).

Table 1 shows that controlling for unobserved heterogeneity yields some improvement in model fits (see columns 5 and 6). But, importantly, estimates of the theoretically meaningful coefficients remain remarkably similar to those obtained without controlling for unobserved heterogeneity. And while controlling for unobserved heterogeneity significantly improves the fit of our main model (with bounded drift) for brewpubs (likelihood-ratio test, $p = 0.03$), the improvement is not significant for microbreweries (likelihood-ratio test, $p = 0.3$). This suggests that a substantial portion of the cross-organization differences in terms of failure risk is captured by the random walk model with bounded drift. Allowing systematic differences between organizations to play a role suggests that these have only limited to moderate explanatory power.

This empirical illustration shows that considerations of age-varying resource accumulation can be integrated in a random walk and that the specification can be estimated with ecological data. Even with no data about resource flows and stocks, estimating extensions of the random walk allows for robust inferences about the dynamics of organizational capital and, ultimately, the dynamics of organizational failure.

DISCUSSION AND CONCLUSION

What happens to the life chances of organizations as they age? We have proposed that the dynamics of the risk of failure depends crucially on early performance and on the

9
The BIC score allows for a comparison of non-nested models. According to this criterion, the model with the lowest BIC score is to be preferred.

10
These parameters imply that the mean of the heterogeneity distribution is 1 and the variance is $1/k$. With this choice, the specification of the random walks tells what holds for cases close to the mean of the distribution of unobserved heterogeneity.

historical pattern of resource accumulation. In addition to leading to the formulation of new predictions, our perspective has allowed us to integrate existing theorizing about age dependence in a single, unified, theoretical framework: the depletion of initial endowment argument and the increasing capability argument are seen as two complementary aspects of a more general model of resource accumulation and its effects on organizational failure. In this section, we discuss the relation of our theory to previous unification attempts, the assumptions of our model and what happens if we relax them, and the phenomenon of organizational obsolescence. We conclude by making some conjectures and suggestions for future research.

Relation to Previous Unification Attempts

Several prior theoretical efforts attempted to combine the arguments about improving capabilities and the depletion of endowment (Hannan, 1998; Pólos and Hannan, 2002, 2004). These have assumed that whether the first argument or the second argument should hold depends on the age under consideration and (for the second and third unification) the existence of common durations of the periods for all organizations in the population. They suggested that, within endowment periods, the argument about depleting endowment holds because it is more specific, and specificity matters in nonmonotonic inference. But these reformulations proposed that the dynamics of aging change considerably when the initial endowment has been depleted. The general tendency for organizations to build capabilities as time passes no longer gets countered by the effect of a diminishing endowment, and, therefore, the hazard of failure should then decrease with age because of a continued improvement in capabilities.

Hannan, Pólos, and Carroll (2007: chap. 7) identified the distinction between two periods, before and after the depletion of the endowment, with different dynamics as problematic for true theoretical unification. Supposing that two contradictory theories (one positing positive age dependence and another one positing negative age dependence) apply to different periods does not really amount to theoretical unification. To address this challenge, these authors made an intense use of nonmonotonic logic to integrate the different theory fragments in a unified framework that could describe the relation between age and the hazard of failure for all stages of the life of an organization. Our theory (in this phase) does not make use of specificity considerations and nonmonotonic logic to achieve theoretical integration.[11] Instead, it introduces two crucial new distinctions: whether initial fitness lies above or below a threshold tied to cost structure and whether the intensity competition prevailing in the category allows for fitness to rise above that threshold. Our model allows for within-population heterogeneity, for the effect of age on the failure hazard to differ among the organizations in a population. A detailed discussion of the relations between our theory and Hannan, Pólos, and Carroll's (2007) recent unification attempt can be found in Le Mens, Hannan, and Pólos (2010).

11
Despite this, we formulated our postulates, propositions, and our theorem using the nonmonotonic logic formalism because we wanted to allow for the possibility of some of our claims being overridden in more specific cases. We used this possibility in a companion paper (Le Mens, Hannan, and Pólos, 2011) to make predictions about the evolution of the failure hazard when organizations become old (obsolescence) and the environment drifts.

The Assumptions of the Model

The main result of the paper, Theorem 1, relies on the assumption of competitive stability, on the assumption of a fixed cost structure, and on the lack of a feedback loop between organizational capital and actual appeal.

Competitive stability. Although our invoking an assumption of competitive stability might appear surprising, the issue primarily concerns levels of analysis. We focus on organization-level processes, but variations in competitive pressure occur primarily at the population level. Furthermore, allowing for freely time-varying competitive pressure would lead to overdetermination because it would add one too many degrees of freedom in the modeling framework. That said, we do not claim that competitive stability holds in most populations, and it is certainly interesting to explore what the predicted evolution of the hazard of failure would be in conditions of time-varying competitive pressure.

Suppose, for example, that competitive intensity increases over time. In our framework, increasing competition translates into a decreasing upper bound on the fitness of the local producer. In line with the notation used above, let us denote this by $\vec{\phi}_x(t)$. We have

$$\vec{\phi}_x(t) = \frac{\vec{\alpha}_x}{\vec{\alpha}_x + \mathrm{A}_x(t)},$$

where $\vec{\alpha}_x$ is the limiting actual appeal of the offering of the focal producer and $\mathrm{A}_x(t)$ the time-varying total appeal of other producers in the category. If the competitive pressure rises, $\vec{\phi}_x(t)$ likely falls below the threshold $\mathfrak{f}_x$ provided that it does not already lie below it. This implies a modification of Theorem 1.A. Even if a producer's initial fitness exceeds the threshold and it learns from experience, so that its actual appeal increases with age, its hazard will start to increase after some age if competition intensifies over time. In other words, producers will experience a liability of newness followed by a liability of aging. And the liability of aging would arise because the focal producer cannot keep up with its competitors. Its actual appeal is bounded by $\vec{\alpha}_x$, but the actual appeal of competitors is not bounded in this scenario. This can correspond to a situation in which inertia limits the possibilities for improvement by the focal producer, but the competitors taken together do not face such limits. One reason is that even if some competitors are also becoming increasingly inert, and potentially fail, they might be replaced by new organizations, better adapted to the current competitive environment.

Similarly, increasing competitive intensity would also affect the argument behind Theorem 1.B. If competition increases so fast that the focal producer's fitness never exceeds the threshold, the result is a liability of aging (cf. Theorem 1.C.). In terms of the extended random walk, this translates into a drift parameter that is potentially negative. If competition increases more slowly, then a focal producer's fitness might first surpass the threshold and later drop below it. In that setting, the predicted pattern would be a liability of adolescence followed by a liability of aging.

It is possible to proceed with similar thought experiments, assuming decreasing competition or nonmonotonic patterns of time-varying competition. But, again, we see such exercises as somewhat inconsistent with our approach because they focus on a different level of analysis.

Fixed cost-structure parameter. We introduced the auxiliary assumption of a fixed cost-structure parameter (Aux. 2) to guarantee that the fitness threshold is fixed over time, which helps simplify the exposition of the core intuition underlying our theory. Though this limits the scope of applicability of the theory in its current form, as discussed above, it is possible to see what happens if we allow the cost parameter to decline over time, as when there are scale economies or increases in operational effectiveness. In that case, the flow of resources $\delta_x(t)$ is positive when

$$\phi_x(t) > h^{-1}(\mathfrak{c}_x(t)/R).$$

Suppose that initial fitness is so low that, right after founding, the flow of resources is below this threshold. Two effects help the flow of resources improve: increasing fitness and a decreasing cost parameter. In such situations, it is more likely that the flow of resources will become positive at some point than when there are just improvements in the fitness level. As such, when the cost structure parameter declines over time, the liability of adolescence becomes a more likely pattern of age dependence in failure hazards at the expense of the liability of aging. It is much more difficult to make predictions when the cost-structure parameter follows other patterns (such as a non-monotonic evolution over time).

Organizational capital does not affect appeal. Our model does not incorporate a "feedback loop" linking organizational capital and actual appeal. We readily acknowledge that this assumption simplifies reality. Organizations with extensive organizational capital would surely have more resources to devote to the improvement of their offerings. But assuming that actual appeal depends on organizational capital leads to almost intractable theoretical complications.

Suppose that actual appeal depends only on organizational capital and not on age. Then the rate of accumulation increases with the stock of organizational capital. When would the growth of organizational capital stop? A model with such positive feedback generally implies an unlimited growth for producers that have accumulated a certain stock of organizational capital. Such producers would therefore become completely immune to failure. This seems to contradict the evidence that hardly any firm can sustain a very long-term competitive advantage (Rosenzweig, 2007). Besides, even without positive feedback from organizational capital to resource flows, models of resource accumulation like the one we propose lead to strong persistence in resource heterogeneity, consistent with empirical observations (Denrell, 2004).

More important, how would producers with little organizational capital ever gain actual appeal and start to accumulate resources? A possibility would be that new producers would start off by targeting social positions that are not targeted by strong incumbents, in a process akin to resource partitioning (Carroll, 1985). But we have assumed that producers in our model only target one social position. This suggests that making actual appeal depend on organizational capital would require a much more complicated model.

If actual appeal depends on organizational capital, it should also depend on the time dimension (to capture experiential learning). Then how would these two factors interact in determining actual appeal? Current research does not offer an answer to this puzzle. For example, Hannan and colleagues (1998) studied the interaction between size and age, and the picture that emerged was that this interaction does not lend itself to a simple conceptualization. We therefore believe it best to ignore the causal influence of organizational capital on actual appeal at this stage of theory development.

Obsolescence

The model developed in this paper predicts that old organizations have a low failure hazard, unless competition in the category is intense. This seems to contradict a perspective that claims that older organizations experience obsolescence (e.g., Barron, West, and Hannan, 1994; Carroll and Hannan, 2000; Sørensen and Stuart, 2000; Hannan, Pólos, and Carroll, 2007). The basic story holds that organizations are preselected at the time of founding to fit the prevailing environmental conditions and have an extremely limited ability to adapt to changing environments, particularly to the drift of audience expectations and tastes. As a result, appeal to audiences begins to decline with age at some age. According to this kind of argument, failure hazards increase with age for old organizations when inertial forces increase with age and the category audiences' tastes drift over time.

The construction leading to Theorem 1 could not address the obsolescence phenomenon, because it claims that actual appeal increases with age without being explicit about audience tastes. In our theoretical framework, the fit of the producer's offering with audience tastes is captured by the construct of intrinsic appeal. In this paper, intrinsic appeal played a very minor role, and we have not explicitly modeled its evolution over time. A companion paper (Le Mens, Hannan, and Pólos, 2011) develops an argument for obsolescence that focuses explicitly on the dynamics of audience tastes. The mechanism delineated in that paper relies on much weaker assumptions than does the previous theorizing. It also develops a (partial) integration with the results presented above.

Some Conjectures

Our main theorem makes predictions about patterns of age dependence in the risk of failure on the basis of the conditions at founding and hypotheses about the limiting behavior of organizational fitness. In this paper, we did not make predictions about the expected patterns of age effects in specific empirical contexts. We now do so in the form of a set of conjectures.

Theorem 1 makes different predictions according to whether initial fitness is above or below the threshold. We believe, however, that initial fitness will fall below the threshold in most empirical settings, such that organizational capital tends to decline in the early stages of the life of the producer. An interesting set of questions pertains to how long this stage of depletion lasts and what affects its duration. We refer to the age at which the stock of organizational capital starts to be accumulated as the "break-even" age. When the limiting fitness is below the threshold, the break-even age is never reached.

Density at founding. Numerous studies have shown evidence for a deleterious influence of the density at founding on the life chances of organizations (see the review by Carroll and Hannan, 2000). The argument that gave rise to studies of this effect relies on the observation that the resources available at the organizing stage of the creation of the producer tend to be slimmer when the density at founding is high and competition is expected to be intense (Carroll and Hannan, 1989). In terms of our model, this suggests that the initial stock of organizational capital will likely be lower when the density at founding is high. And because this initial stock has a long-term effect on the subsequent abundance of resources, this could explain why the density at founding has often been found to have a long-term effect on the hazard of failure.

More relevant to the density-at-founding argument, if high density at founding signals intense competition, we should expect the fitness threshold to be harder to reach for producers founded under conditions of high density. This, in turn, implies that the initial net flow of resources (the initial drift of the random walk) should be lower when the density at founding is high. In cases in which producers tend to achieve a positive net flow of resources after some time, this break-even age should be higher when the density at founding is higher. Another related conjecture is that the limiting net flow of resources should be lower when the density at founding is higher (and more likely to be negative when the density at founding is higher).

Characteristics of the founding team and early funding. The founding team brings human capital, reputation, and social networks (Shane and Stuart, 2002). Characteristics of the founding team have been shown to have durable effects on the fates of organizations (Burton, Sørensen, and Beckman, 2002). In terms of our framework, having a more competent founding team should lead to both a larger initial stock of organizational capital and high initial fitness relative to the threshold. We thus expect the break-even age to occur earlier when the founding team consists of more competent or experienced members or when the team has diverse social networks.

Initial organizational design. The degree of alignment of organizational designs with environments varies among newly founded organizations. For instance, Hannan et al. (2006) analyzed the effect of the founders' models of the employment relationship on the success and failure of high-technology firms in the Silicon Valley. They found that firms that chose an employment model that fits poorly with the industrial context impedes product development, at least partly because it diminishes the chances of recruiting and retaining highly talented scientists and engineers. In terms of our model, such poor initial alignment will increase the break-even age.

Future Research

We have emphasized the role of the conditions at founding, in the form of initial fitness, and of the competitive conditions, insofar as they affect fitness. To guarantee theoretical tractability, we made several simplifying assumptions. In particular, we assumed that producers targeted only one social position. A potentially interesting next step would be to extend our model to allow producers to target several audience segments. This would allow for consideration of the effect of the entry mode, such as de-novo entry, de-alio entry, or entry by spin-off or merger (Carroll and Hannan, 2000). Such an extended model could also be integrated with the most recent formulation of niche theory (Hannan, Pólos, and Carroll, 2007).

The importance of competitive conditions in our theory has some implications for the scope of applicability. The notion of a producer facing an audience and competitors has implications for many kinds of situations. For instance, it applies to firms competing in consumer markets, in labor markets, and capital markets, as well as to social movement organizations and political parties competing for members and political support. In each case, we can identify the relevant audience (consumers, occupational groups, intermediaries in capital markets, activists, voters, etc.), and we can characterize the intensity of competition in these "markets." Although this broad scope of application makes our approach appealing, the scope does not obviously generalize unchanged to a class of situations considered by some earlier work on age dependence: social relations such as marriage, friendship, or employment spells (Fichman and Levinthal, 1991). Their basic intuition that social relations can be seen as regulated by a stock of assets that can be depleted or replenished as a function of the quality of the joint experiences is similar to our view that the dynamics of the stock of organizational capital is regulated by organizational fitness. And the intuition that the expected quality of experiences increases over time thanks to learning and adaptation corresponds to our assumption of increasing appeal. In that sense, an analysis akin to ours could lead to basic predictions about age dependence in the termination hazard of social relations. It is harder to envision, however, what would correspond to the intensity of competition in an analysis of social relations as regulated by a stock of assets. We regard this limit on application to be the price paid for building the specificity needed to generate more precise theoretical implications about age dependence in organizational failures and provide an integrated theory for use in future research.

REFERENCES

Barron, D. N., E. West, and M. T. Hannan
1994 "A time to grow and a time to die: Growth and mortality of credit unions in New York City, 1914–1990." American Journal of Sociology, 100: 381–421.

Brüderl, J., and R. Schüssler
1990 "Organizational mortality: The liabilities of newness and adolescence." Administrative Science Quarterly, 35: 530–547.

Burton, M. D., J. B. Sørensen, and C. M. Beckman
2002 "Coming from good stock: Career histories and new venture formation." In M. Lounsbury and M. Ventresca (eds.), Research in the Sociology of Organizations, 19: 229–262. New York: Elsevier/JAI.

Carroll, G. R.
1983 "A stochastic model of organizational mortality: Review and reanalysis." Social Science Research 12: 303–329.
1985 "Concentration and specialization: Dynamics of niche width in populations of organizations." American Journal of Sociology, 90: 1262–1283.

Carroll, G. R., and J. Delacroix
1982 "Organizational mortality in the newspaper industries of Argentina and Ireland: An ecological approach." Administrative Science Quarterly, 27: 169–198.

Carroll, G. R., and M. T. Hannan
1989 "Density delay in the evolution of organizational populations: A model and five empirical tests." Administrative Science Quarterly, 34: 411–430.
2000 The Demography of Corporations and Industries. Princeton, NJ: Princeton University Press.

Carroll, G. R., and Y. P. Huo
1988 "Losing by winning: The paradox of electoral success by organized labor parties in the Knights of Labor era." In G. R. Carroll (ed.), Ecological Models of Organizations: 175–193. Cambridge, MA: Ballinger.

Carroll, G. R., and A. Swaminathan
2000 "Why the microbrewery movement? Organizational dynamics of resource partitioning in the U.S. brewing industry." American Journal of Sociology, 106: 715–762.

Delaney, K. J.
1999 Strategic Bankruptcy: How Corporations and Creditors Use Chapter 11 to Their Advantage. Berkeley and Los Angeles: University of California Press.

Denrell, J.
2004 "Random walks and sustained competitive advantage." Management Science, 50: 922–934.

Fichman, M., and D. Levinthal
1991 "Honeymoons and the liability of adolescence: A new perspective on duration dependence in social and organizational relationships." Academy of Management Review, 16: 442–468.

Freeman, J., G. R. Carroll, and M. T. Hannan
1983 "The liability of newness: Age dependence in organizational death rates." American Sociological Review, 48: 692–710.

Gimeno, J., T. B. Folta, A. C. Cooper, and C. Y. Woo
1997 "Survival of the fittest? Entrepreneurial human capital and the persistence of underperforming firms." Administrative Science Quarterly, 42: 750–783.

Greene, W. H.
2003 Econometric Analysis, 5th ed. New York: Macmillan.

Hajivassiliou, V. A., and P. A. Ruud
1994 "Classical estimation methods for LDV models using simulation." In R. F. Engle and D. L. McFadden (eds.), Handbook of Econometrics, 4: 2383–2441. Amsterdam: Elsevier.

Hambrick, D. C., and R. A. D'Aveni
1988 "Large corporate failures as downward spirals." Administrative Science Quarterly, 33: 1–23.

Hannan, M. T.
1998 "Rethinking age dependence in organizational mortality: Logical formalizations." American Journal of Sociology, 104: 126–164.

Hannan, M. T., J. N. Baron, G. Hsu, and Ö. Koçak
2006 "Organizational identities and the hazard of change." Industrial and Corporate Change, 15: 755–784.

Hannan, M. T., G. R. Carroll, S. D. Dobrev, and J. Han
1998 "Organizational mortality in European and American automobile industries part I: Revisiting the effects of age and size." European Sociological Review, 14: 279–302.

Hannan, M. T., L. Pólos, and G. R. Carroll
2007 Logics of Organization Theory: Audiences, Codes, and Ecologies. Princeton, NJ: Princeton University Press.

Le Mens, G., M. T. Hannan, and L. Pólos
2010 "On the dynamics of organizational mortality: Age dependence revisited." Research Paper 2062, Stanford Graduate School of Business.
2011 "Drifting tastes, inertia, and organizational mortality." Working paper, Department of Economics and Business, Universitat Pompeu Fabra.

Levinthal, D.
1991 "Random walks and organizational mortality." Administrative Science Quarterly, 36: 397–420.

Nelson, R. R., and S. G. Winter
1982 An Evolutionary Theory of Economic Change. Cambridge, MA: Harvard University Press.

Pólos, L., and M. T. Hannan
2002 "Reasoning with partial knowledge." Sociological Methodology, 32: 133–181.
2004 "A logic for theories in flux: A model-theoretic approach." Logique et Analyse, 47: 85–121.

Ranger-Moore, J.
1997 "If bigger is better, is older wiser? Organizational age and size in the New York life insurance industry." American Sociological Review, 58: 901–920.

Rosenzweig, P.
2007 The Halo Effect. New York: Free Press.

Ryder, N.
1965 "The cohort as a concept in the study of social change." American Sociological Review, 30: 843–861.

Shane, S., and T. E. Stuart
2002 "Organizational endowments and the performance of university start-ups." Management Science, 48: 154–170.

Sørensen, J. B., and T. E. Stuart
2000 "Aging, obsolescence, and organizational innovation." Administrative Science Quarterly, 45: 81–112.

Stinchcombe, A. L.
1965 "Organizations and social structure." In J. G. March (ed.), Handbook of Organizations: 142–193. Chicago: Rand-McNally.

Tuma, N., and M. T. Hannan
1984 Social Dynamics: Models and Methods. Orlando, FL: Academic Press.

Vaupel, J. W., K. G. Manton, and E. Stallard
1979 "The impact of heterogeneity in individual frailty on the dynamics of mortality." Demography, 16: 439–454.

APPENDIX A: Proofs

Testing what follows from the premises, in a stage of a theory in the nonmonotonic logic we use, operates on representations of arguments in the form of "rule chains." The links in these chains are strict rules, definitions, auxiliary assumptions, and causal stories. The chains start with the subject of the argument and terminate with the purported conclusion of the argument (the consequence to be derived). In nonmonotonic inference, different rule chains, each representing an argument embodied in the state of the theory, might lead to opposing conclusions. The testing procedure determines whether any inference can be drawn at all and, if so, which one. Such testing requires standards for assessing whether a pair of relevant rule chains is comparable in specificity and for determining specificity differences for comparable chains. In this paper, the available premises and definitions all point in the same direction; we do not see any rule chains that point to opposing conclusions. Thus all that is required is that we establish a rule chain that connects the antecedent and consequent in a claimed theorem or proposition.

Proposition 1. The only applicable rule chain in this stage of the theory for Case A uses Post. 2, Aux. 1, Aux. 2 and the fact that the function $h(.)$ is increasing. Case B follows immediately from Case A and Def. 1. Case C follows immediately from Case B and Post. 1.

Proposition 2. This is an immediate consequence of Post. 3 and Aux. 3.

Proposition 3. Part A follows from the rule chain supporting Prop. 2 and the definition of fitness (Def. 2). Part B follows from Prop. 2 and Def. 3.

Theorem 1

Case A. The antecedent gives $\phi_x(\tau_x) > f_x$. Let t_1 and t_2 be two time points such that $\tau_x \leq t_1 < t_2$. Prop. 3 implies that for all $t > \tau_x$, $\phi_x(t) > f_x$. Then, Prop. 1 implies that $\kappa_x(t_1) < \kappa_x(t_2)$ and $\omega_x(t_1) > \omega_x(t_2)$.

Case B. The antecedent in this case gives $\phi_x(\tau_x) < f_x$ and $\vec{\phi}_x > f_x$. Prop. 3 implies that fitness presumably grows with age and becomes higher than f_x. Therefore there exists $q > \tau_x$ such that, if $t < q$, then $\phi_x(t) < f_x$ and if $t > q$, then $\phi_x(t) > f_x$. Let t_1 and t_2 be two time points such that $\tau_x \leq t_1 < t_2 < q$. Prop. 1 implies $\kappa_x(t_1) > \kappa_x(t_2)$ and $\omega_x(t_1) < \omega_x(t_2)$. Let t_3 and t_4 be two time points such that $q < t_3 < t_4$. In this case, Prop. 1 implies $\kappa_x(t_3) < \kappa_x(t_4)$ and $\omega_x(t_3) > \omega_x(t_4)$.

Case C. Now the antecedent gives $\phi_x(\tau_x) < f_x$ and $\vec{\phi}_x \leq f_x$. Prop. 3 implies that fitness presumably remains lower than f_x; and Prop. 1 then implies that organizational capital declines with age and the hazard increases with age.

APPENDIX B: The Likelihood Function and the Unobserved Heterogeneity Distribution

Computation of the Likelihood Function

Knowledge of the probability density of failure times of a random walk with time-varying drift is necessary to proceed to maximum likelihood estimation of the parameters of our model (described by equations 1 and 4). Because we are not aware of any analytical formulation of this density, we used computer simulations instead. The idea is to simulate the evolution of the stock of organizational capital for a large number of independent producers and to compute the probability density of the failure times on the basis of these simulations. In this appendix, we provide some details about the stages of the simulations for our model with time-varying and bounded drift (characterized by equations 1 and 4). A similar procedure is used for the simpler model formulations. This goes as follows:

1. Generate a table of random shocks $\varepsilon_{i,t}$, with $1 \leq i \leq 10{,}000$ and $1 \leq t \leq 200$. The random shocks are independent realizations of a random variable with

normal distribution with a mean of 0 and a variance of 1. Each row of this table is a 200 period-long time series of random shocks for one simulated producer. There are 10,000 independent simulated producers.

2. Set the model parameters ($\kappa_x(\tau_x)$, $\mu_x(\tau_x)$,γ, β) at some specific values.
3. For simulated producer *i*, compute the evolution of the stock of organizational capital according to the random walk specified by equations 1 and 4, using the sequence of random shocks $\varepsilon_{i,t}$, for $1 \le t \le 200$. The failure date for producer *i* is the first period in which organizational capital reaches zero or a lower value. These simulations produce a vector of 10,000 failure times. This vector of failure times is used to estimate the probability density of failure times given the model parameters (by computing the histogram of the simulated failure times) as well as the survivor function given the parameters. The (simulated) probability density is denoted by $f(\cdot \mid \kappa_x(\tau_x), \mu_x(\tau_x), \gamma, \beta)$, and the (simulated) survivor function is denoted by $G(\cdot \mid \kappa_x(\tau_x), \mu_x(\tau_x), \gamma, \beta)$.
4. Use the simulated density of failure times to compute the likelihood of the observations. Suppose producer *j* (in the data) fails at age a_j. Then the likelihood of this observation is $f(a_j \mid \kappa_x(\tau_x), \mu_x(\tau_x), \gamma, \beta)$. If producer *j* does not fail, let a_j be the number of years *j* is in the data, then the likelihood of this observation is $G(a_j \mid \kappa_x(\tau_x), \mu_x(\tau_x), \gamma, \beta)$. If there are *n* observations in the data, the likelihood of the data is the product of the individual contributions of the *n* observations in the data.

To implement the maximum simulated likelihood (MSL) procedure (Hajivassiliou and Ruud, 1994), we used "fminsearch" in Matlab R14 (a nonlinear unconstrained optimization routine based on the simplex search method) to minimize the opposite of the simulated likelihood function. We computed the asymptotic covariance matrix of the estimates as the outer product of gradients, an estimator that is otherwise known as the BHHH estimator (see Greene, 2003: 481). Following standard practice, we used the simulated log-likelihood as a basis to compute the gradients.

The following additional comments might be helpful to readers who would like to adopt the above procedure for their own model estimations:

1. To help stabilize the procedure, the table of random shocks (stage 1 above) was generated only once, before running the optimization routine to fit the parameters of the various model specifications. Generating a new table of random shocks at each evaluation step would make the gradient search used by most optimization routines rather unstable, which in turn would make the convergence of the optimization routine problematic.
2. Using simulations to compute the likelihood function can lead to computational problems due to outliers. For example, despite the fact that the distribution of failure times has infinite support for most parameter values, the simulated likelihood function has finite support. This implies that outliers in the data might have a simulated likelihood of 0. This is an issue because, when this happens, the likelihood of the dataset is zero, given the multiplicative nature of the likelihood function. To avoid the risk of running into this problem, we allocated a very small but positive likelihood to observations that would have been assigned a simulated likelihood value of 0. More precisely, we assigned a likelihood of 10^{-8} to the observations that corresponded to a simulated likelihood value of 0. Changing this adjustment to higher or lower values did not substantially affect model estimations, however. Besides, further examination of the likelihood vector showed that for most model specifications, the simulated likelihood of each data point is different from zero.
3. Because it is likely that the simulated likelihood function has several local optima, it is useful to check the robustness of the results by running the optimization routine from multiple initial parameter values. To select initial parameter values in a systematic fashion, we started by running a grid search over a set of potential parameter values. Then, we used the top 1 percent parameter combinations as starting points for the optimization routine. The parameter estimates reported in table 1, in the text, are those that lead to the maximal value of the likelihood. As an illustration, we provide additional details about the procedure we used to estimate the parameters of the model with bounded drift (column 4 in table 1). We started by running a grid search with the following parameters: $\kappa_x(\tau_x) \in \{1,1.5,2,...,7\}$; $\mu_x(\tau_x \in \{-3,-1.75,-1.5,...,1\}$; $\gamma \in \{-1,-1.8,-1.6,...,2\}$; $\beta \in \{0,.2,.4,...,2\}$. Then we ranked the parameter combinations in terms of the corresponding likelihood, and we selected the top 1 percent combinations

(about 380 combinations). We used each of these parameter combinations as a starting point for the Matlab nonlinear optimization routine. Running the optimization routine produced a new set of parameters and associated likelihood. The parameter estimates reported in table 1 are the parameters that produced the maximal likelihood value.

Controlling for Unobserved Heterogeneity

To compute the likelihood function when assuming gamma-distributed unobserved heterogeneity, we build on the approach used by Levinthal (1991; see also Vaupel, Manton, and Stallard, 1979). Let $h(t \mid \mathcal{P})$ denote the simulated hazard function, given the parameters of the model and no control for heterogeneity and let $H(t \mid \mathcal{P}) = \int_{u=0}^{t} h(u \mid \mathcal{P}) du$ be the cumulative hazard function. The probability density function of failure times, with gamma-distributed heterogeneity, with parameters *k* and 1/*k*, is given by

$$f(t \mid \mathcal{P}, k) = h(t \mid \mathcal{P}) \left(\frac{k}{k + H(t \mid \mathcal{P})} \right)^{k+1}.$$

At each estimation step, we first computed the likelihood of the data as follows. We started by computing the probability density of failure times, given the parameters of the model ($\mathcal{P}$), but without controlling for unobserved heterogeneity. This was done using computer simulations, as described in the previous section. Based on this simulated density, we computed the simulated hazard function $h(t \mid \mathcal{P})$ and the simulated cumulative hazard function $H(t \mid \mathcal{P})$. Then we used the above formula to compute the simulated probability density of failure times, controlling for gamma-distributed heterogeneity. This density was then used to compute the simulated likelihood of the data.