Minnesota Journal of Business Law and Entrepreneurship

Volume 2026, No. 3

ISSN: 1540-3270

The Ethics of Scale: Legal Boundaries of Hypergrowth in Venture-Backed Startups

[1]Dr. Kusum Gupta, [2]Dr. Sunil Kumar, [3]Ms. Smita Kaushik, [4]Dr. Geeta Rani

[1]Professor, Jagan Institute of Management Studies

[2]Assistant Professor, Jagannath University, Bahadurgarh

[3]Assistant Professor, Jagannath University, Jaipur

[4]Associate Professor, Jagannath University, Bahadurgarh

Abstract

This growth is generally talked about as a strategy that some businesses undertake at the wrong level and the associated failures are blamed on the nature of specific founders. In this paper, I propose that the pathologies of scale are not dispositional but structural, and are the result of three mechanisms that are neither intentional nor a matter of will. The first is the variance mandate – the optimal level of firm risk from a venture fund perspective is higher than under a founder perspective, and higher still than if a third-party had to decide, because the outcomes for the firm are concentrated at both extremes of the range, and it doesn't have to accept the risk. The second is control dilution; if the number of employees grows exponentially, then the percentage of staff with any significant tenure is a mathematical function of the rate of growth and under that growth regime, a firm with an annual growth rate of 2 is in constant transition, with a majority of workers in it being there for less than a year, and the informal norms that replaced process no longer pass through it. The third is harm scaling: an error rate that is ethically unremarkable when it is made on a thousand users is not the same when it is made on a hundred million when the users do not choose to be there. We then discuss the poor nature of the legal boundary as a constraint on hypergrowth, in that the legal instruments it embodies are ex post and enforced on a case-by-case basis, while growth is exponential, which means the accountability gap expands mechanically. Seven propositions are provided, along with an analytical protocol. No artificial statistics are given.

Keywords: hypergrowth; venture capital; business ethics; risk imposition; corporate accountability; organisational control; judgment-proofing; startup governance

1. Introduction

Much of the literature of scale consists of exhortatory literature. Growth is said to be necessary, speed a competitive requirement, and an openness to “chaos” within the organisation a sign of seriousness; the strongest arguments for this position suggest that in what are called “winner take most” markets the risks of growing too slowly outweigh the dangers of growing too quickly (Hoffman & Yeh, 2018). On the other hand is a list of common and recurring failures: safety incidents in companies that grew more quickly than their control systems, jobs that would never have been offered to workers in companies of a tenth the size, products shipped to populations that would never have been approved by a mature company, and financial statements provided to private investors that would never have been allowed to be public. Character is always the reason given for the catalogue. Some founders are reckless, some cultures are toxic, and some investors indulge them all. This explanation is comforting as it suggests a solution of choice: find better people. It is also, we argue, for the most part incorrect, or at least incorrect as far as the mechanics. The pathologies are constantly repeated in different companies, industries and various jurisdictions in a pattern that can be no coincidence and that can only be attributed to an individual disposition in the founders. When a pattern holds up to changes in character of the people who create it, the explanation is structural.

In this paper three constructions are identified. The first is related to the capital. The empirical literature on venture returns and, more explicitly, in practitioners' reports on the nature of how they make decisions, reports that venture

funds have extremely few outcomes that yield returns to their investors (Cochrane, 2005; Gompers et al., 2020; Korteweg & Sorensen, 2010). A portfolio company that steadily ascends towards a reasonable end result to a moderate result adds almost nothing to the fund's performance - from the fund side it is virtually indistinguishable from failure. It follows immediately that the best instruction for a portfolio company to be given by the fund is to follow a strategy which maximises the right tail; that this is optimal at a variance level well above what an undiversified founder would wish to, and much higher than what an affected third party would wish to. This is what we refer to as the variance mandate and is, we believe, what most of what is said around the term founder recklessness is based on.

The second is related to organisation. The scaling literature itself has a corresponding arithmetic implication as far as the composition of the labor force.There is an arithmetic implication for the composition of the labor force as far as growth of the kind that is celebrated in the scaling literature. As the number of people grows exponentially, the percentage of people who have been present for more than a certain period decreases with the higher the growth rate and about half the people in an organisation with annual doubling have been present for less than a year at any time. Much of what regulates the conduct of a small firm is communicated informally, by example, and as the number of experienced people to inexperienced people grows, the ratio of inexperienced to experienced becomes so great that it begins to erode the control environment, not because of a lack of management attention, but because of a mathematical fact. The managerial constraint on growth rates was first recognized by Penrose more than 60 years ago, and the scaling literature of the present day has basically assumed the constraint to be lifted (Penrose, 1959).

It is the third that has to do with the moral meaning of number. An error rate is not a constant value from an ethical point of view. A "user problem" affecting 1 out of 1000 users is a "support problem" for a firm with 1000 users and a "population problem" for a firm with 100 million; and this change is not just quantitative: it's a risk imposed upon people who have no practical alternative, no knowledge of the imposition, no capacity to bargain about it. The difference between risks taken by a person and risks imposed on the person is a crucial one in risk ethics, as is the difference between the necessity to justify risks imposed on the person, and risks that the person takes (Hansson, 2003, 2013). It is a system of structuring the venture that systematically shifts it from the first to the second as the firm grows, as those willing to take the risk are the founders and investors and those bearing the realisations, are not. The paper then develops each mechanism; why the legal boundary set does not perform well against each mechanism; and finally what would actually constrain the behaviour. The obvious objection to this analysis is that this analysis is proof too much, and will prevent ordinary commercial risk-taking, and it is addressed in Section 7.

2. The Variance Mandate

2.1 From fund structure to firm instruction

Venture funds are organized as partnerships with a fixed lifespan, with the manager taking a percentage on returns above a threshold (the manager's return) alongside a management fee, and have been extensively researched in terms of their economics and the nature of the manager's return (Metrick & Yasuda, 2010; Sahlman, 1990), with the share of profits having option-like properties with respect to the portfolio outcome. There is extreme skewness in returns; that is, a few investments account for the bulk of a fund's return; the rest make up a small fraction of the fund's total return. The concentration under a declared power-law parameter is shown in Figure 1, panel (a).

Minnesota Journal of Business Law and Entrepreneurship

Volume 2026, No. 3

ISSN: 1540-3270

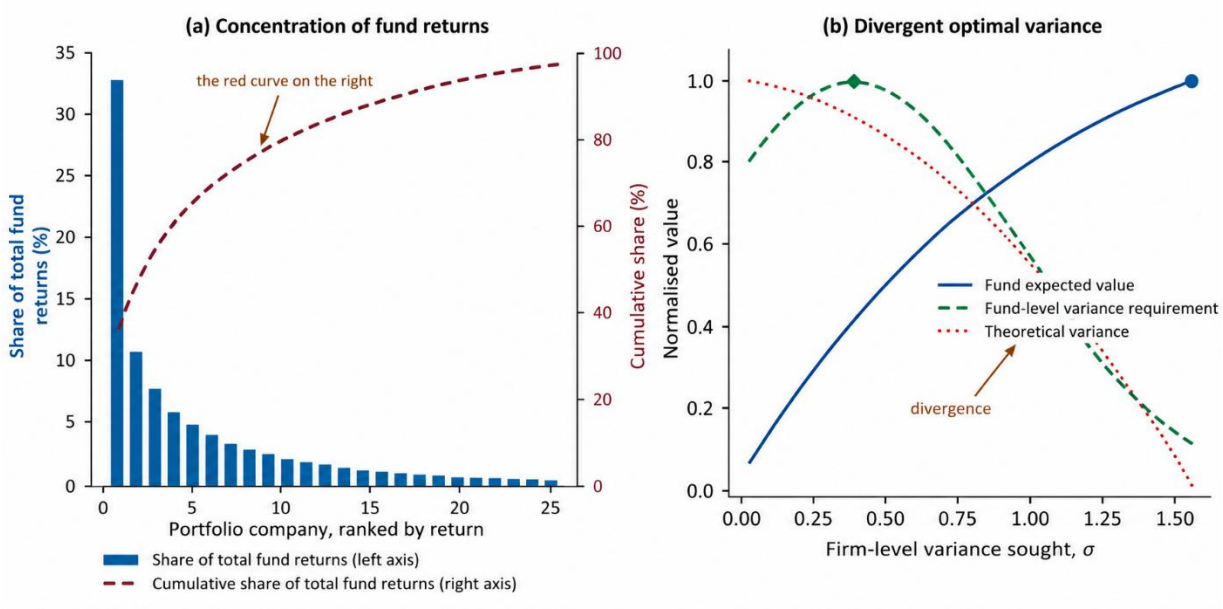

Figure 1. The top graph shows the concentration of fund returns for a portfolio with a known power-law parameter. Affected third parties and the fund's and the founder's optimal firm level variance are shown in panel (b). All the values assumed are those that are declared and represent the structure of the argument only and not the actual values, so these are not empirical values.

The consequence for the individual portfolio company is the paper's starting point. The payoff of the fund is convex in the payoff of any individual investment, so that an increase in variance of the payoff of any individual investment increases the expected value of the fund even if it decreases the expected payoff of the company. In contrast, the founder's position is risk averse, because their human and financial capital is also invested in one undiversified position, and the large non-diversifiable risk of the entrepreneurs has been estimated in the finance literature (Hall & Woodward, 2010). The divergence is shown in the diagram in Panel (b) of Figure 1. Third parties sit on a position even further out, as they don't have the upside and share the downside, their preferred variance is lower than both.

None of this is done with bad faith. Investors who are calling on the company to accelerate, go to an unconsolidated market before the first market is consolidated, or delay investment in compliance infrastructure are giving good advice from their own vantage point. The founder who resists is not being prudent in a manner that the fund can reward, since in the fund's returns it is not possible to distinguish a failure and a success that is not good. The translation from the fund structure to the operating instruction is mapped in Table 1.

Table 1. From fund structure to firm-level instruction

Structural feature	Mechanism	Instruction transmitted to the firm	Whose interest it serves	Countervailing pressure available
Concentrated returns	Fund performance depends on the extreme tail	Pursue outcomes that could be very large, accept high failure probability	Fund; not the founder or third parties	None internal to the structure
Fixed fund life	Returns must be realised within a defined horizon	Compress the timeline; prefer speed to consolidation	Fund and its own investors	Secondary markets; continuation vehicles

Structural feature	Mechanism	Instruction transmitted to the firm	Whose interest it serves	Countervailing pressure available
Carried interest	Manager payoff is convex above a threshold	Increase variance rather than expected value	Fund manager specifically	Clawback provisions; reputation
Staged financing	Capital released against milestone attainment	Optimise for the metric that unlocks the next round	Fund, as a control device	Milestone design; board judgement
Follow-on concentration	Later capital flows to apparent winners	Demonstrate growth even where growth is unprofitable	Fund and the company jointly, in the short run	Discipline in capital markets, when it appears
Portfolio diversification	Fund holds many positions; founder holds one	Accept firm-level risk the founder would decline	Fund	Founder secondary sales, which reduce the asymmetry

2.2 What the mandate does not explain

There are two points in the argument which cannot be exaggerated. First, the same structure generates the risk appetite that enables one to fund truly worthwhile experimentation. That entrepreneurship is an activity involving experimentation under uncertainty is the reason why it is socially useful to have a capital source as tolerant as venture capital, whose failure rate is too high for other types of capital, but whose expected return is positive. Any doctrine that outlaws variance-seeking would outlaw the process that underpins most radical innovation.

Secondly, the requirement is for a “level of risk” and not for an incidence of risk. There is no objection to a variance looking to a fund that the downside is on the fund and the consenting parties, namely the fund's founders. The issue is where the strategies that lead to right-tail outcomes are strategies whose downside lies in another place: with consumers that are exposed to an under-tested product; with workers that are exposed to practices that were adopted in order to avoid compliance; with counterparties that were led to believe that a firm was competent when it actually wasn't; or with the public, where a firm becomes systemically significant before being competent. It is the presence of distinction between risk taken and risk imposed, as developed in section 5, that constitutes the ethical work, not the presence of risk.

3. Control Dilution

3.1 The arithmetic

Think of a company that is experiencing a steady growth rate in the number of employees. The proportion of workers employed for more than a certain period T is an exponential function decreasing in the product of the growth rate and T in steady state. The message is clear and is summarized in figure 2. Immediately, about half the organisation has been around for less than a year, on a permanent basis, at annual doubling. Approximately 2/3 at tripling. This is not a state of transition to one where hiring will "catch up" with it, but rather a state of the growth rate itself.

Minnesota Journal of Business Law and Entrepreneurship

Volume 2026, No. 3

ISSN: 1540-3270

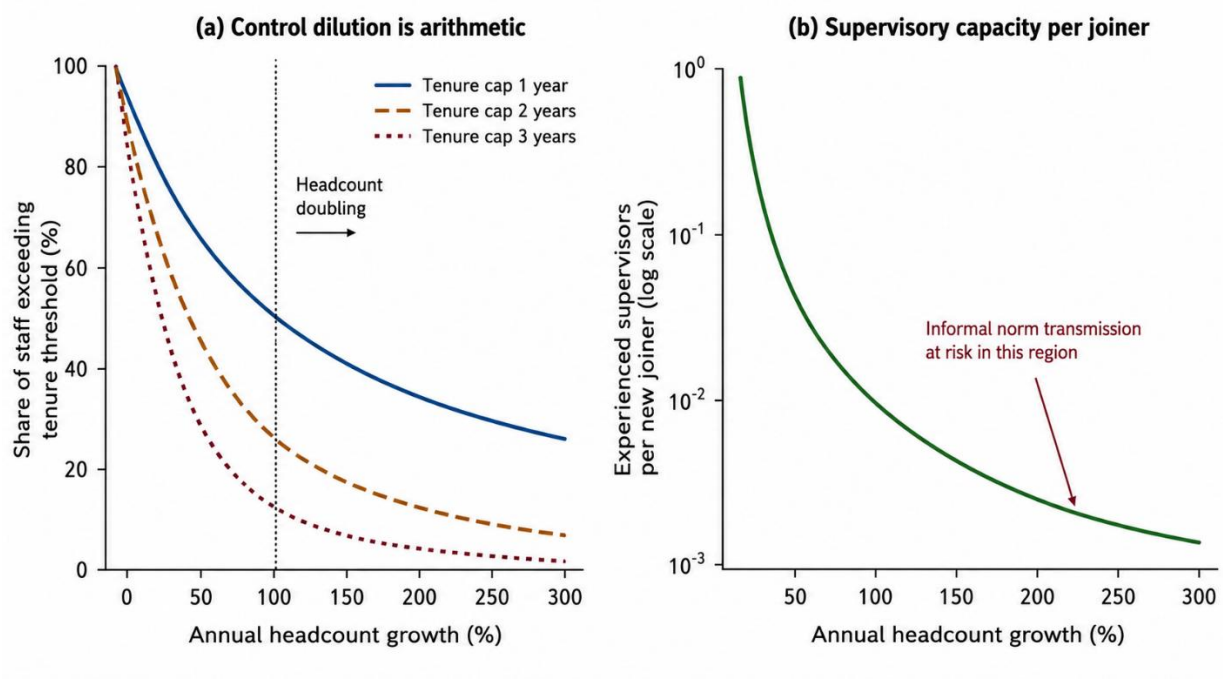

Figure 2. As a steady exponential growth rate, panel (a) is the proportion of staff who have a tenure above a threshold as a function of the annual increase in the number of staff. Panel (b) - Person years available to new joiners. Values are calculated from the given growth relation and from declared parameters, they are only examples and not empirical estimates.

Why this matters is a question concerning the way people actually were controlled in small businesses. Most of what makes for the prevention of misconduct in a small organisation is observation: one sees how things are done, absorbs what is acceptable, and is corrected swiftly when one deviates, and the behaviour of the founders is immediately visible to all. This is a cost-effective and effective control system with strict scaling limitations as it needs ratio of experienced persons/novices to pass the information. The rate of decline of available supervisory capacity per joiner is plotted in panel (b) of Figure 2. Thereafter, the informal system ceases to operate and if no formal system has been developed during this period of time there is nothing in place to regulate the system, only the formal system that existed when the company was smaller.

Table 2. Mechanisms of control dilution under exponential growth

Mechanism	What degrades	Why growth causes it	Observable symptom	Countermeasure and its cost
Norm transmission failure	Informal understanding of acceptable conduct	Ratio of experienced to new staff collapses	Divergent practice across teams doing the same work	Explicit codification; slows onboarding
Supervisory dilution	Depth of review of individual decisions	Managers supervise more people, most of them new	Decisions taken without review that previously received it	Managerial hiring ahead of growth; expensive and unpopular
Process debt	Fit between procedure and current scale	Procedures designed for a firm an order of magnitude smaller	Workarounds become the actual process	Periodic re-engineering; disruptive

Mechanism	What degrades	Why growth causes it	Observable symptom	Countermeasure and its cost
Escalation path loss	Reliability of concerns reaching decision-makers	Layers added faster than escalation norms adapt	Concerns raised locally and not transmitted	Independent reporting channel; cultural resistance
Founder visibility loss	Direct observation of the founders' own conduct	Most staff never interact with the founders	Culture described in terms nobody can operationalise	Explicit statement of standards; loses authenticity
Normalisation of deviance	The boundary of acceptable practice	Repeated uneventful deviation redefines the norm	Prior exceptions treated as precedent	Deviation logging and review; requires slack

3.2 Normalisation of deviance at speed

The last line of Table 2 specifies how control dilution can turn into catastrophe - not disorder. The Challenger decision is described by Vaughan (1996) as a series of increasingly redefined deviations, none of which had an adverse consequence, to arrive at a point where an organisation makes a fatal decision, in each of which they could have offered reasonable arguments and reasons that they were not being reckless. The disaster literature in general focuses on the way that such failures have incubation periods over which warning signs build up and are reinterpreted as a problem (Perrow, 1984; Turner & Pidgeon, 1997).

Each of these elements is enhanced by hypergrowth. More people do more things without much guidance and, hence, the deviations are more frequent. With a higher volume, there are more uneventful outcomes. What Figure 2 reveals is the scarcity of institutional memory to recognise a deviation as a deviation. Additionally, an organisation's learning depends on the outcomes it observes, and rapidly growing firms observe an unrepresentative sample since the time between a decision and the result often entails more time than it takes for the organization to change shape (Levinthal & March, 1993; March, 1991). The adage literature exacerbates this by learning from survivors only and by the observation that firms that scale aggressively tend to be more successful, which is confounded by the fact that those firms may no longer exist to research (Denrell, 2003).

3.3 Attrition and the two-population problem

Figure 2 is conservative as it makes no allowance for attrition. Fast growing companies tend to have a lot of voluntary departures and departures are not random; they target those who joined in the beginning, where equity is vested, and where job roles have evolved from what they had originally envisioned when they joined the company. Attrition thus skews the realised composition of that group for which the informal control system relies, and the growth rate is not sufficient to capture the severity of the real effect.

The other, lesser-known effect is the formation of two groups inside of one firm. The early cohort has a common understanding of the reasons for doing things the way they do and knows the reasons behind rules for which the reason was never recorded. The latter group stores the rules without the rationalization, and in many cases, without the rule. These populations are different systematically in terms of equity, access to founders, and in terms of what the firm is for – and they are often used on the same decisions. The latter is not just an inconsistency: it is a specific failure mode, a rule that has no obvious purpose and will be ignored by a good human being when it gets in the way of an obvious good, and it is this that is the entry point for the normalisation process described above.

Organisational research has long identified qualitatively different stages of the life cycle, with the coordination mechanisms that are successful in one stage becoming the main problem of the next, and the transitions as crises, not smooth adaptations (Greiner, 1972). The amount of time between hypergrowth changes is what is different. An organisation that traverses the same set of transitions in eighteen months will hit each of these transitions

without having consolidated the previous one, and the burden of newness that accompanies young organisations is increased by quick expansion rather than it being solved, according to the literature (Aldrich & Ruef, 2006).

4. Harm Scaling

The third mechanism is related to the moral connotations of number and it is this that is most frequently discussed as if it were only a matter of degree. The point is that they are unequal, as outlined in Table 3, meaning that a firm's overall harm exposure doesn't depend on the number of users.

Table 3. How harms scale with the number of people exposed

Harm category	Scaling with population	Why	Implication for a growing firm	Example domain
Individual defect harm	Linear	Each affected person is harmed independently	Predictable; provisionable	Product malfunction; billing error
Aggregate exposure to a rare event	Linear in expectation, but tail risk grows	More draws from the same distribution	Rare events become certainties at scale	Safety incidents in physical services
Network-mediated harm	Superlinear	Harm propagates between affected parties	Escalates faster than the user base	Misinformation; fraud on a platform
Market-structure harm	Threshold	Harm appears only past a share or dependence level	Nothing observable until a boundary is crossed	Below-cost pricing; dependence of suppliers
Norm displacement	Superlinear and often irreversible	Practice becomes the expected standard	Firm changes the baseline against which it is judged	Employment classification; data practices
Systemic dependence	Threshold, then catastrophic	Others build on the firm's continued operation	Failure becomes a public rather than private event	Payments; infrastructure services

Three consequences follow. First the remediation capacity of the firm itself is not linked with any of these. The capacity to provide remediation increases more or less in proportion to the resources applied, and the resources grow in a linear manner, at best, as the number of users grows. The accountability gap is the result and is shown in Fig. 3, and we stress that this is a mechanical rather than a moral phenomenon, in that a gap opens even though no one chose to let it.

Minnesota Journal of Business Law and Entrepreneurship

Volume 2026, No. 3

ISSN: 1540-3270

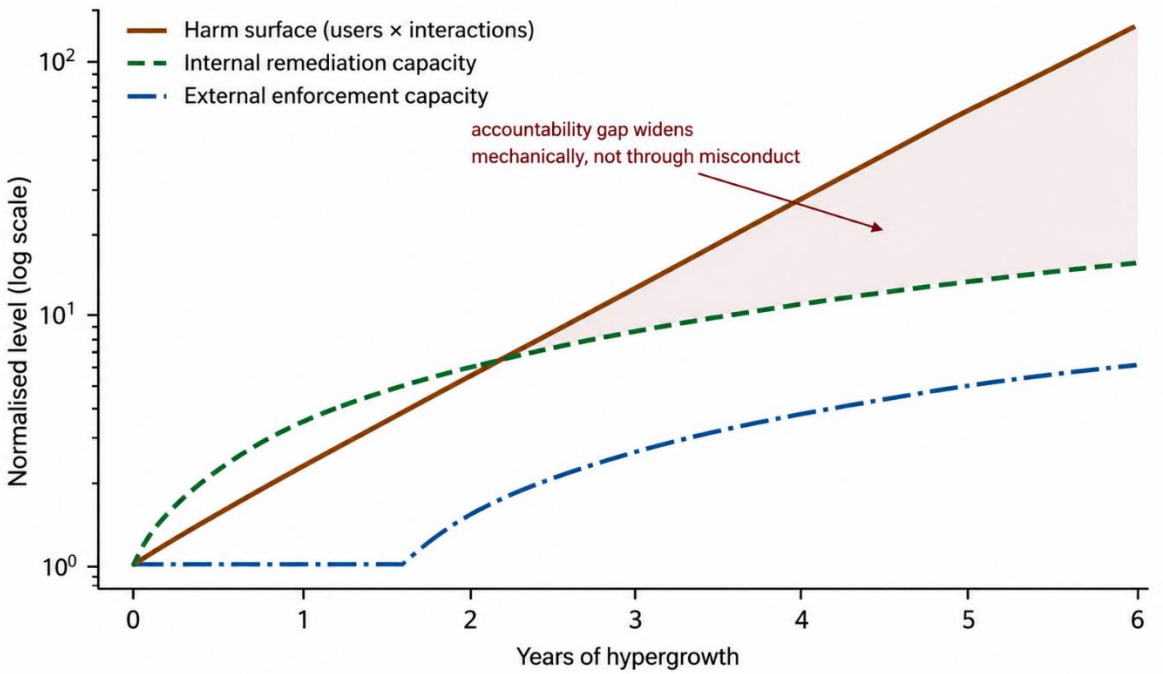

Figure 3. The accountability gap. Harm surface grows exponentially with scale while internal remediation capacity and external enforcement capacity grow at best linearly, so the gap widens without any misconduct. Values are generated from declared parameters and illustrate the structure of the argument only; they are not empirical estimates.

A fourth consequence is related to the right decision rule in uncertain situations. When a harm is reversible, an error is a cost, or a permanent transfer, in the latter case, the asymmetry calls for a different approach to uncertainty. The maximization of precautionary reasoning has been criticized with force on incoherence as such a general principle, because nearly any precaution has risks that would be equally condemned by the precautionary principle (Sunstein, 2005). The criticism does not do the job here of killing a narrower claim: that some are irreversible harms and some are not, that it is often possible to tell in advance that something is an irreversible harm, and that a decision procedure that does not give equal consideration to irreversible harms will systematically underweight the harms that are irreversible. In a hypergrowth company, the answer to this question can be called into question at any time: if this is not right, can it be reverted, and by whom. Second, the categories that scale superlinearly or by threshold are exactly the ones least apparent from the usual metrics. If a firm is keeping a close eye on the defect rate, it will become apparent when the threshold is breached and will see the linear category; it will not see norm displacement or systemic dependence until the threshold is breached, after which the response options will be more constrained.

Third, the degree of reversibility of a harm is more important than its level and decreases as scale increases. What's one time a practice adopted by one firm with a thousand users can be unwound, but a hundred million users is a different story—it's a public policy matter, not a product decision. This is the empirical fact that has been stated in the philosophical literature on technology, namely that the scope of contemporary action exceeds the scope of the moral framework constructed for actions which are local and reversible (Jonas, 1984).

5. Risk Taken and Risk Imposed

The ethical analysis is based on a distinction that is usually not questioned in commercial formulations. No one has to say why someone should take a risk with his own resources, when they are at risk, knowing the odds. A person who imposes a risk on others needs a justification and the justification shouldn't be provided by the fact

that the overall value of the risk to all parties is positive, since the overall value also covers the gain to the other parties to which the risk is not imposed (Hansson, 2003; 2013). Table 4 does that for the categories of party that is affected by hypergrowth.

Table 4. Consent, benefit, and control across affected parties

Party	Consented to the risk?	Shares the upside?	Can exit?	Can observe the risk?	Ethical status of imposition
Founders	Yes, explicitly	Yes, substantially	Formally yes; practically constrained	Yes	Unproblematic
Investors	Yes, professionally	Yes	Yes, through secondary sale	Largely yes	Unproblematic
Early employees	Partially; often without accurate information	Nominally, through illiquid equity	Yes, at the cost of the equity	Partially	Requires disclosure to be legitimate
Later employees	No; they joined a different firm than the one described	Marginally	Yes	Rarely	Requires justification
Users	No; terms of service are not consent to product risk	Only through the service itself	Depends on substitutes	Almost never	Requires justification and often fails it
Third parties and the public	No	No	No	No	Illegitimate absent a strong public-benefit justification
Suppliers and dependent firms	Partially, by contract	Through the commercial relationship	Only at the cost of the relationship	Partially	Requires proportionality in dependence

There are two additional ethical frameworks that add to the challenge. The first is the problem of many hands: When an outcome is a result of a large organisation, which itself is the product of the joint actions of many people, none of whom could individually cause the outcome, nor could any individually prevent it, then it is hard to pinpoint individual moral responsibility for the outcome, though it is easily attributed to the organisation as a whole (Thompson, 1980). This is the usual state, and hypergrowth will maximise the number of hands and minimise the time spent per hand.

The second is judgment-proofing. However, a firm's ability to repair the damage it inflicts is limited by its assets, with damages exceeding assets going to the victim rather than the firm, so that damages do not deter at the point where they become large (Shavell, 1986). Venture-backed firms are unique in that they often are judgment-proof, but also have very high nominal valuations, since the valuation is based on the option value on future outcomes and not the amount of assets which can be used to satisfy a judgment. The erosion of effectual liability as a result of asset structuring arrives at a similar conclusion regarding the limits of the tort system as a restraint on large-scale harm (LoPucki, 1996), and the literature on imposing penalties on corporations through the tort system has

long pointed to the problem of crafting sanctions that are both penal and feasible for enforcement in an entity that cannot be imprisoned and may not be able to pay (Coffee, 1981). This structure is the one to which proposals to extend tort liability beyond the corporate boundary respond (Hansmann & Kraakman, 1991).

6. Why the Legal Boundary Set Underperforms

Table 5 sets out the principal legal constraints on hypergrowth and assesses each against the mechanisms above. The pattern is consistent: the instruments are predominantly ex post, individually enforced, and calibrated to firms whose scale changes slowly.

Table 5. The legal boundary set assessed against hypergrowth

Instrument	Timing	Enforcement architecture	Scaling property	Why it underperforms against exponential growth
Tort and product liability	Ex post	Individual claims	Claims scale linearly with claimants	Slower than exposure growth; defeated by judgment-proofing
Consumer protection	Mixed	Regulator plus individual complaint	Regulator capacity roughly fixed	Complaint volume swamps capacity; triage becomes the policy
Employment and classification law	Ex post	Individual or representative claims	Claims linear; classification disputes slow	Practice becomes entrenched before adjudication concludes
Sector licensing	Ex ante	Entry condition	Binary and effective where it applies	Frequently avoided by characterising the service as outside the sector
Competition law	Ex post	Agency, on a long cycle	Investigations take years	Market structure is settled before a remedy issues
Securities and disclosure	Ex ante for public issuers only	Regulator plus private action	Does not attach while private	Firms now reach systemic scale without ever triggering it
Contract with counterparties	Ex ante	Private enforcement	Follows bargaining power	Weakest parties obtain the weakest terms
Criminal law	Ex post	Prosecution	Very high threshold	Reaches only deliberate wrongdoing, not structural harm

It is worth special mention of the competition row, which is the most used and the most mismated time instrument. The idea of below-cost expansion that is financed by capital markets is uncomfortable with the elements of predation and recoupment that are central to the classical doctrine: is the price below some appropriate measure of cost and can losses be recouped (Areeda & Turner, 1975)? Where the strategy is financed by investors willing to suffer short-term losses in exchange for long-term gains, both limbs are open to debate, and it has been extensively explored in recent literature (Khan, 2017). Of course, whatever is decided on the merits, the timing is independent of the merits: an investigation is a remedy for a counterfactual that no longer exists once the market structure has been established.

Three observations follow. The first is a time mismatch. Instruments that operate ex post demand a completed harm, an identified claimant and a concluded process, and the time between the harm and the remedy, is typically much longer than the time between the conduct and the hypergrowth firm's change of shape. When an employment classification is adjudicated, it is the industry standard practice for years and the remedy is for a firm that is no longer in existence in the form that adopted it. This is how strategies of conscious legal confrontation work—even when confrontation is not a conscious strategy (Pollman & Barry, 2017).

The second is an architecture mismatch. The process capacity to hear the claimants is not scaled with claimants or harm, but with individually enforced remedies. A regime that is effective on individual claims will therefore be the least effective where there is the greatest harm, and that's exactly the opposite of what it's supposed to do. Collective mechanisms do exist, and these are exactly the answer to that problem, but they are not always available, and are often limited by the kinds of harms where they are most needed.

Thirdly, there is a threshold mismatch. The one most important ex ante regime, securities disclosure, attaches on public listing, and firms tend to grow to the scale, systemic significance, and breadth of investors that the regime was designed for while remaining private and therefore outside the regime. The regulatory boundary follows a financing decision, and not a fact of economics, yet this is a defensible design merely in the world where the two go together.

7. The Objection, and What Would Actually Constrain

7.1 Does this prove too much?

The obvious objection is that the argument, if it is accepted, would condemn the commercial world in general, because every commercial business puts some risk on those who did not give their consent for that, and every expanding business has more new employees than it would like to have. The objection deserves a straight answer, not an hedges.

The answer is, that it is not about any of the elements, but about a conjunction. Risk-taking is not a problem, risk-imposition on non-consenting parties is only justifiable if it has a good reason to be imposed, and the structure of the venture is unique in three ways that make the combination non-incidental. It will sequentially choose the strategy with the largest spread, since it is the one with the best return profile. It matches that selection with a growth rate that predictably undermines the ability of the organisation to detect what it's doing. But it often leaves the firm judgment-proof, in relation to the harm that it can impose, and so the mechanism that usually works to price imposed risks back to the imposer does not. A company that's growing slowly, is self financed, and has assets that could be open to judgment does not have any of these features. The argument is thus more limited than the objection posits and doesn't need to involve any objection to growth or risk as such.

A second argument is that such a regulation would stifle experimentation – and experimentation is the basis of innovation – and this one carries some weight. The experimental account of entrepreneurship is that value is just about cheap and fast funding of projects with high-risk of failure (Ewens et al., 2018; Kerr et al., 2014). The answer is that this is a reason to tolerate failure (as those who carry the cost are the ones who should) but not to tolerate imposition (as those who bear the cost would not want to be imposed on). And it is an operational one, because a failed experiment that costs its backers is a working experiment, one that costs its users or the public is a disguised transfer.

7.2 Responses that address the mechanisms

Table 6 sets out responses matched to the three mechanisms. Two design principles run through them.

Table 6. Responses matched to the identified mechanisms

Response	Mechanism addressed	How it operates	Principal objection	Assessment
Scale-triggered disclosure	Legal threshold mismatch	Obligations attach at economic scale rather than at listing	Burden on private firms; confidentiality	Strong; addresses the clearest structural gap
Mandatory reserves or insurance for foreseeable harm	Judgment-proofing	Restores the deterrent function of liability	Capital cost falls hardest on the smallest	Strong if scaled; weak if flat
Growth-rate-linked control obligations	Control dilution	Control investment required as a function of headcount growth	Novel and hard to specify	Promising; the arithmetic is observable and auditable
Founder and investor secondary-sale disclosure	Variance mandate	Reveals when insiders have de-risked while others have not	Privacy of individual holdings	Moderate; changes incentives at the margin
Collective redress availability	Enforcement architecture	Aggregates dispersed low-value harms into viable claims	Litigation industry concerns	Strong; directly addresses the scaling mismatch
Board-level scale-risk review	All three	Independent review triggered by growth thresholds	Boards in these firms are not independent	Weak alone; depends on prior governance reform
Ex ante approval for defined domains	Harm scaling	Licensing where harms are irreversible or systemic	Slows entry; incumbents benefit	Justified only where irreversibility is genuine

The first is that thresholds should be based on economic significance, not legal form. All the mismatches found in Section 6 are for obligations keyed to a status (e.g. list, licence) that are linked to a fact (e.g. number of people relying on the firm). Secondly, the arithmetic of control dilution is observed, that is to say, it is a very governable object. An obligation for a company to prove its conduct investment in relation to its headcount growth rate is auditable in a manner which is not the case for most conduct obligations.

There is one more answer to the questions in Table 6 that runs counter to the others and deserves to be mentioned. While the prospect of expanding corporate liability may be seen as a way to enhance outcomes, it does not guarantee it, and liability rules can backfire and actually hurt the monitoring they hope to foster if they are not accompanied by a reward for detecting and reporting on violations of the rules (Arlen, 1994). This is relevant here specifically as a failure of detection inside as opposed to a failure of intent. It suggests that any requirement to limit control investment should be matched by a corresponding requirement to provide actionable mitigation for enterprises that uncover and report their own failures and that an investment incentive for an apparent clean investment record will exacerbate the control investment problem, not resolve it. The principle of "who can afford to pay the least" in the realm of accident law suggests the same (Calabresi 1970).

Some observations on framing stakeholders are in order. The typical corporate-governance answer to such worries is to expand the scope of directors' duties, which the critical literature on stakeholder governance contends does little to actually increase accountability or reduce discretion for directors.The usual corporate-governance answer to such concerns is to expand the scope of directors' duties, which the critical literature on stakeholder governance

argues does little in practice to increase accountability or decrease discretion for directors. This is for the reason that the answers in Table 6 are formulaic and not exhortative because they alter what might be disclosed, what might be reserved, and what might be demonstrated, not what might be thought.

8. Propositions and Analytical Protocol

Table 7 states seven propositions with testing strategies. Several are testable from data that registries, filings, and employment records already contain.

Table 7. Propositions and testing strategy

Proposition	Statement	Testing strategy	Parameter of interest
P1	Firm-level risk-taking increases with the concentration of the lead investor's expected returns, controlling for sector and stage.	Match portfolio outcomes to observable risk indicators	Coefficient [b1]
P2	The share of staff below one year of tenure predicts control failures better than headcount, firm age, or sector.	Panel of incidents on tenure composition	Contrast [b2t], [b2h]
P3	Incident rates rise disproportionately in the periods immediately following step changes in headcount growth.	Event study around growth accelerations	Effect [b3]
P4	Founder secondary sales predict subsequent increases in firm-level risk-taking, because they reduce founder exposure to the downside.	Secondary transaction records matched to risk indicators	Coefficient [b4]
P5	Time from conduct to legal remedy exceeds the firm's doubling time in the large majority of contested cases.	Case timelines matched to firm growth records	Ratio [r5]
P6	Potential liability exceeds distributable assets for a substantial share of firms at the point at which harms occur.	Balance-sheet analysis against claimed damages	Coverage ratio [c6]
P7	Norm-displacement harms are not detected by internal metrics before the displacement is complete.	Retrospective analysis of internal reporting against external timelines	Detection lag [t7]

There are four cautions for measuring. However, the most important threat is all throughout: survivorship is the main issue in most datasets, and its absence biases all estimated relationships between growth and outcome in a favorable direction; this is the specific pathology revealed in the learning literature, and it is true of research as well as practice (Denrell, 2003). Incident data are also the result of internal reporting systems whose sensitivity depends on the control environment being measured; the less the control environment is robust, the less sensitive the internal reporting system will be, which biases P2 and P3 toward zero for a firm with a weaker control environment. There is an incomplete and non-random availability of secondary transaction records. Then, there is no direct measure of risk-taking and so operationalisation of risk-taking is a combination of observable proxies whose validity needs to be argued explicitly rather than assumed.

Minnesota Journal of Business Law and Entrepreneurship

Volume 2026, No. 3

ISSN: 1540-3270

9. Discussion

9.1 Implications

The main takeaway is that the pathologies of scale are more spread across the landscape than is generally recognized, and that they are extended to capital. If the variance mandate is playing out as described, then it's not just that the investor recommending the acceleration is taking the risk on himself; it's that he's creating the risk, and the allocation of legal and reputational liability, which's almost entirely on the firm and the founders, doesn't align with the causal structure. While this does not mean that investors will be held liable for the actions of their portfolio companies, it does mean that disclosure of an investor's influence over the actions is a reasonable governance objective. A second implication is that of the founders themselves, which is more actionable. In Section 3, all of the arithmetic is known in advance to all of the founders. A growth rate is selected and once selected, selection of that rate results in a tenure composition of the organisation being formed and thus the feasibility of informal control. No matter if they have decided to build a control system or not, if they intend to grow at a rate that it will leave most of the firm below one year in the job, they have essentially chosen to use a formal control system. This reframing of the decision problem makes it a resourcing question and gives a calculative answer. A third relates to the research literature. Prescriptive work on scaling is focused overwhelmingly on survivors, and its core propositions are therefore confuted, albeit understood in principle and often forgotten in practice. The best empirical contribution that is available in this area is systematic study of firms that have scaled aggressively, but have failed which involves gathering data on the population, not the sample that is also visible.

9.2 Limitations

There are three restrictions that should be clearly expressed. The propositions are not tested: This is a conceptual and normative paper, which is to say that the figures are purely manifesting the structure of an argument and do not attempt to estimate any magnitude. The control dilution arithmetic in Figure 2 is correct under the assumptions it makes, which are contained in the remainder of this report, and which are both simplifications: the amount of attrition, even in fast-growing firms, is not negligible; and the precise amounts depend on parameters not estimated here and that are not constant. Secondly, the analysis relies on the common-law and European context, and the enforcement system varies significantly elsewhere, especially on the basis of the availability of collective redress, which the analysis deems as being most significant. Third, the moral stance used, based on the difference between risk taken and risk imposed, is just one of several morally defensible ones, and readers who embrace an aggregate welfare moral stance will come to different decisions regarding situations where large diffuse benefits are accompanied by concentrated imposed risk.

10. Conclusion

The problems that characterise Hypergrowth are its failures are not based on the nature of the people involved. The capital structure encourages variance and makes no judgments about moderate success and failure; it will lead firms toward risk that neither the founders nor the affected parties will choose. The actual growth rate in itself establishes the mix of tenure in the organisation and when the growth rate exceeds a calculable threshold, the informal control system is destroyed without any decision to abandon the system. The fact that people are exposed alters the moral nature of an act from risk they might have taken to risk they don't, at a stage in a firm's life when its ability to compensate for harm is the lowest relative to the harm that it can cause.

Each of the three mismatches of the legal boundary set is particularly problematic, with the dominant instruments being ex post, individually enforced, and keyed to the legal status of the party rather than the economic significance of the event, and all causing significant harm as growth speeds up. The answers that follow are neither glamorous nor of anything particularly individual, and are mostly structural: obligations that bind at scale instead of at listing; reserves that regain the force of deterrence that comes with liability; control over investment that attaches to the observable arithmetic of growth; and, collective mechanisms to link up what individual claims

cannot. None of these are saying that growth is a bad thing and none of these are saying that experimentation should not be allowed, which would be necessary to eliminate if the venture model is to be eliminated. What they would prevent is the special conjunction, as described in this paper, in which the party that chooses the risk would be different from the party that bears the risk and the party that could have detected the risk. Here are seven propositions, and the analytical protocol offered, which are designed to render that claim testable and hence refutable, not just asserted.

References

[1] Aldrich, H. E., & Ruef, M. (2006). *Organizations evolving* (2nd ed.). Sage.

[2] Areeda, P., & Turner, D. F. (1975). Predatory pricing and related practices under Section 2 of the Sherman Act. *Harvard Law Review, 88*(4), 697-733.

[3] Arlen, J. (1994). The potentially perverse effects of corporate criminal liability. *Journal of Legal Studies, 23*(2), 833-867.

[4] Bebchuk, L. A., & Tallarita, R. (2020). The illusory promise of stakeholder governance. *Cornell Law Review, 106*(1), 91-178.

[5] Calabresi, G. (1970). *The costs of accidents: A legal and economic analysis*. Yale University Press.

[6] Cochrane, J. H. (2005). The risk and return of venture capital. *Journal of Financial Economics, 75*(1), 3-52.

[7] Coffee, J. C. (1981). "No soul to damn: No body to kick": An unscandalized inquiry into the problem of corporate punishment. *Michigan Law Review, 79*(3), 386-459.

[8] Denrell, J. (2003). Vicarious learning, undersampling of failure, and the myths of management. *Organization Science, 14*(3), 227-243.

[9] Ewens, M., Nanda, R., & Rhodes-Kropf, M. (2018). Cost of experimentation and the evolution of venture capital. *Journal of Financial Economics, 128*(3), 422-442.

[10] Gompers, P. A., Gornall, W., Kaplan, S. N., & Strebulaev, I. A. (2020). How do venture capitalists make decisions? *Journal of Financial Economics, 135*(1), 169-190.

[11] Greiner, L. E. (1972). Evolution and revolution as organizations grow. *Harvard Business Review, 50*(4), 37-46.

[12] Hall, R. E., & Woodward, S. E. (2010). The burden of the nondiversifiable risk of entrepreneurship. *American Economic Review, 100*(3), 1163-1194.

[13] Hansmann, H., & Kraakman, R. (1991). Toward unlimited shareholder liability for corporate torts. *Yale Law Journal, 100*(7), 1879-1934.

[14] Hansson, S. O. (2003). Ethical criteria of risk acceptance. *Erkenntnis, 59*(3), 291-309.

[15] Hansson, S. O. (2013). *The ethics of risk: Ethical analysis in an uncertain world*. Palgrave Macmillan.

[16] Hoffman, R., & Yeh, C. (2018). *Blitzscaling: The lightning-fast path to building massively valuable companies*. Currency.

[17] Jonas, H. (1984). *The imperative of responsibility: In search of an ethics for the technological age*. University of Chicago Press.

[18] Kerr, W. R., Nanda, R., & Rhodes-Kropf, M. (2014). Entrepreneurship as experimentation. *Journal of Economic Perspectives, 28*(3), 25-48.

[19] Khan, L. M. (2017). Amazon's antitrust paradox. *Yale Law Journal, 126*(3), 710-805.

[20] Korteweg, A., & Sorensen, M. (2010). Risk and return characteristics of venture capital-backed entrepreneurial companies. *Review of Financial Studies, 23*(10), 3738-3772.

[21] Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal, 14*(S2), 95-112.

[22] LoPucki, L. M. (1996). The death of liability. *Yale Law Journal, 106*(1), 1-92.

[23] March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science, 2*(1), 71-87.

[24] Metrick, A., & Yasuda, A. (2010). The economics of private equity funds. *Review of Financial Studies, 23*(6), 2303-2341.

[25] Nanda, R., & Rhodes-Kropf, M. (2013). Investment cycles and startup innovation. *Journal of Financial Economics, 110*(2), 403-418.

[26] Penrose, E. T. (1959). *The theory of the growth of the firm*. Oxford University Press.

[27] Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. Basic Books.

[28] Pollman, E., & Barry, J. M. (2017). Regulatory entrepreneurship. *Southern California Law Review, 90*(3), 383-448.

[29] Sahlman, W. A. (1990). The structure and governance of venture-capital organizations. *Journal of Financial Economics, 27*(2), 473-521.

[30] Shavell, S. (1986). The judgment proof problem. *International Review of Law and Economics, 6*(1), 45-58.

[31] Sunstein, C. R. (2005). *Laws of fear: Beyond the precautionary principle*. Cambridge University Press.

[32] Thompson, D. F. (1980). Moral responsibility of public officials: The problem of many hands. *American Political Science Review, 74*(4), 905-916.

[33] Turner, B. A., & Pidgeon, N. F. (1997). *Man-made disasters* (2nd ed.). Butterworth-Heinemann.

[34] Vaughan, D. (1996). *The Challenger launch decision: Risky technology, culture, and deviance at NASA*. University of Chicago Press.