Which is king: cash or cash burn?
Investment-liquidity trade-off in VC-backed startups

Peter K. Pham* Jason Zein† Ximu Zeng‡

August 14, 2026

Abstract

Early-stage startups face a fundamental trade-off: investing aggressively to accelerate growth depletes liquidity needed to hedge against refinancing risk. We study this trade-off using hiring intensity – the rate of additional labor-cost commitment to the amount of capital raised in a financing round – as an inverse proxy for post-financing liquidity buffer. Hiring intensity varies substantially across startups and strongly predicts every major dimension of startup success, including next-round continuation, valuation growth, successful exits, product-market success, and long-term employment expansion. A double machine learning analysis shows that these relationships remain robust after flexibly controlling for a rich set of observable startup, founder, investor, and employee characteristics. Our findings are consistent with a signalling model in which high-quality startups credibly reveal their type by committing to operate with minimal liquidity buffers.

Keywords: Venture Capital, Private Firm Valuations, Fundraising, Startup Employment Growth

JEL Codes: G24, G32

*University of Sydney Business School. Email: peter.pham@sydney.edu.au
†UNSW Business School, University of New South Wales. Email: jason.zein@unsw.edu.au
‡University of Sydney Business School. Email: ximu.zeng@sydney.edu.au

1 Introduction

A fundamental tension defines venture capital-backed startups: the choice between burning cash to accelerate growth and conserving cash to extend runway. Early-stage startups generate little or no internal cash flow, so a financing round is effectively their entire stock of liquidity until the next one. While a high cash burn allows capital to be deployed quickly to achieve scale sooner, it shortens the cash runway, tightens the cash reserve, and exposes the startup's next funding round (and survival) to external financing market conditions. This investment-liquidity trade-off is a central challenge in entrepreneurial finance, yet we lack systematic evidence on when cash burn creates versus destroys value.

Venture capital practice often rewards aggressive cash burn and rapid scaling. Startup outcomes are highly skewed ex post and unpredictable ex ante: many portfolio companies fail, a small number of exceptional successes generate a disproportionate share of venture returns, and even experienced investors cannot tell the two apart at the time they invest (Hall and Woodward, 2010; Kerr et al., 2014).[1] Venture capitalists can therefore tolerate substantial losses on individual investments as long as a small number of winners generate outsized returns. This payoff structure has helped sustain a "growth-at-all-costs"[2] mentality among VC-backed startups, under which rapid scaling is viewed as a way to accelerate growth, increase valuation, and improve future funding prospects.

However, conventional wisdom increasingly emphasizes the risks of aggressive cash burn.

[1] Kerr et al. (2014) report that about 55 percent of US venture-backed startups financed between 1985 and 2009 were terminated at a loss, while the 6 percent returning more than five times capital generated roughly half of gross returns. Hall and Woodward (2010) report about 50 percent zero-value exits and a 13 percent IPO rate.

[2] "Growth-at-all-costs" refers to the startup strategy of aggressively pursuing rapid expansion in revenue, users, and market share, typically running significant operational losses funded by venture capital, with the belief that achieving scale and market dominance will eventually lead to profitability.

The data supports this concern: CB Insights 2021 reports that 38% of startups fail from running out of money (compared to the 29% reported in 2019), the first leading cause, followed by lack of product-market fit at 35%.[3] This conservative stance gained momentum since the 2022 market downturn. Sequoia Capital, the leading venture capital firm, warned that "cut costs, preserve cash," prompting investors to deprioritize growth in favor of efficiency, a marked departure from the "growth-at-all-costs" mentality of previous years.[4] This shift is evident in recent funding data: Silicon Valley Bank states in 2025 that "among companies that raised capital in 2024, the typical Series B company only increased its burn 8% YoY. This means that companies are growing, but they aren't growing their burn. Companies that are raising are increasingly efficient. This is vastly different from companies raising in 2021 and 2022 that rapidly grew burn YoY in an environment where capital was easier to come by."[5]

This debate is exemplified by prominent startups that pursued similarly aggressive cash-burn strategies but experienced sharply different outcomes. Uber accumulated more than $10 billion in net losses between 2016 and 2018, yet reached an $82 billion valuation at its 2019 IPO. In contrast, WeWork pursued comparably aggressive expansion, reporting $1.9 billion in losses against $1.54 billion in revenue in its 2019 IPO filing, but subsequently saw its valuation collapse from $47 billion to below $8 billion and ultimately entered Chapter 11

[3]CB Insights 2021 THE TOP 12 REASONS STARTUPS FAIL. Available at: https://www.cbinsights.com/research/report/startup-failure-reasons-top/.

[4]With approximately $85 billion in assets under management as of 2022, Sequoia has invested in thousands of startups, including early-stage investments in Google, Apple, Nvidia, Uber, and Zoom. In Sequoia's presentation to founders on May 16, 2022, Sequoia stated that "With macro uncertainty around inflation, interest rates and war, investors are relooking for companies that can produce near-term certainty. Capital is getting more expensive . . . leading to investors de-prioritising and paying up less for growth . . . The focus on near-term momentum is often shifting toward companies who can demonstrate current profitability." Available at: https://www.sequoiacap.com/article/adapting-to-endure/.

[5]Silicon Valley Bank State of the Markets Report H1 2025. Available at: https://www.svb.com/trends-insights/reports/state-of-the-markets-report/.

bankruptcy restructuring in 2023. Aggressive capital commitments can therefore have very different consequences for financial fragility, depending on the firm's capacity to refinance future liquidity needs. This makes aggressive burn inherently difficult to interpret ex ante.

Academic research provides mixed evidence on this trade-off. A long literature has extensively documented venture capital's positive impact on startup growth. In particular, employment (Puri and Zarutskie, 2010; Davila et al., 2003; Hellmann and Puri, 2002), sales (Bertoni et al., 2011; Hellmann and Puri, 2000), innovation (Lerner and Nanda, 2020; Kortum and Lerner, 2000), and productivity (Croce et al., 2013; Chemmanur et al., 2011; Pellón and Alemany, 2006). On average, these studies have highlighted that VC-financed startups grow faster, patent more, have higher productivity and likelihood to go public than non-VC-backed. However, recent studies suggest that prioritizing growth over efficiency can increase failure rates by limiting experimentation and premature commitment to unvalidated models (Lee and Kim, 2024; Frésard et al., 2023). Yet we lack systematic understanding of how VC-backed startups strategically navigate the investment-liquidity trade-off to optimize both growth and survival prospects.

This paper examines this investment-liquidity trade-off by focusing on how aggressively startups convert raised capital into additional recurring labor-cost commitments. Labor is both a core productive input and a major operating expense for startups, making changes in salary spending an economically meaningful proxy for capital-deployment intensity. This proxy is particularly well suited to our sample, in which approximately 72% of startups operate in technology sectors and primarily employ high-skilled workers.

Our analysis documents 17,583 first-round financing observations for U.S. startups, combining granular monthly employment data from Revelio with comprehensive financing char-

acteristics from PitchBook. We document a substantial "hiring surge" around each financing event: monthly salary spending grows by 159.1% on average from six months before to six months after each financing event, representing 80.6% of the net increase until the startup raises capital again in the next round. Employee count grows by 146.5% over the same window, representing 71.9% of the net increase until the next round. Exploiting this pattern, we construct a measure of the investment-liquidity trade-off called Hiring Burn Rate (HBR), defined as the ratio of the net change in monthly salary spending over the financing window to the amount of capital raised in the round. This measure captures how aggressively a startup invests, at the same time burning through available cash. In other words, the measure reflects how a startup trades up liquidity for growth in a limited and conditional funding environment.

To motivate our empirical analysis, we develop a simple conceptual framework in which a high Hiring Burn Rate serves as a signal for venture quality. We build on the firm liquidity-demand models of Holmström and Tirole (1998, 2000) by introducing asymmetric information. In our model, a high-quality startup, when seeking refinancing, can pledge more of its continuation value to investors, allowing it to accommodate a larger liquidity shock. It can afford to operate with a thin liquidity buffer whereas a low-quality startup operating in the same condition would more often be forced into liquidation. In the resulting separating equilibrium, high-quality startups can thus deliberately commit to a more aggressive investment schedule while retaining less liquidity per dollar deployed, thereby accepting greater liquidation risk in order to credibly reveal their type. A small liquidity buffer is not merely a mechanical reflection of a high-quality startup's better investment opportunities, but a strategic choice that low-quality firms cannot profitably imitate.

Consistent with the model's central prediction that a higher burn rate reveals quality, we find that the Hiring Burn Rate predicts a broad set of financing and real outcomes, with magnitudes that are economically meaningful and stable across specifications. In the financing market, a one-standard-deviation increase in HBR raises the probability of securing a follow-on round by 5.11 percentage points. High-burn firms also raise on better terms, with a one-standard-deviation increase in HBR associated with 12.9% higher valuation growth between rounds, or about $1.99 million in additional value at the average pre-money valuation of $15.41 million. A Cox proportional-hazards specification shows that the effect operates on the timing of financing as well as its occurrence: higher HBR raises the hazard of securing the next round by 15.3% to 17.3%, so high-burn firms refinance both more often and more quickly.

These relationships are not confined to the next financing event, nor do they reflect a one-time burst of hiring. Higher HBR predicts lower long-run failure, with a one-standard-deviation increase associated with a 2.87% to 3.44% reduction in failure probability, and it predicts better real performance conditional on survival: product activity is 1.8% and 3.1% higher over the three- and five-year horizons, respectively, while patent activity is 3.9% and 5.3% higher. The hiring itself persists: firms with higher HBR sustain 4.3% and 3.6% greater employment growth over the same three- and five-year horizons. These findings are robust to defining HBR over the narrower $[M-3, M+3]$ window. These results match the model's prediction that a thin liquidity policy reflects a durable type rather than a transitory effect.

A natural concern is that high-quality startups hire aggressively because they have better prospects, so that HBR proxies for observable quality rather than conveying additional information. We first address this concern by conditioning on startup, investor, founder,

employee, and market controls, together with fixed effects, in the baseline specifications. The fixed effects allow us to compare startups financed in the same year and state and operating within the same industry-by-pre-round-employment-size-by-fund-size cell, while absorbing common financing conditions, local market conditions, and unobservable differences in sector, scale, and fund size.

To further examine the quality-signaling and liquidity-risk channels, we classify financing continuation and firm failure according to whether a startup records a new product launch or patent filing between the focal round and the next financing or failure event, respectively. Higher-burn startups are more likely to raise again when they have no such milestone, and no more likely when they do. The burn rate therefore carries information precisely when investors have no hard evidence of quality to price, and adds nothing once that evidence exists. Failure shows the cost side. Higher HBR lowers the probability of failure overall, but among startups that fail, those with higher HBR are more likely to have recorded a product launch or patent filing than those with lower HBR. Aggressive burn does not make failure more likely; it changes what failure looks like, shifting it away from startups that never had a viable product and toward startups that had one and still ran out of money.

The workforce decompositions show where this salary-based signal is concentrated. Salary-spending HBR remains associated with financing, survival, and real outcomes after partialling out employee-count HBR, controls, and fixed effects. Within salary spending, the strongest follow-on-financing effects are concentrated in product-related and high-wage labor commitments: product HBR has the largest function-level association (3.9% for the Next Round), management HBR has the largest seniority-level association for the Next Round (2.7%) and senior HBR has the largest seniority-level association for valuation growth (7.8%),

and the top wage quartile has the largest wage-level association (3.6% for the Next Round). These results suggest that the HBR signal is concentrated in product-oriented and high-wage talent, consistent with the model's prediction that HBR reflects credible commitments to product development and skilled-talent deployment that are difficult for low-quality startups to imitate.

The preceding tests condition on rich linear controls, but observable quality may still enter in high-dimensional and nonlinear ways. We therefore turn to double machine learning, which is designed for this problem. It uses flexible machine learning to predict both the burn rate and each outcome from the full set of 278 startup, deal, founder, and employee characteristics, removes those predictions to isolate the parts of HBR and of outcomes that observables cannot explain, and then estimates the association between the two residuals, with cross-fitting to keep the flexibility of the prediction step from biasing the estimate. We find that the part of HBR that is orthogonal to all 278 observables continues to predict outcomes across every dimension and horizon we examine, a pattern consistent with private information about quality that insiders act on but the econometrician cannot measure.

These analyses do not make HBR an exogenous treatment. A central empirical challenge is that HBR is an endogenous allocation choice. The baseline specifications and DML analysis discipline the role of observed startup, founder, investor, and employee characteristics, including high-dimensional and nonlinear observables, but they cannot rule out unobserved investment opportunities, managerial expectations, or private information that jointly shape HBR and subsequent outcomes. We therefore interpret the evidence as showing that HBR is a robust predictor of startup performance, consistent with the quality-dominance case in the model, rather than as estimating the causal effect of increasing burn.

This paper contributes to several literatures. First, we provide novel evidence on the investment-liquidity trade-off in entrepreneurial finance, showing how aggressive cash burn can create value through credible signaling. Our use of HBR distinguishes our research by directly measuring cash burn intensity as a proxy for liquidity risk, a fundamental tension in venture capital financing. Our findings relate to prominent empirical work such as Puri and Zarutskie (2010) which shows that venture capital significantly accelerates startup growth in employment and sales at every investment stage, with VC-financed firms exhibiting lower cumulative failure rates, particularly in their initial years after receiving VC funding. However, recent work has questioned the universal benefits of rapid scaling. Lee and Kim (2024) find that scaling early can suppress startup performance by limiting learning through experimentation and committing too soon to unvalidated business models, leading to higher failure likelihood without compensating benefits in exit outcomes. We offer practical insights for this ongoing debate between growth and efficiency in startup strategy, particularly relevant given tightening and volatile capital markets. Second, our decomposition across functions and seniority levels elucidates strategic human capital allocation and scaling priorities, while capturing organizational professionalization dynamics (Hellmann and Puri (2002)) and innovation-to-market transitions in startup development. Furthermore, we extend research on venture capital staging and monitoring by demonstrating how employment scaling serves as both a strategic tool and an information revelation mechanism.

Our study uses 17,583 first financing events of U.S.-based startups linked to monthly employment records over extended observation windows. This granularity contrasts with existing work that relies on annual employment snapshots or short event windows around financing. The sample is concentrated in technology sectors, which account for approximately

72% of observations; SaaS and AI startups account for 32.8% and 19.4%, respectively. The data also capture the human capital behind these VC-backed, high-growth innovation startups: 36,873 founders and 1,474,501 high-skilled startup employees observed from one year before to five years after the first financing round.

Empirically, we construct rich, multidimensional startup characterizations that allow us to condition on a broad set of observables. Beyond traditional econometric approaches, we implement the Double Machine Learning (DML) framework of Chernozhukov et al. (2018) to capture complex non-linear relationships in startup hiring dynamics. DML demonstrates robust out-of-sample predictability and shows that HBR retains predictive content after high-dimensional adjustment.

The remainder of the paper is organized as follows. Section II presents the conceptual framework, Section III describes the data and empirical design, Section IV reports the results, Section V presents the double machine learning analysis, and Section VI concludes.

2 Conceptual Framework

Our starting point is the variable-scale liquidity-demand model of Holmström and Tirole (2000). In that model, an entrepreneur with limited inside wealth finances a project of freely chosen scale, but her effort is not contractible, so she must retain an incentive stake and only part of the project's continuation value, the pledgeable income, can be promised to investors. An interim liquidity shock then arrives that must be covered for the project to continue, met from liquidity set aside in advance and from refinancing raised on the spot up to pledgeable income. The firm chooses its scale and its liquidity buffer jointly: a larger

buffer raises the continuation cutoff, the largest shock the firm survives, but competes with scale for the same pledgeable income, so the firm does not insure fully.

We extend this benchmark by incorporating asymmetric information. Firms come in two privately known types that differ only in pledgeable income, so a higher-quality firm can refinance a larger shock; we call this difference the pledgeability gap. And the market learns a firm's type only imperfectly before refinancing, which we summarize by a single transparency parameter. Together these ingredients turn the liquidity buffer into a signal. A thin buffer is affordable for a good firm and costly for a bad one, because if the market learns the type the good firm's survival threshold exceeds the bad firm's by exactly the pledgeability gap. In the resulting separating equilibrium the good type raises a larger multiple of her assets and retains less liquidity per unit of investment, accepting higher liquidation risk as the price of revealing quality. The signal operates only when the market has some chance of learning the type.

We summarize the setting and key results of our framework below and provide detailed proofs in the Appendix A.

2.1 Setup

There are three dates $t = 0, 1, 2$, universal risk neutrality, a zero interest rate, limited liability, and competitive investors.

At date 0 an entrepreneur privately knows her type $\theta \in \{G, B\}$, good or bad. The market prior is $\alpha = \Pr(\theta = G)$, the population share of good types. She owns assets $A > 0$, her inside capital and the only internal source of funding, and operates a constant-returns

technology in which any investment scale $I \geq 0$ may be implemented at unit cost.

At date 1 a liquidity shock ρI is realized, where $\rho \sim F$ with continuous density $f > 0$ on $[0, \bar{\rho}]$. The quantity ρ is the firm's date-1 funding need per unit of investment, for example a cost overrun or a refinancing requirement, and meeting it is necessary for continuation; an abandoned project yields nothing beyond cash on hand. For simplicity, we assume that the firm has no interim revenue, as typical for early stage startups.We further assume that F is strictly log-concave with density $f > 0$ on $[0, \bar{\rho}]$ and that $\int_0^{\bar{\rho}} F(\rho)\, d\rho > 1$.

Before refinancing, a public signal reveals the type with probability ϕ, independently of ρ, and with probability $1-\phi$ the signal is uninformative, so ϕ measures interim transparency. We assume that $\phi \in (0, 1]$, that the signal is independent of ρ, and that it is realized before date-1 refinancing.

At date 2, conditional on continuation, the entrepreneur chooses unobservable effort. Behaving yields success probability p_H for type G or q_H for type B and no private benefit; misbehaving yields p_L or q_L and a private benefit BI, the entrepreneur's nonpecuniary gain from shirking, proportional to scale. Spreads are equal, $p_H - p_L = q_H - q_L = \Delta p$, so the continuation incentive constraint does not depend on the type. With a success-contingent stake R_b and per-unit stake $b \equiv R_b/I$ (we reserve R_b for the total and use b throughout), the entrepreneur behaves if and only if $(\Delta p)\, b \geq B$. Success pays RI, where R is output per unit of investment, and failure pays 0. We assume that $R > B/\Delta p$, so that a successful unit produces more than the minimum stake the entrepreneur must retain and a strictly positive residual remains to pledge to investors.

Per-unit pledgeable and total continuation values are then

$$\rho_0^G = p_H\left(R - \tfrac{B}{\Delta p}\right) \;>\; \rho_0^B = q_H\left(R - \tfrac{B}{\Delta p}\right) \;>\; 0, \qquad \rho_1^G = p_H R \;>\; \rho_1^B = q_H R. \tag{1}$$

Here ρ_1^θ is what a continued unit is worth, success odds times output, and ρ_0^θ is the part of that value that can be credibly promised to outside financiers once the entrepreneur's minimum stake $B/\Delta p$ is set aside. The difference $\rho_1^\theta - \rho_0^\theta = p_\theta B/\Delta p$ is a non-transferable agency rent. All three objects are denominated in the same units as the shock, namely date-1 resources per unit of investment, and they are best read as landmarks on the shock line. The pledgeability gap $\Delta\rho_0 \equiv \rho_0^G - \rho_0^B = (p_H - q_H)(R - B/\Delta p) > 0$ is the only difference between the two types.

At date 1 a firm meets the shock from a prearranged liquidity facility, either cash hoarded at date 0 or an irrevocable credit line, and from spot refinancing up to perceived pledgeable income. The liquidity policy is the facility per unit of investment, $\ell \geq 0$, and a firm whose contract is attributed to type G continues if and only if $\rho \leq \kappa$, where

$$\kappa \;=\; \ell + \rho_0^G. \tag{2}$$

The floor $\ell \geq 0$ is a renegotiation-proofness requirement rather than an assumption. By the soft-budget logic of Holmström and Tirole (2000), the spot market always rescues a firm whose shock lies below perceived pledgeable income, so no contract can commit to a cutoff below $\rho_0^{\text{perceived}}$, and zero retained liquidity is the maximal credible exposure. Writing $S_\theta(\kappa)$ for the survival probability of a type-θ holder of a G-attributed contract, where the subscript

tracks the holder's true type,

$$S_G(\kappa) = F(\kappa), \qquad S_B(\kappa) \;=\; \phi\, F(\kappa - \Delta\rho_0) + (1-\phi)\, F(\kappa). \tag{3}$$

A true good type has capacity κ and survives with probability $F(\kappa)$. A bad type imitating the contract keeps the facility in all events, but upon revelation, which occurs with probability ϕ, the spot market funds her only to her true pledgeable income, so her capacity falls to $\ell + \rho_0^B = \kappa - \Delta\rho_0$. The difference $S_G(\kappa) - S_B(\kappa) = \phi\,[F(\kappa) - F(\kappa - \Delta\rho_0)]$ is strictly positive if and only if $\phi > 0$, and it is the product of transparency ϕ and the pledgeability gap $\Delta\rho_0$.

2.2 The symmetric-information benchmark

To isolate what asymmetric information adds, we first solve the model when types are public. This also fixes the closed forms used throughout. Define the unit cost of effective investment

$$c(\kappa) \;\equiv\; \frac{1 + \int_0^\kappa \rho f(\rho)\, d\rho}{F(\kappa)}, \tag{4}$$

and let V_θ denote the entrepreneur's gross expected receipts and $W_\theta \equiv V_\theta - A$ her net gain, which equals project net present value when investors break even on the true type.

Lemma 1. *On the competitive contract frontier, defined by the binding break-even condition, maximal scale, and the stake at its incentive floor $b = B/\Delta p$, for a type θ priced truthfully at any cutoff κ with $F(\kappa) > 0$, the equity multiplier and values are*

$$I = k_\theta(\kappa) A, \quad k_\theta(\kappa) = \frac{1}{F(\kappa)\,(c(\kappa) - \rho_0^\theta)}, \qquad V_\theta(\kappa) = \frac{\rho_1^\theta - \rho_0^\theta}{c(\kappa) - \rho_0^\theta}\, A, \quad W_\theta(\kappa) = \frac{\rho_1^\theta - c(\kappa)}{c(\kappa) - \rho_0^\theta}\, A.$$

The multiplier k_θ is strictly decreasing in κ on $\{\kappa > \rho_0^\theta\}$, and c is strictly quasi-convex with unique minimizer ρ^* solving $\int_0^{\rho^*} F = 1$ and $c(\rho^*) = \rho^*$.

The benchmark cutoff ρ^* depends only on F. With ρ^* defined, we can state the remaining restrictions. We assume that one of two creditworthiness regimes holds. In the funded cross-section regime, denoted (F) and the main case, $\rho^* < \rho_1^B < \rho_1^G$, so both types are creditworthy. In the exclusion regime, denoted (E) and the limiting case, $\rho_1^B < \rho^* < \rho_1^G$, so only the good type is creditworthy. We also assume $\rho_0^G < \rho^* < \bar{\rho}$, so that the symmetric-information buffer is strictly positive, and we assume $c(\rho_0^G) < \rho_1^G$, so that the good project has positive net value even at zero retained liquidity.

Lemma 2. *With types public, any creditworthy type θ chooses the cutoff $\kappa = \rho^*$, which is independent of ρ_0^θ and ρ_1^θ, with liquidity policy $\ell_\theta^* = \rho^* - \rho_0^\theta$, multiplier $k_\theta^* = 1/[F(\rho^*)(\rho^* - \rho_0^\theta)]$, and net value $W_\theta^* = A(\rho_1^\theta - \rho^*)/(\rho^* - \rho_0^\theta)$. In regime (F) both types are funded, and the good type holds the thinner policy and the larger multiple,*

$$\ell_G^* = \ell_B^* - \Delta\rho_0, \qquad k_G^* > k_B^*,$$

with no excess liquidation risk, since both continue if and only if $\rho \le \rho^*$. In regime (E) the bad type is unfunded and consumes A.

Remark 1. Lemma 2 produces a correlation between quality and thinness even without any signaling. A better firm runs a thinner buffer because its spot capacity ρ_0^θ is larger, so it finances the same survival target ρ^* with less hoarded cash. We refer to this as the mechanical component. It is survival-neutral, since both types continue with the same probability $F(\rho^*)$.

Any difference in survival or downstream outcomes across the burn distribution therefore cannot be mechanical, and is what asymmetric information adds in Section 2.3.

2.3 Signaling through the liquidity buffer

Because frontier contracts carry the same stake and scale for whoever holds them, the comparison that governs imitation reduces to a scale-free ratio,

$$\Omega(\kappa) \;\equiv\; \frac{V_B^{\mathrm{mim}}(\kappa)}{V_G(\kappa)} \;=\; \frac{q_H\, S_B(\kappa)}{p_H\, S_G(\kappa)} \;=\; \frac{q_H}{p_H}\left[1-\phi\Big(1-\tfrac{F(\kappa-\Delta\rho_0)}{F(\kappa)}\Big)\right]. \tag{5}$$

Lemma 3. *Given we assume that F is strictly log-concave and that $\phi > 0$, the ratio Ω is strictly increasing on $[\rho_0^G, \rho^*]$, so reducing the buffer is relatively cheaper for the good type. The mechanism is the interaction of transparency ϕ and the pledgeability gap $\Delta\rho_0$.*

Corollary 1. *If $\phi = 0$ then $\Omega(\kappa) \equiv q_H/p_H$ for all κ, the liquidity margin provides no differential screening, pooling survives the Intuitive Criterion, and any separation is value-burning with no liquidity-specific content. Strictly positive interim information is therefore necessary for the buffer to screen types.*

A separating equilibrium also requires that the good type not present herself as a bad type in order to obtain a fat-buffer contract at bad-type pricing. We rule this out by assuming a no-masquerade condition at the equilibrium cutoff: with $\bar{\Psi} \equiv \frac{p_H}{q_H}[1-\phi+\phi F(\rho_0^G)/F(\rho_0^B)]$,

$$\frac{\rho^* - \rho_0^B}{c(\hat{\kappa}) - \rho_0^G} \;\geq\; 1-\phi+\phi\,\frac{F(\rho_0^G)}{F(\rho_0^B)}.$$

The condition holds when the bad type's franchise value at bad-type pricing is small relative

to the survival advantage a good type retains under such a contract, and it is vacuous in the Exercise 6.10 limit. We present the funded cross-section regime (F), which is the operative case for the data because the sample conditions on funded firms; regime (E) follows as the limit in which the bad type's franchise collapses to autarky, and the two paste continuously at $\rho_1^B = \rho^*$. Imitation is deterred when $\Omega(\kappa)V_G(\kappa) \leq V_B^*$, which is independent of A and equivalent to

$$S_B(\kappa)\, k_G(\kappa) \;\leq\; F(\rho^*)\, k_B(\rho^*). \tag{IC-F}$$

The condition states that the bad type's survival probability times the scale she would command under the good type's contract must not exceed her own survival times her own scale. The left side is strictly increasing on $(\rho_0^G, \rho^*]$, which defines the threshold κ_{IC}^F and the equilibrium cutoff $\hat{\kappa}_F = \min\{\rho^*, \kappa_{IC}^F\}$.

Proposition 1. *In regime (F), given the no-masquerade condition at $\hat{\kappa}_F$ and that (IC-F) is attainable on $(\rho_0^G, \rho^*]$, there is a perfect Bayesian equilibrium on the frontier in which the bad type takes her benchmark contract, with fat buffer ℓ_B^* at bad-type pricing, and the good type takes the distorted thin contract at $\hat{\kappa}_F$ at good-type pricing. On the equilibrium path:*

(i) both types are funded, and the good type earns $V_G(\hat{\kappa}_F) > V_B^*$;

(ii) imitation in both directions is deterred;

(iii) the good type retains less liquidity, $\hat{\ell}_F = \hat{\kappa}_F - \rho_0^G < \ell_B^*$, with the gap

$$\ell_B^* - \hat{\ell}_F \;=\; \underbrace{\Delta\rho_0}_{\textit{mechanical}} \;+\; \underbrace{(\rho^* - \hat{\kappa}_F)}_{\textit{strategic}},$$

and operates at a larger multiple, $k_G(\hat{\kappa}_F) > k_G^* > k_B^*$, so she both raises a larger

multiple of her assets and retains less liquidity per unit of investment;

(iv) the thin firm carries more date-1 liquidation risk, $F(\hat{\kappa}_F) \leq F(\rho^*)$, strictly when $\kappa_{IC}^F < \rho^*$, so the dissipative cost of the distortion falls on path and is concentrated in the high-burn good firms;

(v) the equilibrium is renegotiation-proof for both contracts.

Proposition 2. *Given $\phi > 0$, among the separating equilibria on the frontier only the least-cost one, $\hat{\kappa}_F$, survives the Intuitive Criterion, and pooling fails whenever the good type's least-cost separating value exceeds her pooling value. When $\phi = 0$ this selection fails, by Corollary 1.*

The relevant proofs are provided in Appendix A.

2.4 Empirical predictions

The model implies that, conditional on continuing, a high-burn firm is of higher quality and should therefore perform better on the dimensions that quality affects, such as the probability of raising another round, valuation growth, exit, and longer-run real outcomes. This relationship is conditional on a the presence of sufficient quality difference between startup types . We take this quality-dominance case as the natural one in the startup economy, under which case the model predicts a positive association between the burn rate and the above observable outcomes.

A second implication concerns the source of this predictive content. Because higher-quality firms can pledge more, they could retain less liquidity for purely mechanical reasons, with no signaling at all, but that channel operates through observable differences in quality

and investment opportunity. The model instead predicts that a component of the burn rate forecasts outcomes even after such observables are held fixed, because the thin buffer is a strategic choice rather than a by-product of higher pledgeable value. This is the prediction that our controls and the double machine learning analysis are designed to test, by isolating the part of the burn rate that is orthogonal to a large set of firm, founder, and employee characteristics. The same logic answers why low-quality firms do not simply imitate the policy: in the separating equilibrium a thin buffer is bearable only for high-quality firms, since only they can refinance the liquidity shocks to which it exposes them.

A third implication concerns persistence. The burn rate reflects a durable type rather than a transitory burst of hiring, so its informativeness should not fade with time, and it should continue to predict performance at long horizons, including employment growth several years after the round. We take these implications to the data in the sections that follow.

3 Data, Measurement, and Empirical Design

The model centers on a financing-allocation decision: at a financing event, the capital raised funds current investment and a liquidity buffer retained against future refinancing needs. In this section, we combine PitchBook financing data with Revelio workforce data to construct the empirical counterpart to this decision. We first describe the data and sample construction. We then develop the Hiring Burn Rate (HBR), a salary-based measure of the new human-capital investment committed per dollar raised. Next, we establish the conditions under which a higher HBR maps into a thinner implied liquidity buffer and assess

whether HBR exhibits the timing and economic properties required for this interpretation. Finally, we present the predictive specifications used to evaluate the framework's empirical implications.

3.1 Data Sources and Sample Construction

The analysis sample consists of U.S.-headquartered startups whose first financing round is classified as Seed or Early Stage VC and occurs between 2010 and March 2023. Each startup contributes one financing-round observation and is followed from that event through subsequent outcomes.

PitchBook provides round-level data on startup financings, including deal type, deal date, round sequence, amount raised, valuations, exit events, and firm characteristics. Revelio provides individual-level human-capital data on startup founders and employees, including employment dates, job functions, seniority, salary estimates, career histories, demographic characteristics, and education-and-network prestige scores. Combining these sources allows us to measure startups' workforce investment around the focal raise relative to the capital raised and to trace its association with subsequent financing, survival, and real outcomes.

We link PitchBook companies to Revelio using LinkedIn profiles or company websites. We then restrict the analysis to startups whose first financing event is classified as a Seed or Early Stage VC round. We further require that the round (i) has a valid financing date and a positive deal size; (ii) has at least one employee with a nonmissing salary estimate during the observation window; and (iii) has nonmissing information on startup age, industry, financing year, and state. Finally, we truncate the sample in March 2023 to ensure at least 36 months

of post-financing data for measuring the primary follow-on financing outcome.

We aggregate Revelio employment records to the startup-month level to construct monthly measures of employee counts, salary spending, and workforce composition by function and seniority. The monthly panel allows us to trace net changes in employment and salary spending from one year before each financing event to five years afterward. The individual career histories also allow us to characterize founder and employee human capital at each financing round. Product activity is measured using Capital IQ Key Developments classified as new-product-launch events by FinBERT, a finance-domain language model used for text classification. Patent activity is measured using matched USPTO filings. These measures capture product launches and patent filings recorded in the matched data sources.

The final analysis sample contains 17,583 startups, each contributing one first-financing observation. The matched records cover 36,873 founders and 1,474,501 employees, observed from one year before to five years after the financing event. These are predominantly white, highly skilled workers: about two-thirds are classified as White, over half majored in STEM, and one in eight founders holds an Ivy-8 or Ivy-Plus undergraduate degree. They are also entrepreneurially selected: 37% of founders and 19% of employees previously worked at a VC-backed startup, and one in ten founders had founded one themselves. Appendix B Tables XI and XII report the full distributions. Approximately 72% of the startups operate in technology sectors. The data therefore provide an individual-level view of the early-stage workforce of VC-backed startups, covering highly skilled talent in innovation-oriented firms.

3.2 Measuring the Investment–Liquidity Trade-off

3.2.1 Hiring Surges Around Financing

In the model, the financing decision allocates newly raised capital between current investment and a liquidity buffer. An empirical counterpart to this decision should therefore capture new investment undertaken in connection with the financing event, relative to the amount raised. We begin by documenting the timing of startups' human-capital investment. Figure 1 plots the monthly change in salary spending, scaled by the amount raised, in Panel A and the monthly change in employee count per $1 million raised in Panel B. Both are shown from 12 months before to 12 months after the financing event, separately by financing stage. Month 0 denotes the financing month. We winsorize both measures at the 1st and 99th percentiles before averaging.

Monthly changes in salary spending and employee count are positive throughout the event window, indicating continued workforce expansion on average. For both measures, the increase begins about six months before financing, accelerates immediately before the event, peaks in month 0, and declines toward its pre-financing level over the following six months. The pattern is strongest for Seed rounds, followed by Early Stage VC rounds. Later Stage VC rounds show a modest spike in the financing month but no comparable run-up. The pre-financing buildup is also consistent with the timing of the broader fundraising process, in which investor screening, meetings, and diligence occur before closing. Survey evidence indicates that the average VC deal takes approximately 83 days to close, with shorter timelines for early-stage and information-technology investments (Gompers et al., 2020).

Appendix B Table I quantifies this concentration. During the 12-month window from six months before to six months after financing, monthly salary spending increases by an average of $74,000 and employee count increases by 7.9. Among startups with available longer-horizon measures, the increase during this window accounts, on average, for 77.0% of the net increase in salary spending and 71.2% of the net increase in employee count over a 24-month horizon beginning six months before financing and truncated at the month before the next round. The corresponding shares over a similarly constructed 60-month horizon are 74.7% and 69.4%. Thus, the expansion of early-stage workforces is concentrated around the focal financing event rather than occurring at a uniform rate over time. The change in salary spending over this window therefore captures the human-capital investment associated with the focal raise.

3.2.2 Hiring Burn Rate

We define the salary-based Hiring Burn Rate (HBR) as the net increase in monthly salary spending over the financing-centered 12-month window, scaled by the amount raised in the focal round. For startup i in financing round $n-1$,

$$\mathrm{HBR}^{\mathrm{sal}}_{i,n-1} = \frac{\Delta\mathrm{SalarySpending}_{i,n-1,[M-6,M+6]}}{\mathrm{AmountRaised}_{i,n-1}}, \tag{6}$$

where M denotes the financing month. The numerator is the net change in monthly salary spending between $M-6$ and $M+6$, and the denominator is the amount raised in the focal round.

HBR measures how intensively a startup converts a new financing round into new human-

capital investment. For example, an HBR of 0.10 means that the startup adds $100,000 in monthly salary spending for every $1 million raised. The measure is therefore about the new labor investment associated with the round, rather than the startup's absolute payroll level.

The use of a change rather than a level is deliberate. Pre-financing payroll reflects investment commitments made before the focal round. It is observable to investors when the round is negotiated and does not describe how the newly raised capital is allocated. The change in payroll around the financing event instead captures the new human-capital investment undertaken with the focal raise. We control separately for pre-round salary runway throughout the analysis.

Human capital is a central input to early-stage startup investment. This is particularly relevant in our sample, where 72% of startups operate in technology sectors and rely heavily on skilled labor. For Internet and software startups, cloud computing allows computing and storage capacity to be rented in small increments rather than requiring large upfront hardware investments (Ewens et al., 2018). Early-stage physical-capital requirements are therefore limited, making human capital a particularly important component of startup investment. Salary spending captures both the number of new hires and the cost of those hires, whereas employee count captures only the quantity margin. This distinction is economically important because the new payroll becomes a recurring cash obligation once employees are hired. Consistent with this interpretation, Berman and Hernández-Lagos (2019) model startups as allocating a limited budget between human and non-human inputs, with workforce spending reflecting that investment decision.

Industry evidence also indicates that personnel spending is the main component of startup operating burn. People-related costs account for roughly three quarters of operating costs

at venture-backed startups, and headcount accounts for approximately 70% of operating expenses at B2B SaaS firms.[6] Even as Salesforce and NetSuite approached their IPOs, each spent approximately $0.80 on payroll per dollar of revenue.[7] In practice, investors and founders commonly approximate startup burn using headcount multiplied by a fully burdened employee cost.[8]

Neither total startup expenditure nor the model's liquidity policy is consistently observed across private firms. In the model, the amount raised is allocated between current investment and a liquidity buffer. If total investment were observable, the portion of the raise not committed to investment would identify the implied liquidity buffer. We therefore focus on new salary spending, a central component of startup investment that our data measure precisely and uniformly at the employee-month level. This focus does not minimize the importance of nonlabor expenses, including marketing, cloud services, software, and rent. Because HBR excludes these expenses, it understates total cash-burn intensity, and its implied runway overstates the startup's full cash runway. HBR nevertheless provides a clean measure of the new hiring investment commitment made with the focal raise.

3.2.3 Mapping HBR to Liquidity Policy

We connect this measure to the conceptual framework, which lets us read HBR as a measure of the model's liquidity buffer ℓ, or equivalently the continuation cutoff $\kappa = \ell + \rho_0^G$. In the model the date-0 raise is the investment outlay plus the prefunded buffer. Using the

[6] Kruze Consulting (2024), "Startup Payroll Costs and Tax Budgeting for New Hires," reports that payroll-related expenses account for 76.4% of operating costs among more than 100 funded startups. OpenView (2023), "The 3 Part Framework for Designing Efficient B2B SaaS Organizations," reports that headcount accounts for approximately 70% of B2B SaaS operating expenses.

[7] Tunguz (2016), "What Percentage of Revenue Should SaaS Startups Spend on Payroll?"

[8] Wilson (2011), "Burn Rates: How Much?," proposes approximately $10,000 per employee-month as a fully burdened burn benchmark.

equilibrium scale $I = k_G(\kappa)A$ of Lemma 1,

$$\text{Raise} \;=\; \underbrace{(I - A)}_{\text{outlay net of founder equity}} \;+\; \underbrace{\ell\, I}_{\text{prefunded buffer}} \;=\; A\big[k_G(\kappa)\,(1+\ell) - 1\big]. \tag{7}$$

Two features of our setting, established above, make the mapping immediate. Labor is the investment good: personnel expenditure is first-order, so investment scale I and monthly salary spending are linked through $\text{SalarySpending} = I/12$. And the surge window captures the dominant share of the deployment of the round, so the window increase in monthly salary spending approximates the round's monthly labor-cost commitment, $\Delta\text{SalarySpending} \approx I/12$. The empirical funding variable is the gross raise in (7) under the prefunded implementation, which is the realistic one for venture rounds; the prefunding-equivalence result of A, that a prefunded cash buffer and a committed credit line break even on identical terms at a zero interest rate, is what allows us to use this gross object. We assume the measurement-monotonicity condition $f(\kappa)\,\ell_G(\kappa) < 1$ on $(\rho_0^G, \rho^*]$, where $\ell_G(\kappa) = \kappa - \rho_0^G$; it holds for the uniform distribution and whenever the shock density near the cutoff is not too high, and it restricts the measurement mapping rather than the equilibrium.

Combining $\Delta\text{SalarySpending} = k_G(\kappa)A/12$ with (7) and cancelling assets gives the burn rate as a function of the liquidity cutoff alone,

$$\text{HBR}^{\text{sal}}(\kappa) \;=\; \frac{\Delta\text{SalarySpending}}{\text{Raise}} \;=\; \frac{1}{12\left[\,1 + \ell(\kappa) - \dfrac{1}{k_G(\kappa)}\,\right]}, \qquad \ell(\kappa) = \kappa - \rho_0^G. \tag{8}$$

This expression makes precise the sense in which HBR is informative about the model. It is independent of founder assets A, so it reflects the per-unit liquidity policy rather than

firm size, and under the measurement-monotonicity condition it is strictly decreasing in the cutoff κ. The signaling distortion of Section 2, a lower cutoff $\hat{\kappa}_F < \rho^*$ for the good type, therefore raises observed HBR, and reading across the funded cross-section the good type's burn rate exceeds the bad type's. This is the mapping from the model's "raise big, run thin" prediction to a higher measured burn rate. We treat it as motivation for the measure rather than as a formal result; the underlying break-even and multiplier algebra is in A.

The measure is equivalent to an implied cash runway. Monthly incremental salary spending is ΔSalarySpending, so the number of months of spending that the raise covers is

$$\text{Runway (months)} = \frac{\text{Raise}}{\Delta\text{SalarySpending}} = \frac{1}{\text{HBR}^{\text{sal}}} = 12\left[1 + \ell(\kappa) - \frac{1}{k_G(\kappa)}\right], \qquad (9)$$

the reciprocal of salary HBR, and increasing in the buffer ℓ under the same condition. A high burn rate is thus the same statement as a short implied runway, and the prediction that quality is increasing in burn is equivalently a prediction that quality is decreasing in runway. Implied runway gives a robustness companion denominated in survival time rather than salary-spending intensity, which severs any mechanical link between burn and employment growth.

3.3 Economic Magnitudes

Table I reports the distribution of HBR in the baseline sample. Salary-spending HBR has a mean of 0.062, a median of 0.022, and a standard deviation of 0.153. Its 25th and 75th percentiles are 0.005 and 0.054, respectively. The distribution is right-skewed and displays substantial variation in the rate at which startups commit a financing dollar to new payroll.

The mean and median correspond to an additional $62,000 and $22,000 of monthly salary spending per $1 million raised, respectively. The median startup has five employees before financing and raises $1.97 million in the focal round.

The economic implications of this variation are substantial. At the mean HBR of 0.062, a startup commits $62,000 of monthly salary spending per $1 million raised, corresponding to approximately 16 months of salary-based implied runway. At the median HBR of 0.022, the corresponding implied runway is approximately 45 months. A startup one standard deviation above the mean has less than five months of implied runway. Thus, startups differ markedly in how quickly they commit a financing round to new hiring.

The observed salary increments also have an economically plausible per-hire scale. During the financing-centered window, monthly salary spending increases by $74,000 on average while employee count increases by 7.9, implying approximately $9,400 per added employee-month. This figure is close to the $10,000 fully burdened monthly employee-cost benchmark reported by Wilson (2011), providing an external check on the economic scale of the salary measure.

3.4 Empirical Approach

Under the quality-dominance case in Section 2, in which the information about startup quality embedded in HBR dominates the associated liquidity risk, the framework predicts that higher HBR forecasts stronger financing and real outcomes. Specifically, higher HBR should be associated with a higher probability of securing the Next Round, greater valuation growth, lower firm failure, and stronger long-run employment, product, and patent outcomes.

Our baseline specifications test these predictions using regressions of the form:

$$Y_i = \alpha + \beta\, \mathrm{HBR}^{\mathrm{sal}}_{i,n-1} + X'_{i,n-1}\gamma + \lambda_{j(i)\times e(i)\times f(i)} + \tau_{t(i)} + \mu_{s(i)} + \varepsilon_i, \tag{10}$$

where i indexes startups and $n-1$ denotes each startup's first financing round. Y_i denotes the subsequent outcome of interest, $\mathrm{HBR}^{\mathrm{sal}}_{i,n-1}$ is salary-spending HBR, and $X_{i,n-1}$ is the vector of controls.

The control vector varies across four progressively richer baseline specifications and includes startup, investor, founder, and employee characteristics that may jointly predict HBR and subsequent outcomes. Specification (1) controls for pre-round salary runway and startup age. Specification (2) adds industry hotness, the log number of experienced current-round investors, and the log number of investors in the current-round syndicate. Industry hotness is the log number of U.S. startups receiving a first Seed or Early Stage VC financing in the same industry during the preceding two quarters. An investor is classified as experienced if it has completed at least five VC investments in the same industry during the preceding two years. Specification (3) adds founder prior-work experience, the founder prestige score, and serial-founder status. Specification (4) adds employee prior-work experience, the employee prestige score, and the share of employees with prior big-firm experience. Founder and employee prestige scores are Revelio-calculated measures of education-and-network prestige, initialized using world-university rankings and updated through worker-mobility networks.

The sequential specifications assess whether the relation between HBR and subsequent outcomes persists after conditioning on increasingly rich observable startup, investor, founder, and employee characteristics. This approach reduces observable confounding and allows us

to assess the stability of the HBR coefficient across specifications.

The fixed effects $\lambda_{j(i)\times e(i)\times f(i)}$ interact industry group, pre-round employment-size quartile, and focal-round fund-size quartile. The size quartiles are constructed from the pooled distribution of all financing rounds. These fixed effects allow us to compare startups in the same industry and the same pre-round employment-size and focal-round fund-size quartiles, while absorbing unobservable factors associated with industry and the joint scale of the startup and financing round. Financing-year fixed effects $\tau_{t(i)}$ absorb market-wide changes in financing conditions, while startup-state fixed effects $\mu_{s(i)}$ account for persistent geographic differences in startup and labor-market conditions.

We restrict the sample to each startup's first financing round to avoid conditioning the analysis on endogenous selection into subsequent rounds and to limit confounding from prior financing, hiring, and performance histories.

These design choices reduce observable heterogeneity and dynamic selection, but HBR may still reflect unobservable investment opportunities or managerial expectations that also predict future performance. We therefore interpret β as a conditional predictive association rather than a causal effect. Section 5 allows the adjustment for observable characteristics to be high-dimensional and nonlinear while retaining the same predictive interpretation.

We estimate equation (10) using ordinary least squares. Continuous variables are winsorized at the 1st and 99th percentiles, and standard errors are clustered at the startup level.

4 Empirical Results

This section tests the empirical implications of the framework in Section 2. We first examine the quality-dominance prediction and its persistence. In the separating equilibrium, high-quality startups strategically choose a thinner liquidity buffer at the focal round to signal their type; their greater continuation capacity therefore predicts that high-HBR startups subsequently perform better. We show that higher HBR predicts a greater probability of securing the Next Round. Conditional on continuation, higher HBR also predicts higher valuation growth, lower subsequent firm failure, and stronger long-run real outcomes. We then examine the model's liquidation-risk prediction using product and patent milestones realized after the focal round. HBR strongly predicts financing continuation even when the startup's underlying quality has not yet materialized in a recorded milestone. Moreover, HBR is much more strongly associated with lower failure when no milestone is recorded than with lower failure despite recorded progress, consistent with the greater liquidation exposure created by a thinner buffer. Finally, we show that HBR's predictive content across the workforce is concentrated in product-oriented, senior, and high-wage roles. Section 5 then tests the framework's second implication by asking whether HBR retains predictive content after flexibly removing the observable mechanical component.

4.1 Quality-Dominance Implication

4.1.1 Financing Success

We first examine whether HBR predicts financing continuation. Table II reports estimates of equation (10), with Next Round$_{i,n}$ as the dependent variable.

The variable equals one if the startup's first subsequent relevant event is a follow-on VC round within three years of the focal round, a PE Growth/Expansion financing, or a successful exit, and zero otherwise.[9] We truncate the focal-round sample in March 2023 so that every round has a three-year observation window.

HBR positively predicts the Next Round in every specification. A one-standard-deviation increase in HBR is associated with a 4.9-percentage-point increase in the probability of securing the Next Round in the baseline specification, equivalent to 8.0% of the unconditional probability of 60.8%. The estimate remains significant at the one percent level and stable as we add current-investor, founder, and employee controls, reaching 5.11 percentage points in the full specification.

This result is consistent with the separating equilibrium. High-quality startups strategically choose a thinner buffer at the focal round, while low-quality startups find imitation too costly. HBR therefore identifies the high-quality type, which has greater pledgeable income and continuation capacity. The positive association with the Next Round shows that this quality advantage dominates the additional liquidation risk created by the thinner buffer.

We next examine whether high-HBR startups that continue also obtain better financing terms. Table III defines valuation growth as the log pre-money valuation of the subsequent round minus the log post-money valuation of the focal round.

HBR positively predicts valuation growth across all specifications, significant at the one percent level. In the full specification, a one-standard-deviation increase in HBR is associated with 12.9% higher valuation growth, equivalent to approximately $1.99 million at the average

[9]Successful exits include IPO, M&A, Buyout/LBO, Management Buyout, Reverse Merger, Merger of Equals, PIPE, and Public Investment Secondary Offerings.

pre-money valuation of $15.41 million.

Together, the financing results show that high-HBR startups are more likely to continue and, conditional on continuing, raise capital at higher valuations. Both results support the quality-dominance implication.

4.1.2 Firm Failure

We next examine whether the quality advantage associated with HBR extends beyond financing continuation to firm survival. Table IV reports estimates with Firm Failure as the dependent variable. We trace each startup from the focal round to its first completed exit event. Firm Failure is assigned as one if the event is a failed exit,[10] or if the startup meets any of the following inferred-failure criteria: (1) the pre-money valuation in its last observed VC/PE financing falls below 10% of prior capital raised, indicating severe value destruction; (2) employment after the last VC/PE financing falls below a calibrated distress threshold, indicating operational collapse;[11] or (3) the startup does not secure a successful VC/PE refinancing or exit outcome within five years after its last VC/PE financing. Firm Failure is assigned as zero if the event is a successful exit.[12] Firms whose last VC/PE financing is too recent to observe a full five-year post-financing window are left missing.

HBR is negatively associated with Firm Failure across all specifications. A one-standard-deviation increase in HBR is associated with a 2.9 to 3.4 percentage point lower probability of failure, about 4.3% to 5.7% of the sample mean. The estimate remains statistically

[10]Failed exits include Out of Business, Bankruptcy, and Restart events in PitchBook.

[11]The employment distress threshold is the 75th percentile of employment declines among startups with observed failure events. In the current sample, this cutoff is an 83.3% drop from peak employment, equivalent to remaining employment of no more than 16.7% of peak, measured using average employment in the failure month and prior two months.

[12]Successful exits include IPO, M&A, Buyout/LBO, Management Buyout, Reverse Merger, Merger of Equals, PIPE, and Public Investment Secondary Offering.

significant as we add current-investor, founder, and employee controls.

Firm failure provides a sharper test of quality-dominance because the model contains an opposing force. Proposition 1(iv) shows that the high-quality type accepts greater liquidation risk when it chooses a thinner buffer. Nevertheless, high HBR is associated with lower, rather than higher, firm failure. The quality advantage associated with the high-HBR type therefore dominates the liquidation-risk cost on average.

Appendix B Table IV complements the linear probability estimates with Cox proportional-hazards models that retain censored startups. A one-standard-deviation increase in HBR raises the hazard of securing the Next Round by approximately 15.3% to 17.3% and lowers the hazard of firm failure by approximately 10.7% to 11.5%. These timing results reinforce the baseline incidence estimates: higher-HBR startups are not only more likely to obtain subsequent financing and less likely to fail, but they also do so more quickly along the financing margin and more slowly along the failure margin.

4.1.3 Long-Term Real Outcomes

We finally examine whether HBR predicts persistent differences in real activity. If HBR identifies a durable type rather than a transitory hiring burst, its predictive content should remain visible several years after the focal round. Table V reports employment growth and product and patent activity over the three- and five-year periods following the HBR measurement window.

A one-standard-deviation increase in HBR is associated with stronger real activity after the financing period. Employment growth is 4.3% and 3.6% higher over the three- and five-year horizons, respectively. Product activity is 1.8% and 3.1% higher, while patent activity

is 3.9% and 5.3% higher. The employment and product estimates are significant at the 5% level or better at both horizons; the patent estimates are significant at the 10% level over three years and the 5% level over five years.

These long-run results support two implications of the framework. First, they extend quality-dominance from financing and survival to real operating outcomes. Second, their persistence over three to five years indicates that HBR identifies a durable startup type rather than a temporary increase in hiring around the focal round.

4.2 Realized Milestones, Signaling, and the Liquidity Cost

Product and patent milestones let us separate what the burn rate reveals from what a startup has already demonstrated. HBR is fixed at the focal round, whereas a product launch or patent filing may occur between that round and the subsequent outcome. The split is therefore not an additional signaling decision. It asks whether the information in HBR operates through observable progress or around it, and what happens to startups whose progress is realized but who fail nonetheless.

Table VI reports fixed-effects multinomial logit estimates. Panel A decomposes financing continuation, with no subsequent financing as the omitted category. In the full specification, a one-unit increase in HBR raises the probability of raising again with no recorded milestone by 39.5 percentage points, significant at the one percent level. Where a milestone is recorded, the estimate is 4.8 percentage points and is not statistically distinguishable from zero. The two differ at the one percent level, and the pattern holds in every specification.

The burn rate therefore predicts financing only for startups with nothing visible to show.

Once a startup has launched a product or filed a patent, investors can observe quality directly and the burn rate adds nothing to what they already see. Where it has not, the burn rate is what separates the startups that raise again from those that do not. This is the quality-signaling channel in its sharpest form: a costly action carries information exactly where hard evidence is missing, and becomes redundant once that evidence arrives.

Panel B applies the same decomposition to firm failure, with successful exit as the omitted category. A one-unit increase in HBR lowers the probability of failing with no recorded milestone by 19.5 percentage points, again significant at the one percent level. Where a milestone is recorded, the estimate is 1.8 percentage points and is not statistically distinguishable from zero. The two differ at the five percent level in the full specification.

A high burn rate therefore protects startups that never recorded a milestone, and does not protect those that did. Among startups that fail, those with higher HBR are more likely to have recorded a product launch or patent filing than those with lower HBR. Aggressive burn does not make failure more likely; it changes the composition of failure, moving it away from startups that never had a viable product and toward startups that had one and still could not finance the next stage. This is the cost that Proposition 1(iv) attributes to the thin buffer. The quality advantage that comes with a high burn rate reaches ordinary failure but not the liquidity margin the buffer itself creates, so the dissipative cost falls on the firms whose quality would otherwise have carried them through.

Two limits apply. The difference in Panel B is not statistically significant in specification (3), so the composition result belongs to the full specification. Milestone status is also realized after the focal round and is plausibly affected by the burn rate itself, so these estimates describe which startups fail rather than how the burn rate changes the cause of

failure.

4.3 Workforce Composition

We next examine which part of the hiring commitment carries HBR's predictive content. Appendix B Table V first separates employee quantity from recurring labor cost. The table compares salary-spending HBR with employee-count HBR and then residualizes salary-spending HBR on employee-count HBR, controls, and fixed effects. Salary-spending HBR continues to predict financing continuation, valuation growth, firm failure, and long-run real outcomes after employee-count HBR is removed. Employee-count HBR, by contrast, is not significantly related to the main financing and survival outcomes. The baseline relation therefore reflects the recurring salary commitment, not employee quantity alone.

We then examine where the salary-spending commitment is concentrated within the workforce. Function categories[13] capture the operational skill domain of the salary commitment, seniority categories[14] capture the organizational layer through which startups deploy and integrate workers, and wage quartiles capture the cost and skill intensity of the labor commitment.

Appendix B Table VI decomposes salary-spending HBR by employee function, seniority, and wage quartile. The financing relations are broad across sales, product, R&D, and operations, but the largest one-standard-deviation effects are concentrated in product-function HBR (3.9 percentage points for the Next Round, 7.3% for valuation growth, and -3.5 percentage points for firm failure). The seniority and wage panels point to a similar concentration in

[13]Revelio reports seven job categories. We collapse Engineer into Product, Scientist into R&D, Sales and Marketing into Sales, and Admin, Finance, and Operations into Operations.

[14]Revelio reports seniority on a seven-point scale. We classify scores 1–2 as Junior, 3–4 as Senior, 5–6 as Management, and 7 as Leadership.

higher-cost labor commitments: management and senior HBR have the strongest seniority-level associations with financing and valuation outcomes (2.7 percentage points for the Next Round for management HBR and 7.8% valuation growth for senior HBR), while top-quartile wage HBR has the largest wage-level associations with the Next Round, valuation growth, and later employment growth (3.6 percentage points, 10.9%, and 8.4%, respectively). The real-outcome columns line up more closely with the type of labor added: product HBR is most strongly related to observed product activity, while R&D HBR is most closely related to patent activity and later employment.

Overall, HBR's predictive content is not distributed evenly across payroll expansion. It is concentrated in product-oriented, senior, and high-wage commitments that are central to scaling the startup and costly for a lower-quality firm to sustain with a thin liquidity buffer. This concentration accords with the separating equilibrium's imitation constraint: the most informative commitments are those that are hardest for the low-quality type to mimic.

4.4 Robustness

Appendix B Table IX considers alternative constructions of HBR. It first shortens the baseline $[M-6, M+6]$ salary-spending window to $[M-3, M+3]$, which concentrates on the months immediately surrounding the raise. The narrower measure continues to predict next-round financing, valuation growth, lower firm failure, and three-year employment growth at the one percent level. Its association with product activity remains positive and significant at the ten percent level, whereas the patent estimate is positive but imprecise. The table then constructs HBR from changes in employee count rather than changes in salary spending,

using both windows. Employee-count HBR predicts the financing and failure outcomes under both windows, but its associations with the real outcomes are weaker and less uniform. Thus, the main financing and survival patterns do not depend on the baseline window, while salary spending provides the more informative measure of the hiring commitment.

Appendix B Table X repeats the baseline analysis in a matched sample. For each focal round, we pool startups in the same industry whose first financing is also a Seed or Early Stage VC round, occurs within five years of the focal round, and whose firm age is within 50 percent of the focal startup's age. We retain candidate pairs whose business-description similarity exceeds the 80th percentile of the pooled eligible-pair distribution. Pre-round employment and deal-size ranges begin at 50 to 150 percent of the focal values and widen jointly, only when fewer than five retained peers are available, to a maximum of 50 to 250 percent. Within the narrowest range containing at least five retained peers, or the maximum range otherwise, we select up to five peers with the highest business similarity. With matched-set fixed effects and the full control set, higher HBR continues to predict higher next-round financing, valuation growth, employment growth, product activity, and patent activity, and lower failure. All six estimates are statistically significant at the one percent level. The baseline pattern therefore persists among closely matched startups.

5 Double Machine Learning Evidence

Section 4 shows that HBR predicts financing, survival, and long-run real outcomes after a parsimonious set of startup, investor, founder, and employee controls. This section tests the framework's second implication: whether HBR contains predictive information beyond

observable startup quality and investment opportunities. We use double machine learning to remove nonlinear functions and high-order interactions of 278 startup, deal, founder, and employee characteristics from both HBR and the outcomes. The machine-learning models explain approximately 45 percent of HBR out of sample. Nevertheless, residual HBR continues to predict financing continuation, valuation growth, firm failure, and product activity as the information set expands. Three-year employment growth is the exception: its association with residual HBR is not statistically significant once pre-round workforce characteristics are included.

The conceptual framework makes this distinction substantive. Remark 1 shows that a better firm can retain less liquidity for purely mechanical reasons: greater pledgeable income allows it to finance the same survival target with less cash. This component operates through observable differences in quality and investment opportunity and is survival-neutral. By contrast, Proposition 1(iii) adds a strategic component to the thin buffer. If the baseline relations reflect only the mechanical component or a restrictive functional form, a flexible adjustment for observables should absorb them. If the strategic component contains information about latent type that observables do not span, residual HBR should continue to predict outcomes as the information set expands. Because quality is embedded in founder human capital, workforce composition, deal terms, and market conditions, these relations may be nonlinear and high-dimensional. We therefore implement the test with the double/debiased machine-learning (DML) framework of Chernozhukov et al. (2018), which permits flexible machine-learning adjustment while preserving valid inference on the low-dimensional HBR coefficient.

5.1 Orthogonalized Design

We adopt the partially linear model of Chernozhukov et al. (2018). Let $D_i \equiv \mathrm{HBR}^{\mathrm{sal}}_{i,n-1}$ denote salary-based HBR, let Y_{ik} denote outcome k, and let X_i^m denote the covariates in information set $m \in \{\mathrm{M1}, \mathrm{M2}, \mathrm{M3}\}$. For each outcome and information set,

$$Y_{ik} = \theta_k^m \, D_i + g_k^m(X_i^m) + u_{ik}, \qquad \mathbb{E}[u_{ik} \mid X_i^m, D_i] = 0, \tag{11}$$

$$D_i = q^m(X_i^m) + v_i, \qquad \mathbb{E}[v_i \mid X_i^m] = 0, \tag{12}$$

where g_k^m and q^m are unrestricted functions of the covariates and θ_k^m is the coefficient of interest on HBR. Equation (12) is an auxiliary prediction problem: it asks how much of HBR the covariates can explain, and its residual v_i is the component of HBR that they cannot. By the logic of partialling out (Frisch and Waugh, 1933; Lovell, 1963), θ_k^m is recovered by regressing the outcome net of its conditional mean on this HBR residual.

Estimation proceeds in three steps. First, machine-learning models estimate the two conditional means implied by the model, $\ell_k^m(x) = \mathbb{E}[Y_{ik} \mid X_i^m = x]$ and $q^m(x) = \mathbb{E}[D_i \mid X_i^m = x]$, using training cohorts only. Second, letting $s(i)$ denote the prediction cohort containing startup i, the out-of-fold residuals are

$$\widehat{v}_i^m = D_i - \widehat{q}^{\,m,(-s(i))}(X_i^m), \tag{13}$$

$$\widehat{u}_{ik}^m = Y_{ik} - \widehat{\ell}_k^{\,m,(-s(i))}(X_i^m), \tag{14}$$

where the superscript $(-s(i))$ indicates that the prediction comes from models estimated without startup i's cohort. Third, the second stage regresses the outcome residual on the

HBR residual without an intercept,

$$\widehat{\theta}_k^m = \frac{\sum_i \widehat{v}_i^m \, \widehat{u}_{ik}^m}{\sum_i (\widehat{v}_i^m)^2}. \tag{15}$$

Two ingredients make this construction reliable, and each addresses a distinct threat. Residualizing both the outcome and HBR yields a Neyman-orthogonal score, so estimation error in the two prediction models does not propagate first order into $\widehat{\theta}_k^m$; this guards against regularization bias. Cross-fitting ensures that the predictions used for each startup come from models that never saw that startup in training; this guards against own-observation overfitting (Chernozhukov et al., 2018). Our implementation follows the practical guidelines in Feyzollahi and Rafizadeh (2024). Standard errors are clustered by company and, because the sample contains one focal financing per startup, coincide with heteroskedasticity-robust standard errors. The coefficient $\widehat{\theta}_k^m$ is an orthogonalized conditional association: the procedure protects against overfitting and regularization bias, but it does not make HBR exogenous with respect to characteristics outside the conditioning set.

5.2 Nested Information Sets

The three information sets are strictly nested. M1 contains 151 firm-level variables covering the startup, the focal deal, the financing market, and the size and composition of the workforce in place at month $M - 7$. M2 adds 66 founder variables. M3 adds a further 61 employee variables, for 278 in total. The two added layers measure individual-level human capital, including education, prior work and founding histories, seniority, prestige, and demographics. Moving from M1 to M2 asks whether founder quality accounts for the baseline

relation; moving from M2 to M3 asks whether the quality of the workforce already hired does. Founder-roster identity is treated as time invariant; all time-varying education and career characteristics are measured no later than $M - 7$, before the HBR window opens. Deal size and focal-round terms are included deliberately: deal size is the denominator of HBR, so conditioning on it strips scale variation out of the residual, although their inclusion means the covariates are not strictly pre-treatment. Appendix D lists the variables by tier.

5.3 Out-of-Time Prediction and Model Selection

Predictions are made strictly out of time. We begin by fitting the prediction models to the 2010 through 2014 financing cohorts and use them to generate out-of-fold predictions for startups financed in 2015. The training window then rolls forward one year, so models fitted to the 2011 through 2015 cohorts predict the 2016 cohort, models fitted to the 2012 through 2016 cohorts predict the 2017 cohort, and we repeat this process through 2023. Any startup appearing in a prediction cohort is excluded from that cohort's training window, so no startup is ever predicted by a model that saw it in training. Missing-value imputation and the ensemble weights described below are likewise estimated within each training window and applied unchanged to the prediction cohort. This chronological separation of training and prediction follows the standard protocol in the finance machine-learning literature for avoiding look-ahead bias (Gu et al., 2020; Amini et al., 2021), and the resulting out-of-fold residuals implement the cross-fitting that Chernozhukov et al. (2018) require for valid inference. Because predictions require five years of training history, the second stage covers 13,065 financing rounds, with smaller outcome-specific samples held fixed across M1–M3.

Model choice is empirical. The learner library contains a linear model, LASSO, and six tree-based learners: Random Forest (Breiman, 2001), Extra Trees (Geurts et al., 2006), and four implementations of gradient boosting (Friedman, 2001), namely XGBoost (Chen and Guestrin, 2016), LightGBM (Ke et al., 2017), histogram-based gradient boosting, and CatBoost (Prokhorenkova et al., 2018). Tree ensembles accommodate the sparse, nonlinear, and interactive structure of the feature set, and in related corporate-finance prediction problems they outperform linear models out of sample even when the linear models are augmented with polynomial and interaction terms (Amini et al., 2021). The reported nuisance predictions combine the six tree learners with nonnegative-least-squares weights estimated on time-ordered validation splits inside each outer training window, so no prediction-cohort observation enters tuning or weighting. We report the ensemble because the combination reduces dependence on any single learner; it is not claimed to dominate the best individual model, and Appendix B Table VIII reports each learner separately. Base-learner hyperparameters are fixed before outer evaluation and shared across tiers, targets, and prediction years. Binary outcomes are predicted as conditional probabilities and evaluated with squared error on the probability scale, which matches the first stage's purpose: estimating conditional means for residual construction, not classification.

5.4 How Much of HBR Do Observables Explain?

The machine-learning first stage achieves substantial gains over the linear benchmark in predicting HBR. Appendix B Table VIII reports ensemble out-of-sample R^2 statistics of 0.408, 0.444, and 0.449 under M1, M2, and M3, compared with roughly 0.27 for the linear

model in each tier. The ensemble's root mean squared errors are 10 to 14 percent below those of the linear model, and the differences are significant at the one percent confidence level. The ensemble performs comparably to the best individual learner in each tier. We combine learners rather than select one because no learner is best in every tier, and choosing one on out-of-sample fit would import information from the prediction cohorts into the design. Nearly half of the variation in hiring intensity is thus predictable from pre-round observables, and a substantial share of that predictability reflects nonlinear structure that the linear model misses.

Appendix B Table VII decomposes the gains across layers. Founder characteristics raise the HBR out-of-sample R^2 by 0.036, a significant improvement over startup fundamentals alone. The employee layer adds a further 0.005, an increment that is not statistically significant. The outcome models show a different pattern: founder information improves the prediction of every outcome, whereas employee information matters most for subsequent employment growth, where it contributes an increment of 0.013. The exception becomes relevant below. These fit statistics describe what information the nuisance models remove; they are not themselves estimates of an economic mechanism.

We now turn to which covariates carry the prediction of HBR. Appendix B Figure I illustrates each feature's contribution with out-of-fold TreeSHAP values for the twenty leading predictors in each tier and, to examine whether the drivers are stable over time, tracks their importance ranks across the 2015 through 2023 prediction years. Deal size is the dominant predictor in every tier, as expected from its role as the denominator of HBR; pre-round runway risk and startup age follow. Once the founder layer enters, founding-team size, whether a founder holds the CEO title at founding, and founder prestige join the leading ranks. Few

employee variables do: consistent with the small M2-to-M3 gain in fit, the M3 ranking is nearly identical to M2, with rank correlations of mean absolute SHAP across common features of 0.98 between M1 and M2 and 0.98 between M2 and M3. SHAP values explain the fitted prediction, can divide importance among correlated features, and do not identify causal determinants of payroll policy. Their message for the design is narrower and more useful: the prominence of deal size and pre-round runway risk shows that the first stage absorbs the mechanical, scale-related component of the burn–quality relation, which is precisely its role.

5.5 Residual HBR and Outcomes

Table XI reports the residual-on-residual estimates. Residualized HBR is positively associated with successful continuation, valuation growth, and product activity and negatively associated with resolved failure in every information tier, and the coefficients change only modestly as the conditioning set expands. The continuation coefficient declines from 0.484 under M1 to 0.444 under M2 and 0.435 under M3, an attenuation of about ten percent that remains significant at the one percent confidence level. The valuation coefficient moves from 1.290 to 1.307, the failure coefficient from -0.220 to -0.194, and the product coefficient from 0.227 to 0.218. Under the two implications set out above, this stability indicates that nonlinear functions of the measured startup, founder, and employee information do not account for these four relations.

The magnitudes remain economically meaningful. A 0.10 increase in residual HBR corresponds to \$100,000 of additional monthly salary commitment per \$1 million raised beyond what the information set predicts. Under M3, such a deviation is associated with a 4.4-

percentage-point higher continuation probability, approximately 7 percent of the 60.8 percent unconditional continuation rate; 0.131 additional log points of valuation growth, approximately 14 percent; a 1.9-percentage-point lower probability of failure relative to successful exit in the resolved-outcome sample; and 0.022 additional units of $\log(1 + \text{product count})$. These are associations between out-of-fold residuals, not effects of mechanically increasing payroll.

Three-year employment growth is the exception, and an informative one. Its coefficient is 0.102 under M1, 0.050 under M2, and 0.040 under M3, none statistically distinguishable from zero, so the positive employment association in the conventional specification does not survive the flexible adjustment. The first stage explains why: employment growth is the outcome for which the pre-round workforce layer adds the most predictive content, so the characteristics of the organization in place before financing, rather than residual hiring intensity, carry the employment prediction. This selectivity also argues against a mechanical artifact of the residual-on-residual design, which would manufacture associations in all five outcomes or none. The DML table covers these five outcomes; patent and five-year outcomes are outside the estimation, and the conclusions here do not extend to them.

5.6 Interpretation and Limits

The DML evidence narrows the set of explanations for the burn–quality relation in a specific way. Nonlinear combinations of 278 startup, deal, founder, and employee characteristics, including the deal-size denominator and the pre-round runway risk that dominate the prediction of HBR, do not account for the relation between HBR and continuation, val-

uation, failure, or product activity. The residual component of HBR is, by construction, the part of hiring intensity that no measured variable explains. Because hiring intensity is ultimately an entrepreneurial decision, this residual may reflect private, quality-related information, such as assessments of product-market fit, undisclosed traction, or strategic opportunities observable to insiders but not to the econometrician, or it may reflect luck. The cross-outcome consistency disciplines the interpretation: idiosyncratic shocks would not systematically predict refinancing, revaluation, survival, and product shipment at once. The pattern is instead what the separating equilibrium of Section 2 predicts, in the spirit of Holmström and Tirole (1998): the willingness to operate with a thin liquidity buffer reveals information about venture quality that observables do not span.

These results do not establish that the surviving component is private information alone, and we do not claim a causal effect of payroll expansion. Residual HBR may still contain unobserved financing terms, founder expectations, investor support, local labor-market conditions, or measurement error in estimated salary spending, and any of these could correlate with later outcomes. The first-stage fit implies that roughly half of HBR variation remains unexplained by observables, and HBR is measured through $M + 6$ while continuation and failure can occur inside that window, so the post-window outcomes provide the cleanest temporal separation. Subject to these limits, the appropriate conclusion is that the component of the burn rate that observables cannot explain continues to predict which startups refinance, revalue, survive, and ship products, which is what a costly signal of quality, broadly defined, would look like in these data.

6 Conclusion

This paper studies the investment-liquidity trade-off in venture capital-backed startups. Startups can use newly raised capital to expand quickly, but doing so through recurring labor-cost commitments leaves a thinner liquidity buffer against refinancing risk. We measure this trade-off with Hiring Burn Rate (HBR), the increase in monthly salary spending around the financing event scaled by the amount raised. Using 17,583 first Seed and Early Stage VC financings of U.S.-headquartered startups linked to monthly employment records, we examine whether this thin liquidity policy predicts subsequent financing, survival, and real outcomes.

The evidence shows that HBR is a robust predictor of startup performance. Startups with higher HBR are more likely to secure the next financing round, raise at higher valuations, and fail less often. The relation also extends beyond the financing event: higher-HBR startups exhibit stronger longer-run employment growth and more observed product and patent activity after the HBR window. These patterns are consistent with the quality-dominance case in the conceptual framework, in which high-quality startups can credibly sustain a thinner liquidity buffer because they have greater pledgeable continuation value.

Additional tests help clarify what HBR reveals. The HBR relation is stronger for startups whose subsequent financing or failure outcomes are not accompanied by a recorded product launch or patent filing, suggesting that HBR contains predictive content beyond observable product and patent milestones. The relation is also not explained by employee-count growth alone: the predictive content is concentrated in salary-spending HBR and, within salary spending, in product-oriented and high-wage labor commitments. The double

machine learning analysis further shows that the component of HBR not explained by 278 startup, deal, founder, and employee characteristics continues to predict continuation, valuation growth, lower failure, and product activity. Together, these findings point to HBR as an informative measure of costly workforce commitments that observable startup characteristics do not fully span.

We do not interpret the estimates as the causal effect of increasing salary spending. High-HBR startups may differ from low-HBR startups in unobserved investment opportunities, managerial expectations, investor support, or private information, and neither the baseline controls nor the DML design makes HBR exogenous. The contribution is instead to introduce HBR as a new, high-frequency measure of how startups trade liquidity for growth and to show that it strongly predicts which startups subsequently refinance, obtain higher valuations, survive, and generate real activity. This predictive content extends beyond standard observable characteristics and is concentrated in costly workforce commitments that are difficult for weaker startups to sustain, consistent with the model's costly-signal interpretation.

Overall, the results show that cash burn is not uniformly a sign of wasteful overexpansion, nor is cash conservation uniformly optimal. In early-stage venture finance, the same payroll commitment that shortens runway can also reveal information about continuation prospects. This paper contributes to entrepreneurial finance by providing a direct empirical measure of the investment-liquidity trade-off and by showing that workforce commitments can serve as an information-revelation mechanism in staged financing markets. A natural next step is to examine how the value and risk of this signal vary across financing cycles, especially when external capital becomes scarce.

References

Shahram Amini, Ryan Elmore, Özde Öztekin, and Jack Strauss. Can machines learn capital structure dynamics? *Journal of Corporate Finance*, 70:102073, 2021.

Ron Berman and Pablo Hernández-Lagos. Entrepreneur's knowledge, firm survival and the normalized burn rate. SSRN Working Paper 3020881, Social Science Research Network (SSRN), 2019. Available at SSRN.

Fabio Bertoni, Massimo G. Colombo, and Luca Grilli. Venture capital financing and the growth of high-tech start-ups: Disentangling treatment from selection effects. *Research Policy*, 40(7):1028–1043, 2011.

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

Thomas J. Chemmanur, Karthik Krishnan, and Debarshi K. Nandy. How does venture capital financing improve efficiency in private firms? a look beneath the surface. *The Review of Financial Studies*, 24(12):4037–4090, 2011.

Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. ISSN 1368-4221. doi: 10.1111/ectj.12097.

In-Koo Cho and David M. Kreps. Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102(2):179–221, 1987. doi: 10.2307/1885060.

Annalisa Croce, José Martí, and Samuele Murtinu. The impact of venture capital on the productivity growth of european entrepreneurial firms: 'screening' or 'value added' effect? *Journal of Business Venturing*, 28(4):489–510, 2013.

Antonio Davila, George Foster, and Mahendra Gupta. Venture capital financing and the growth of startup firms. *Journal of Business Venturing*, 18(6):689–708, 2003.

Michael Ewens, Ramana Nanda, and Matthew Rhodes-Kropf. Cost of experimentation and the evolution of venture capital. *Journal of Financial Economics*, 128(3):422–442, 2018. doi: 10.1016/j.jfineco.2018.03.001.

Maryam Feyzollahi and Nima Rafizadeh. Double/debiased machine learning for economists: Practical guidelines, best practices, and common pitfalls. SSRN Working Paper 4703243, Social Science Research Network (SSRN), January 2024. Available at SSRN.

Jerome H. Friedman. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232, 2001.

Ragnar Frisch and Frederick V. Waugh. Partial time regressions as compared with individual trends. *Econometrica*, 1(4):387–401, 1933. doi: 10.2307/1907330.

Laurent Frésard, Loriano Mancini, Enrique J. Schroth, and Davide Sinno. How do firms choose between growth and efficiency? Technical Report 23-37, Swiss Finance Institute, June 2023.

Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. *Machine Learning*, 63(1):3–42, 2006.

Paul A. Gompers, Will Gornall, Steven N. Kaplan, and Ilya A. Strebulaev. How do venture capitalists make decisions? *Journal of Financial Economics*, 135(1):169–190, 2020. ISSN 0304-405X. doi: 10.1016/j.jfineco.2019.06.011.

Shihao Gu, Bryan Kelly, and Dacheng Xiu. Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5):2223–2273, 2020.

Robert E. Hall and Susan E. Woodward. The burden of the nondiversifiable risk of entrepreneurship. *American Economic Review*, 100(3):1163–1194, 2010. doi: 10.1257/aer.100.3.1163.

Thomas Hellmann and Manju Puri. The interaction between product market and financing strategy: The role of venture capital. *The Review of Financial Studies*, 13(4):959–984, 2000.

Thomas Hellmann and Manju Puri. Venture capital and the professionalization of start-up firms: Empirical evidence. *Journal of Finance*, 57(1):169–197, 2002.

Bengt Holmström and Jean Tirole. Private and public supply of liquidity. *Journal of Political Economy*, 106(1):1–40, February 1998. doi: 10.1086/250001.

Bengt Holmström and Jean Tirole. Liquidity and risk management. *Journal of Money, Credit and Banking*, 32(3):295–319, 2000. doi: 10.2307/2601167.

Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

William R. Kerr, Ramana Nanda, and Matthew Rhodes-Kropf. Entrepreneurship as experimentation. Technical Report 15-005, Harvard Business School, July 2014. URL https://ssrn.com/abstract=2473226.

Samuel Kortum and Josh Lerner. Assessing the impact of venture capital on innovation. *RAND Journal of Economics*, 31(4):674–692, 2000.

Saerom Lee and J. Daniel Kim. When do startups scale? large-scale evidence from job postings. *Strategic Management Journal*, 2024. Forthcoming.

Josh Lerner and Ramana Nanda. Venture capital's role in financing innovation: What we know and how much we still need to learn. *Journal of Economic Perspectives*, 34(3): 237–261, 2020.

Michael C. Lovell. Seasonal adjustment of economic time series and multiple regression analysis. *Journal of the American Statistical Association*, 58(304):993–1010, 1963. doi: 10.1080/01621459.1963.10480682.

José Martí Pellón and Luisa Alemany. Productivity growth in spanish venture-backed firms. In G. Gregoriou, M. Kool, and R. Kraeussl, editors, *Venture Capital in Europe*, chapter 7. Elsevier, 2006. URL https://ssrn.com/abstract=1087954.

Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: Unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

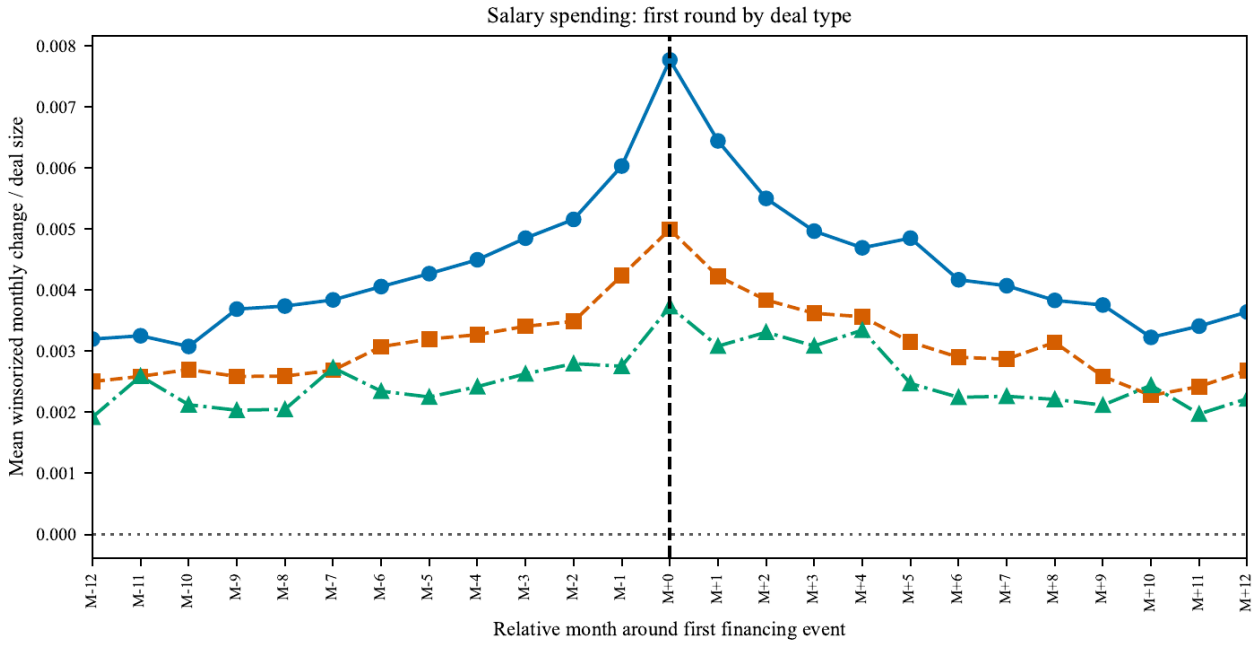
Manju Puri and Rebecca Zarutskie. On the lifecycle dynamics of venture-capital-and non-venture-capital-financed firms. *The Journal of Finance*, 65(6):2247–2293, 2010.

Jean Tirole. *The Theory of Corporate Finance*. Princeton University Press, Princeton, NJ, 2006.

Figure 1. Hiring Surges Around the First Financing Round

This figure plots monthly hiring activity around the first financing event. Panel A reports monthly changes in salary spending and Panel B monthly changes in employee count, each per million dollars raised. The window spans 12 months before to 12 months after financing, with month 0 denoting the financing month, and patterns are shown separately by deal type. Both measures are winsorized at the 1st and 99th percentiles.

Panel A. Surges in Salary Spending



Panel B. Surges in Employee Count

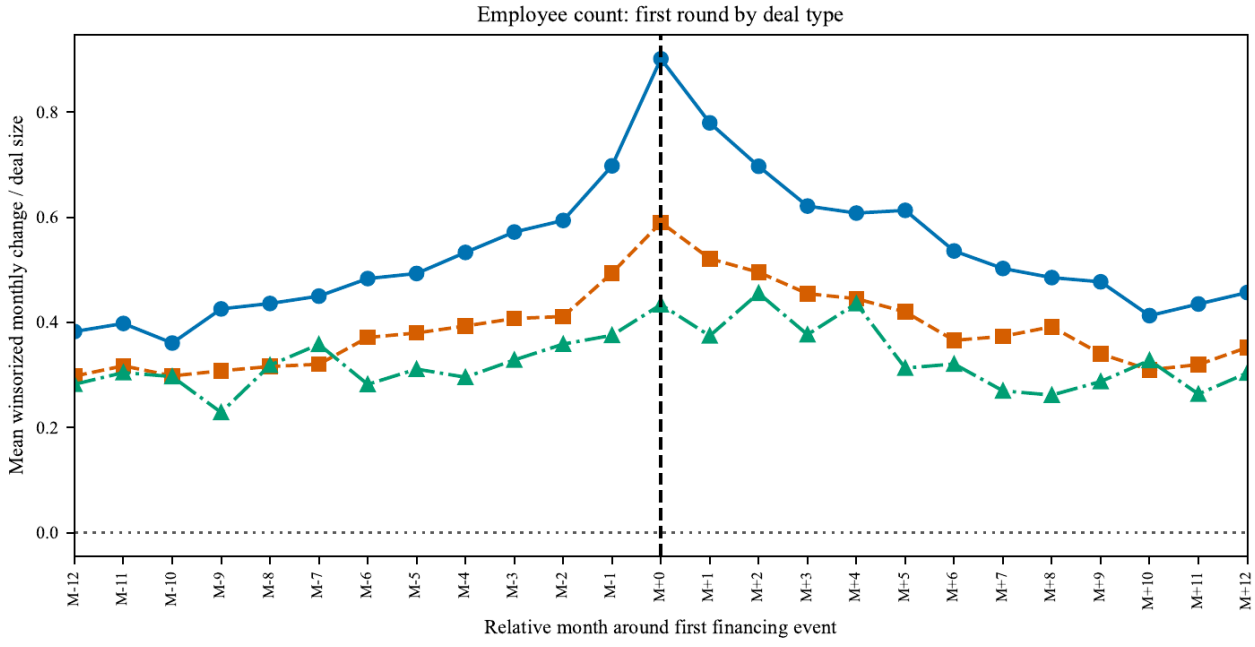




Table I. Summary Statistics

This table presents summary statistics for the variables used in the analysis. Panel A presents the outcome variables. Panel B presents Hiring Burn Rate (HBR) measures constructed using changes in monthly salary spending or employee count over the indicated windows, per million dollars raised. Panel C presents the control variables used in the regression specifications. Panel D presents financing-round characteristics, SaaS and AI indicators, and product and patent activity. Expected Runway is the focal amount raised divided by monthly total compensation spending at the financing month, expressed in months. Realized Runway is the financing-clock duration from the focal round to the subsequent financing or successful exit, expressed in months and reported for observations with Next Round equal to one. Ln(Products), Prior 3 Years and Ln(Patents), Prior 3 Years measure product and patent activity during the three years before the focal round. Product Observed at Any Time (Patent Observed at Any Time) is an indicator equal to one if at least one product (patent) is observed for the startup between its founding date and 2026, and zero otherwise. Each observation is a financing round. Salary-spending and amount-raised variables are expressed in millions of dollars. Continuous variables are winsorized at the 1st and 99th percentiles, except outcome variables.

Variable	Obs	Mean	Median	Std. Dev.	P25	P75
Panel A: Outcomes						
Next Round	17,583	0.608	1	0.488	0	1
Δ Ln(Valuation)	5,715	0.715	0.647	0.684	0.284	1.088
Firm Failure	8,153	0.667	1	0.471	0	1
Δ Ln(Employment)$_{t,t+3}$	16,908	0.221	0.262	0.933	-0.288	0.821
Δ Ln(Employment)$_{t,t+5}$	12,406	0.254	0.288	1.186	-0.526	1.066
Ln(Product)$_{t,t+3}$	5,392	0.291	0	0.447	0	0.693
Ln(Product)$_{t,t+5}$	4,368	0.448	0	0.547	0	0.693
Ln(Patent)$_{t,t+3}$	4,074	0.864	0.693	0.862	0	1.386
Ln(Patent)$_{t,t+5}$	3,372	1.155	1.099	0.992	0.693	1.792
Panel B: Hiring Burn Rate						
HBR$_{\text{[M-6,M+6], Salary-spending}}$	17,583	0.062	0.022	0.153	0.005	0.054
HBR$_{\text{[M-3,M+3], Salary-spending}}$	17,583	0.035	0.011	0.095	0.001	0.030
HBR$_{\text{[M-6,M+6], Employee-count}}$	17,583	7.217	2.200	18.599	0.500	6
HBR$_{\text{[M-3,M+3], Employee-count}}$	17,583	4.101	1.111	11.639	0	3.333
Panel C: Main Controls						
Runway Risk PreRound	17,583	0.091	0.028	0.221	0.009	0.071
Ln(Founder Prior-Work Exp.)	17,583	2.211	2.526	1.015	1.946	2.904
Serial Founder	17,583	0.366	0.333	0.391	0	0.667
Founder Prestige Score	17,583	0.445	0.509	0.308	0.205	0.695
Ln(Employee Prior-Work Exp.)	17,583	1.737	2.079	1.060	1.099	2.526
Employee Prestige Score	17,583	0.304	0.314	0.274	0	0.519
Employee Prior Big-Firm Share	17,583	0.528	0.600	0.394	0	0.909
Industry Hotness	17,583	4.933	4.875	1.440	3.892	6.277
Ln(Experienced Investors)	13,003	0.549	0	0.656	0	1.099
Ln(Investor Syndicate)	13,003	1.385	1.386	0.656	0.693	1.792
Ln(Startup Age)	17,583	0.914	0.693	0.625	0.693	1.386
Panel D: Financing Round Characteristics						
Employee Count Pre-Round	17,583	9.606	5	14.976	2	11
Amount Raised	17,583	4.266	1.970	8.093	0.750	4.150
Percent Acquired	9,497	28.158	26.560	14.164	18.920	35.720
Expected Runway (Months)	16,918	28.287	11.457	58.680	5.055	25.085
Realized Runway (Months)	10,691	26.685	20.074	19.793	12.912	34.957
Tech (%)	17,583	72.109	100	44.847	0	100
SaaS (%)	16,279	35.469	0	47.843	0	100
AI (%)	16,279	20.996	0	40.729	0	0
Ln(Product), Prior 3Y	5,494	0.022	0	0.122	0	0
Ln(Patent), Prior 3Y	4,396	0.295	0	0.529	0	0.693
Product Observed at Any Time	17,583	0.312	0	0.464	0	1
Patent Observed at Any Time	17,583	0.250	0	0.433	0	1

Table II. The Probability of Securing the Next Financing Round

This table reports ordinary least squares (OLS) estimates of the relation between HBR and the probability of securing the next financing round. The dependent variable, Next Round, equals one if the startup raises a follow-on financing round after the focal round and zero otherwise. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. Column (1) controls for pre-round runway risk and startup age. Columns (2)-(4) sequentially add focal-round investor characteristics, pre-round founder characteristics, and pre-round employee characteristics. All specifications include industry-by-employment-size-quartile-by-fund-size-quartile fixed effects, financing-year fixed effects, and state fixed effects. Each observation is a financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Standard errors are clustered at the company level and reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	Next Round			
	(1)	(2)	(3)	(4)
Hiring Burn Rate	0.318***	0.336***	0.331***	0.337***
	(0.030)	(0.036)	(0.036)	(0.035)
Runway Risk PreRound	−0.189***	−0.202***	−0.203***	−0.216***
	(0.021)	(0.025)	(0.025)	(0.025)
Ln(Founder Prior-Work Exp.)			0.024***	0.023***
			(0.006)	(0.006)
Founder Prestige Score			0.084***	0.067***
			(0.017)	(0.017)
Serial Founder			-0.014	-0.012
			(0.012)	(0.012)
Ln(Employee Prior-Work Exp.)				-0.000
				(0.006)
Employee Prestige Score				0.087***
				(0.019)
Employee Prior Big-Firm Share				0.033**
				(0.015)
Industry Hotness		−0.051***	−0.050***	−0.050***
		(0.017)	(0.017)	(0.017)
Ln(Experienced Investors)		0.059***	0.060***	0.060***
		(0.010)	(0.010)	(0.010)
Ln(Investor Syndicate)		0.031***	0.031***	0.028***
		(0.008)	(0.008)	(0.008)
Ln(Startup Age)	−0.099***	−0.076***	−0.062***	−0.067***
	(0.006)	(0.008)	(0.008)	(0.008)
Industry x Emp Size Q x Fund Size Q FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Startup controls	Y	Y	Y	Y
Current investor controls	N	Y	Y	Y
Founder controls	N	N	Y	Y
Employee controls	N	N	N	Y
N	17,583	12,962	12,962	12,962
Adj. R^2	0.077	0.088	0.093	0.096

Table III. The Log Valuation Change Between the Focal and Next Financing Rounds

This table reports ordinary least squares (OLS) estimates of the relation between HBR and the log valuation change between the focal and next financing rounds. The dependent variable, Δ Ln(Valuation), measures the difference in ln startup valuation between the two rounds. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. Column (1) controls for pre-round runway risk and startup age. Columns (2)-(4) sequentially add focal-round investor characteristics, pre-round founder characteristics, and pre-round employee characteristics. All specifications include industry-by-employment-size-quartile-by-fund-size-quartile fixed effects, financing-year fixed effects, and state fixed effects. Each observation is a financing round with an observed valuation at the next financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Standard errors are clustered at the company level and reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	Δ Ln(Valuation)			
	(1)	(2)	(3)	(4)
Hiring Burn Rate	1.112***	1.224***	1.234***	1.228***
	(0.141)	(0.163)	(0.165)	(0.164)
Runway Risk PreRound	-0.082	-0.179	-0.196	-0.187
	(0.143)	(0.172)	(0.174)	(0.172)
Ln(Founder Prior-Work Exp.)			−0.078***	−0.071***
			(0.015)	(0.015)
Founder Prestige Score			0.204***	0.178***
			(0.041)	(0.042)
Serial Founder			-0.012	-0.010
			(0.026)	(0.026)
Ln(Employee Prior-Work Exp.)				−0.042***
				(0.015)
Employee Prestige Score				0.107**
				(0.047)
Employee Prior Big-Firm Share				0.023
				(0.037)
Industry Hotness		-0.029	-0.030	-0.032
		(0.043)	(0.043)	(0.043)
Ln(Experienced Investors)		0.040*	0.032	0.031
		(0.022)	(0.022)	(0.022)
Ln(Investor Syndicate)		0.017	0.010	0.011
		(0.019)	(0.019)	(0.019)
Ln(Startup Age)	−0.084***	−0.097***	−0.105***	−0.099***
	(0.019)	(0.020)	(0.021)	(0.021)
Industry x Emp Size Q x Fund Size Q FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Startup controls	Y	Y	Y	Y
Current investor controls	N	Y	Y	Y
Founder controls	N	N	Y	Y
Employee controls	N	N	N	Y
N	5,639	4,693	4,693	4,693
Adj. R^2	0.099	0.113	0.121	0.122

Table IV. The Probability of Firm Failure

This table reports ordinary least squares (OLS) estimates of the relation between HBR and the probability of firm failure. The dependent variable, Firm Failure, equals one if the startup fails and zero if it has a successful exit. Ongoing firms with unresolved outcomes are excluded. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. Column (1) controls for pre-round runway risk and startup age. Columns (2)-(4) sequentially add focal-round investor characteristics, pre-round founder characteristics, and pre-round employee characteristics. All specifications include industry-by-employment-size-quartile-by-fund-size-quartile fixed effects, financing-year fixed effects, and state fixed effects. Each observation is a financing round with an observed firm outcome. Continuous variables are winsorized at the 1st and 99th percentiles. Standard errors are clustered at the company level and reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	Firm Failure			
	(1)	(2)	(3)	(4)
Hiring Burn Rate	−0.176***	−0.207***	−0.199***	−0.207***
	(0.033)	(0.041)	(0.041)	(0.041)
Runway Risk PreRound	0.094***	0.124***	0.123***	0.136***
	(0.022)	(0.027)	(0.027)	(0.027)
Ln(Founder Prior-Work Exp.)			0.002	0.004
			(0.008)	(0.008)
Founder Prestige Score			−0.084***	−0.069***
			(0.024)	(0.025)
Serial Founder			-0.016	-0.018
			(0.018)	(0.018)
Ln(Employee Prior-Work Exp.)				0.003
				(0.009)
Employee Prestige Score				−0.071**
				(0.029)
Employee Prior Big-Firm Share				−0.047**
				(0.023)
Industry Hotness		-0.020	-0.020	-0.019
		(0.028)	(0.028)	(0.028)
Ln(Experienced Investors)		−0.112***	−0.109***	−0.110***
		(0.015)	(0.015)	(0.015)
Ln(Investor Syndicate)		-0.010	-0.008	-0.006
		(0.013)	(0.013)	(0.013)
Ln(Startup Age)	0.036***	0.009	0.001	0.006
	(0.008)	(0.012)	(0.012)	(0.012)
Industry x Emp Size Q x Fund Size Q FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Startup controls	Y	Y	Y	Y
Current investor controls	N	Y	Y	Y
Founder controls	N	N	Y	Y
Employee controls	N	N	N	Y
N	8,090	5,613	5,613	5,613
Adj. R^2	0.124	0.148	0.150	0.153

Table V. Long-Term Impact of Hiring Burn Rate on Employment, Product, and Patent Outcomes

This table reports ordinary least squares (OLS) estimates of the relation between HBR and long-term employment, product, and patent outcomes. Columns (1) and (2) report Δ Ln(Employment) over the three- and five-year periods following the HBR window. Columns (3) and (4) report Ln(Products), and columns (5) and (6) report Ln(Patents), over the same periods. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. All specifications control for pre-round runway risk, startup age, focal-round investor characteristics, pre-round founder characteristics, and pre-round employee characteristics. All specifications include industry-by-employment-size-quartile-by-fund-size-quartile fixed effects, financing-year fixed effects, and state fixed effects. Each observation is a financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Standard errors are clustered at the company level and reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	Δ Ln Employment		Ln Products		Ln Patents	
	$t+3$ (1)	$t+5$ (2)	$t+3$ (3)	$t+5$ (4)	$t+3$ (5)	$t+5$ (6)
Hiring Burn Rate	0.284*** (0.068)	0.231** (0.099)	0.153** (0.070)	0.248*** (0.095)	0.322* (0.170)	0.431** (0.208)
Industry x Emp Size Q x Fund Size Q FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y
Startup controls	Y	Y	Y	Y	Y	Y
Current investor controls	Y	Y	Y	Y	Y	Y
Founder controls	Y	Y	Y	Y	Y	Y
Employee controls	Y	Y	Y	Y	Y	Y
N	12,424	8,836	4,258	3,373	2,924	2,368
Adj. R^2	0.073	0.070	0.085	0.109	0.171	0.191

Table VI. Financing Continuation and Firm Survival by New Product and Patent Activity

This table reports fixed-effects multinomial logit estimates of the relation between HBR and the investment-liquidity trade-off. Panel A decomposes next-round financing outcomes. The omitted category is no subsequent financing; the two reported categories distinguish follow-on financing with and without new product or patent activity between the focal and next financing rounds. Panel B decomposes firm-failure outcomes. The omitted category is no failure (a successful exit); the two reported categories distinguish failures with and without new product or patent activity between the focal round and the observed failure event. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. Column pairs (1)-(2), (3)-(4), (5)-(6), and (7)-(8) correspond to four specifications. The first pair controls for pre-round runway risk and startup age; the remaining pairs sequentially add focal-round investor characteristics, pre-round founder characteristics, and pre-round employee characteristics. All specifications include industry-by-employment-size-quartile-by-fund-size-quartile fixed effects, financing-year fixed effects, and state fixed effects. Each observation is a financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Cells report multinomial-logit coefficients, average marginal effects in brackets, and standard errors in parentheses. The Diff rows report significance stars from tests of equality between the paired HBR coefficients. Standard errors are clustered at the company level. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	w new prod/pat	w/o new prod/pat	w new prod/pat	w/o new prod/pat	w new prod/pat	w/o new prod/pat	w new prod/pat	w/o new prod/pat
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Next Round								
Hiring Burn Rate	1.212***	1.790***	1.547***	2.069***	1.578***	2.073***	1.849***	2.154***
	[0.016]	[0.365***]	[0.027]	[0.397***]	[0.027]	[0.395***]	[0.048]	[0.395***]
	(0.331)	(0.192)	(0.401)	(0.246)	(0.417)	(0.246)	(0.428)	(0.248)
Diff: w = w/o, HBR		***		***		***		***
N		17,644		13,044		13,044		13,044
Pseudo R^2		0.046		0.054		0.063		0.068
Panel B: Firm Failure								
Hiring Burn Rate	−0.526*	−0.944***	−0.805**	−1.028***	−0.763*	−0.973***	−0.745*	−1.134***
	[0.003]	[−0.176***]	[−0.030]	[−0.174***]	[−0.029]	[−0.162***]	[−0.018]	[−0.195***]
	(0.316)	(0.228)	(0.385)	(0.284)	(0.390)	(0.286)	(0.404)	(0.289)
Diff: w = w/o, HBR		**		*				**
N		8,181		5,707		5,707		5,707
Pseudo R^2		0.056		0.076		0.086		0.094
Industry × Emp Size Q × Fund Size Q FE		Y		Y		Y		Y
Year FE		Y		Y		Y		Y
State FE		Y		Y		Y		Y
Baseline controls		Y		Y		Y		Y
Current investor controls		N		Y		Y		Y
Prior product/patent controls		N		N		N		N
Founder controls		N		N		Y		Y
Employee controls		N		N		N		Y

Table XI. Double Machine Learning Second Stage

This table reports second-stage estimates from double machine learning (DML). Columns (1)-(5) report Next Round, Δ Ln(Valuation), Firm Failure, Δ Ln(Employment) over the three-year period following the HBR window, and Ln(Products) over the same period, respectively. In each panel, the out-of-fold residualized outcome is regressed on residualized HBR without an intercept. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. M1 residualizes HBR and each outcome using startup characteristics. M2 adds pre-round founder characteristics, and M3 adds pre-round employee characteristics. Each observation is a financing round. Standard errors are clustered at the company level and reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	Next Round	Δ Ln(Valuation)	Firm failure	Δ Ln(Employment)	Ln(Product)
				$t+3$	$t+3$
	(1)	(2)	(3)	(4)	(5)
Panel A: M1 - Startup ML controls					
Residualized Hiring Burn Rate	0.484***	1.290***	−0.220***	0.102	0.227**
	(0.040)	(0.162)	(0.050)	(0.069)	(0.092)
N	13,065	4,271	4,435	12,390	3,708
R^2	0.012	0.024	0.004	0.000	0.002
Panel B: M2 - Startup + founder ML controls					
Residualized Hiring Burn Rate	0.444***	1.337***	−0.201***	0.050	0.239**
	(0.041)	(0.171)	(0.050)	(0.071)	(0.098)
N	13,065	4,271	4,435	12,390	3,708
R^2	0.010	0.024	0.003	0.000	0.002
Panel C: M3 - Startup + founder + pre-round employee ML controls					
Residualized Hiring Burn Rate	0.435***	1.307***	−0.194***	0.040	0.218**
	(0.041)	(0.175)	(0.050)	(0.070)	(0.097)
N	13,065	4,271	4,435	12,390	3,708
R^2	0.010	0.024	0.003	0.000	0.001
No intercept	Y	Y	Y	Y	Y

Appendix A: Derivations and Proofs

This appendix collects the derivations and proofs for the conceptual framework of Section 2 and for the measurement results of Section 3. Throughout, V_θ denotes the entrepreneur's gross expected receipts and $W_\theta \equiv V_\theta - A$ her net gain, $c(\kappa) = [1 + \int_0^\kappa \rho f]/F(\kappa)$ is the unit cost of effective investment, and the survival probabilities are $S_G(\kappa) = F(\kappa)$ and $S_B(\kappa) = \phi F(\kappa - \Delta\rho_0) + (1-\phi)F(\kappa)$.

12

A.1 Closed forms, break-even, and prefunding equivalence

Competitive investors who attribute the contract to type G break even in expected resource terms,

$$F(\kappa)\, p_H (R-b)\, I \;=\; (I - A) \;+\; \Big(\int_0^\kappa \rho f(\rho)\, d\rho\Big) I, \tag{16}$$

where the right side is the date-0 outlay plus the expected liquidity injections actually consumed.

6

Proof of Lemma 1. Write $N(\kappa) \equiv 1 + \int_0^\kappa \rho f$, so $c = N/F$, and recall $\int_0^\kappa \rho f = \kappa F(\kappa) - \int_0^\kappa F$. Dividing (16) by I gives $F(\kappa)p_\theta(R-b) = N(\kappa) - A/I$ with $p_\theta \in \{p_H, q_H\}$. The gross value $V_\theta = F(\kappa)\, p_\theta\, b\, I$ is increasing in both b and I; substituting the break-even condition, $V_\theta = F(\kappa)[\rho_1^\theta - c(\kappa)]I + A$, which is increasing in I when $\rho_1^\theta > c(\kappa)$, so scale is pushed to the cap imposed by $b \geq B/\Delta p$, that is, $p_\theta(R-b) \leq \rho_0^\theta$. At the cap $b = B/\Delta p$ and $I = A/[N - F\rho_0^\theta] = k_\theta(\kappa)A$, which gives the per-unit stake and multiplier. For monotonicity,

$D_\theta \equiv 1/k_\theta = N - F\rho_0^\theta$ has $D'_\theta(\kappa) = f(\kappa)(\kappa - \rho_0^\theta) > 0$ on the stated domain. The value expressions follow by substitution, using $1 + \frac{\rho_1 - c}{c - \rho_0} = \frac{\rho_1 - \rho_0}{c - \rho_0}$. Finally $c' = (f/F^2)[\kappa F - N] = (f/F^2)[\int_0^\kappa F - 1]$. Given we assume that $\int_0^{\bar\rho} F > 1$, the bracket is strictly increasing with a unique interior zero ρ^*, and $c(\rho^*) = \rho^*$ follows from $N(\rho^*) = \rho^* F(\rho^*)$.

Remark 2. On the frontier the per-unit stake equals the agency floor, so a revealed type-B holder of a type-G contract has spot capacity exactly $\ell + q_H(R - B/\Delta p) = \ell + \rho_0^B$, as used in S_B in (3). The entrepreneur is compensated through scale and not through slack in the stake.

Remark 3. Suppose the facility is prefunded, so that investors wire $(I - A) + \ell I$ at date 0, the buffer ℓI sits in riskless cash, shocks $\rho \le \ell$ are met from the buffer, shocks $\rho \in (\ell, \kappa]$ exhaust it and draw spot funds $(\rho - \ell)I$, and on liquidation, when $\rho > \kappa$, or at date 2 any unused cash reverts to investors. Expected investor receipts are

$$F(\kappa)p_H(R-b)I - \int_\ell^\kappa (\rho-\ell) f\, d\rho\, I + \int_0^\ell (\ell-\rho) f\, d\rho\, I + (1-F(\kappa))\,\ell I = F(\kappa)p_H(R-b)I + \ell I - \int_0^\kappa \rho f\, d\rho\, I,$$

using $\int_0^\ell (\ell - \rho)f - \int_\ell^\kappa (\rho - \ell)f = \ell F(\kappa) - \int_0^\kappa \rho f$. Setting receipts equal to the gross outlay $(I - A) + \ell I$, the ℓI terms cancel and (16) reappears. A prefunded cash buffer and a committed credit line are therefore break-even-equivalent at a zero interest rate. What differs is the gross date-0 raise, Raise $= (I - A) + \ell_0 I$, relative to the economic cost of liquidity, $\int_0^\kappa \rho f\, d\rho \cdot I$ in expectation. The empirical funding variable is the gross raise under the prefunded implementation, as used in Section 3.2.

12

A.2 The symmetric-information benchmark

6

Proof of Lemma 2. $W_\theta(\kappa)$ is a strictly decreasing transformation of $c(\kappa)$ on $\{c > \rho_0^\theta\}$, minimized at ρ^*, which depends only on F; positivity of ℓ_θ^*, finiteness of k_θ^*, and $W_\theta^* > 0$ follow from creditworthiness, $\rho_1^\theta > \rho^*$. The interior optimum requires no binding side constraint, since it is the minimizer of c. In regime (E) the inequality for the bad type follows from $\mathbb{E}[\mathbf{1}_\mathcal{C}(\rho_1^B - \rho)] \leq \int_0^{\rho_1^B}(\rho_1^B - \rho)f = \int_0^{\rho_1^B} F$ together with $\rho_1^B < \rho^*$; in regime (F), $W_B^* > 0$ because $\rho_1^B > \rho^*$, and $\ell_G^* - \ell_B^* = -\Delta\rho_0$. The multiplier ranking $k_G^* > k_B^*$ is immediate from the multiplier at $\kappa = \rho^*$.

12

A.3 The cost of not signaling

The following lemma supplies the signaling motive: separation rescues the good type from exclusion or cross-subsidization.

8

Lemma 4. *Consider any pooling allocation, defined by a common frontier contract held by both types, each type's gross value at least its outside option, and investors breaking even on the pool. In regime (E) there is $\alpha^* \in (0,1)$ such that no pooling allocation exists for $\alpha < \alpha^*$, so the market breaks down because the bad type cannot be funded alone, and for $\alpha \geq \alpha^*$ the good type is cross-subsidized, $V_G^{\text{pool}} - A \leq W_G(C) + \frac{1-\alpha}{\alpha} W_B(C) < W_G^*$. In regime (F) pooling exists for all α, and whenever $W_B(C) \leq W_B^*$ the bad type's participation forces $\pi_B \leq 0$,*

pooled zero profit gives $\pi_G \geq 0$*, and* $V_G^{\text{pool}} - A \leq W_G(C) \leq W_G^*$ *at any contract dominated for each type by its truthful optimum.*

6

Proof. The identity $(V_\theta - A) + \pi_\theta = W_\theta(C)$ holds for each type executing the common contract at its true capacities and success probability. In regime (E), summing with weights $\alpha, 1-\alpha$ and imposing pooled zero profit and both participation constraints gives $\alpha W_G(C) + (1-\alpha)W_B(C) \geq 0$. Per unit of scale the left side is at most $[\alpha\bar{w}_G - (1-\alpha)\delta]I$ with $\bar{w}_G \equiv \int_0^{\rho_1^G} F - 1 > 0$ and $\delta \equiv 1 - \int_0^{\rho_1^B} F > 0$, which is negative for $\alpha < \alpha^* \equiv \delta/(\bar{w}_G + \delta)$; the cross-subsidy bound rearranges the weighted identity using $V_B \geq A$. In regime (F), the bad type's participation, since her truthful frontier contract is available as a bad-type-attributed deviation, gives $V_B^{\text{pool}} - A \geq W_B^*$, so $\pi_B = W_B(C) - (V_B^{\text{pool}} - A) \leq W_B(C) - W_B^* \leq 0$; pooled zero profit gives $\pi_G = -\frac{1-\alpha}{\alpha}\pi_B \geq 0$, and the identity yields $V_G^{\text{pool}} - A = W_G(C) - \pi_G \leq W_G(C)$. Investor receipts are not ranked pointwise across types, so we do not sign π_B for an arbitrary common contract; the hypothesis $W_B(C) \leq W_B^*$ excludes the configurations in which blended financing relaxes the bad type's stand-alone constraint.

12

A.4 Single crossing and zero transparency

6

Proof of Lemma 3. On the range $\kappa - \Delta\rho_0 \geq \rho_0^G - \Delta\rho_0 = \rho_0^B > 0$, so $F(\kappa - \Delta\rho_0) > 0$. Given we assume that F is strictly log-concave, the reverse hazard f/F is strictly decreasing, so

$\frac{d}{d\kappa}\ln[F(\kappa-\Delta\rho_0)/F(\kappa)] = f(\kappa-\Delta\rho_0)/F(\kappa-\Delta\rho_0) - f(\kappa)/F(\kappa) > 0$, and Ω in (5) inherits this property because $\phi > 0$.

6

Proof of Corollary 1. Setting $\phi = 0$ in (5) gives $\Omega \equiv q_H/p_H$. With no revelation both types face identical continuation events under any belief, so $V_B = (q_H/p_H)V_G$ pointwise, at the pooled contract and at any deviation under any belief-induced pricing. Hence for any deviation the bad type gains if and only if the good type gains, the improvement sets coincide, and the Cho and Kreps (1987) test never assigns the deviation to the good type alone, so pooling survives. Deterrence of any separating contract requires $(q_H/p_H)V_G(\hat{\kappa}) \leq V_B^{\text{out}}$, a restriction on the level $V_G(\hat{\kappa})$ only, which any value-reducing instrument satisfies identically.

12

A.5 Off-path deviations, equilibria, and refinement

Deviations are frontier contracts attributed to the bad type off path. A deviating good type holding a bad-type-attributed contract at cutoff κ' has value $\Psi(\ell')\, V_B^{BB}(\kappa')$, where

$$\Psi(\ell) \equiv \frac{p_H[\phi F(\ell+\rho_0^G) + (1-\phi)F(\ell+\rho_0^B)]}{q_H\, F(\ell+\rho_0^B)}, \qquad V_B^{BB}(\kappa') = \frac{\rho_1^B - \rho_0^B}{c(\ell'+\rho_0^B) - \rho_0^B}\, A.$$

Given we assume that F is log-concave, Ψ is maximized at $\ell = 0$, with $\bar{\Psi} \equiv \frac{p_H}{q_H}[1-\phi+\phi F(\rho_0^G)/F(\rho_0^B)]$, and $\sup_{\kappa'} V_B^{BB} = A(\rho_1^B - \rho_0^B)/(\rho^* - \rho_0^B)$, which equals V_B^* in regime (F) and lies below A in regime (E). The no-masquerade condition of Section 2.3 is exactly the

requirement that the good type prefers her separating contract to the most tempting such deviation.

6

Proof of Proposition 1. For the bad type, her on-path contract attains her constrained optimum among bad-type-attributed contracts; deviating to the good type's contract yields $\Omega(\hat{\kappa}_F)V_G(\hat{\kappa}_F) \leq V_B^*$ by construction of $\hat{\kappa}_F$; any other proposal is bad-type-attributed and weakly dominated. For the good type, among good-type-attributed cutoffs, V_G increasing on $(\rho_0^G, \rho^*]$ makes $\hat{\kappa}_F$ optimal within the deterrence set; taking the bad type's contract, which gives capacity $\ell_B^* + \rho_0^G = \rho^* + \Delta\rho_0$ on revelation, yields $\Psi(\ell_B^*)V_B^* \leq \bar{\Psi}V_B^* \leq V_G(\hat{\kappa}_F)$, where the first inequality holds because Ψ is decreasing in ℓ and the second is the no-masquerade condition. Investor optimality and on-path Bayes updating hold by construction. For the cross-section ordering, $k_G(\hat{\kappa}_F) > k_G(\rho^*) = k_G^*$ by Lemma 1 and $k_G^* > k_B^*$ by Lemma 2, and the buffer gap is arithmetic; renegotiation-proofness follows because rescue beyond pledgeable income is infeasible.

In the exclusion regime the bad type's outside option is A, the limit of V_B^* as $\rho_1^B \downarrow \rho^*$, so deterrence is $\Omega(\kappa)V_G(\kappa) \leq A$, equivalently $\Omega(\kappa)(\rho_1^G - \rho_0^G) \leq c(\kappa) - \rho_0^G$; define κ_{IC}^E as the largest admissible cutoff and $\hat{\kappa} = \min\{\rho^*, \kappa_{IC}^E\}$.

8

Proposition 3. *In regime (E), given the no-masquerade condition at $\hat{\kappa}$ and the deterrence constraint attainable, there is a perfect Bayesian equilibrium in which the good type proposes the frontier contract at $\hat{\kappa}$ and the bad type proposes nothing and consumes A. The good type is funded with $V_G(\hat{\kappa}) > A$, imitation is deterred, and when $\kappa_{IC}^E < \rho^*$ the good type strictly*

distorts, with a thinner buffer, higher liquidation probability $F(\hat{\kappa}) < F(\rho^*)$*, and a larger multiple* $k_G(\hat{\kappa}) > k_G^*$.

6

Proof. For the bad type, proposing nothing yields A, the equilibrium contract yields $\Omega(\hat{\kappa})V_G(\hat{\kappa}) \leq A$, and any other bad-type-attributed frontier contract yields at most $A(\rho_1^B - \rho_0^B)/(\rho^* - \rho_0^B) < A$ in regime (E). For the good type, proposing nothing yields $A < V_G(\hat{\kappa})$, where positivity follows from $c(\hat{\kappa}) \leq c(\rho_0^G) < \rho_1^G$ and the gains-from-trade assumption; V_G is increasing on $(\rho_0^G, \kappa_{IC}^E]$, so $\hat{\kappa}$ is best, and any bad-type-attributed deviation yields at most $\bar{\Psi}A(\rho_1^B - \rho_0^B)/(\rho^* - \rho_0^B) \leq V_G(\hat{\kappa})$ by the no-masquerade condition. The remaining claims follow as in the proof of Proposition 1.

6

Proof of Proposition 2. Under belief G, competitive acceptance prices any contract (I, ℓ, b) by (16), and by Lemma 1 the proposer's value at a fixed κ is maximized on the frontier, so it suffices to test frontier deviations. Take a separating equilibrium at $\kappa' < \hat{\kappa}_F$ and a frontier deviation at $\kappa \in (\kappa', \hat{\kappa}_F]$. For the bad type the best conceivable outcome, acceptance at good-type pricing, yields $\Omega(\kappa)V_G(\kappa)$, which is at most her equilibrium payoff because $\kappa \leq \kappa_{IC}^F$, so the deviation is equilibrium-dominated for her; for the good type it is not, since $V_G(\kappa) > V_G(\kappa')$. The criterion forces $\mu(G \mid \kappa) = 1$, investors accept, and the equilibrium at κ' unravels, leaving $\hat{\kappa}_F$, with the analogue $\hat{\kappa}$ in regime (E). For pooling with $\phi > 0$, the deviation at the least-cost cutoff is equilibrium-dominated for the bad type but strictly gainful for the good type, so pooling unravels. At $\phi = 0$ this argument fails by Corollary 1.

12

A.6 Comparative statics

8

Proposition 4. *Let the relevant threshold, κ_{IC}^{E} or κ_{IC}^{F}, be interior. (i) Transparency: Ω is pointwise decreasing in ϕ while the right sides of the incentive constraints do not depend on ϕ, so the threshold rises with ϕ, and the distortion $\rho^{*} - \hat{\kappa}$, the excess liquidation, and the excess scale all decrease in ϕ, vanishing in the Exercise 6.10 limit. (ii) Quality gap: in regime (E) a lower q_H lowers Ω and the distortion falls; in regime (F) it also lowers V_B^{*}, so the net effect is ambiguous. (iii) Size invariance: neither the incentive constraints nor the no-masquerade condition contains A. (iv) Adverse selection: the value of separating falls in α, and in regime (E) below α^{*} separation is the only funded outcome. (v) Shock riskiness: a mean-preserving spread of F lowers ρ^{*}.*

6

Proof. For (i), $\partial\Omega/\partial\phi = -(q_H/p_H)(1 - F(\kappa - \Delta\rho_0)/F(\kappa)) < 0$; the thresholds are crossings of an increasing left side with a constant or decreasing right side, so a pointwise lower left side raises them, and the displayed objects are monotone transformations from Lemma 1. In the Exercise 6.10 limit, $S_B = 0$ at $\hat{\ell} = 0$ while $S_G = 1$, so the constraint is slack at the benchmark. For (ii), in regime (E) a lower q_H lowers the prefactor q_H/p_H and raises $\Delta\rho_0$, both of which lower Ω; in regime (F), $V_B^{*} = A\, q_H(B/\Delta p)/(\rho^{*} - \rho_0^{B})$ also falls, which tightens (IC-F), and the two forces oppose. Part (iii) is inspection. Part (iv) is Lemma 4. For (v), a mean-preserving spread raises $\int_0^x F$ pointwise, which forces ρ^{*} down in $\int_0^{\rho^{*}} F = 1$.

12

A.7 Two channels: pledgeability versus shock exposure

The baseline is a pledgeability-difference model, in which both types draw from the same F and the bad type's disadvantage is what she can pledge if revealed. The alternative is a shock-exposure model with common pledgeability, in which the bad type draws from a worse distribution F_B that is first-order stochastically dominated by F_G, for example a stretch $F_B(\rho) = F_G(\rho/\sigma)$ with $\sigma > 1$. In the shock-exposure model single crossing obtains even at $\phi = 0$, because under any common contract the imitator's survival is $F_B(\kappa) < F_G(\kappa)$, so $\Omega^{\text{shock}}(\kappa) = q_H F_B(\kappa)/[p_H F_G(\kappa)]$ varies with κ through the distribution ratio, and a Spence-type cost argument delivers separation with no interim information. The predictions also differ in sign: the thin good firm faces severe shocks less often, so it continues more often than the fat bad firm, which means high burn predicts lower liquidity-driven distress and there is no excess cash-out failure among good firms. The two channels therefore diverge at the two tests that most cleanly separate them. The no-information result at zero transparency holds under pledgeability with interim information and fails under shock exposure, and the shift in the composition of distress holds under the former and reverses under the latter. We keep the pledgeability channel as the baseline because it is native to Holmström and Tirole (2000) and Exercise 6.10 of Tirole (2006), because it makes the role of transparency explicit, and because it is the more falsifiable of the two. The channels are not mutually exclusive, and the composition-of-distress and transparency-split tests apportion the weight between them.

Appendix B Table I: Additional Summary Statistics

This table presents additional summary statistics for net changes during the HBR window and the share of longer-horizon net change realized during that window. Panel A reports changes in monthly salary spending and employee count from M-6 to M+6. Panel B reports the share of the 24- and 60-month net changes realized during the HBR window. The 24-month horizon runs from M-6 to M+18, and the 60-month horizon runs from M-6 to M+54. If a subsequent financing occurs after M+6 but before the indicated endpoint, the longer horizon ends in the month immediately preceding that financing. Observations with a subsequent financing on or before M+6 do not have valid HBR-window measures. Panel C reports the 24-month salary-spending shares by employee function and seniority. Ratios may exceed 100 percent if part of the initial expansion is subsequently reversed. Each observation is a financing round. Salary-spending changes are expressed in millions of dollars.

Variable	Obs	Mean	Median	Std. Dev.	P25	P75
Panel A. Net Changes During the HBR Window						
Salary Spending	17,583	0.074	0.044	0.098	0.012	0.098
Employee Count	17,583	7.908	5	11.032	1	10
Panel B. Share of Longer-Horizon Net Change Realized During the HBR Window						
24-Month Horizon (%)						
Salary Spending	15,514	77.020	67.520	77.736	41.689	98.560
Employee Count	15,269	71.212	65.217	59.102	38.710	100
60-Month Horizon (%)						
Salary Spending	13,394	74.719	57.023	105.822	28.106	90.057
Employee Count	13,222	69.391	54.545	84.157	25.641	90
Panel C. Share of Net Salary-Spending Change Realized During the HBR Window: 24-Month Horizon (%)						
By Employee Function						
Sales	15,229	67.382	59.212	83.068	0	100
Product	15,283	57.038	46.886	73.663	0	96.871
R&D	16,367	17.365	0	36.730	0	0
Operations	15,582	49.207	35.627	59.995	0	100
By Employee Seniority						
Leadership	15,539	35.218	0	50.645	0	100
Management	15,418	62.118	58.993	65.573	0	100
Senior	15,204	56.144	44.405	73.328	0	99.751
Junior	15,055	64.073	47.831	95.348	0	98.799

Appendix B Table II: Financing Round Distribution by Year

This table reports the number of startups and financing rounds by the indicated grouping. Next Round is the share of observations followed by a subsequent financing round. Firm Failure, IPO, and M&A rates are calculated among observations with nonmissing outcome information.

Year	Startups N	Financings N	Next Round %	Firm Failure %	IPO %	M&A %
2010	307	307	74.267	49.603	2.280	29.642
2011	844	844	68.957	53.177	2.133	30.332
2012	1,029	1,029	66.375	56.769	2.235	26.531
2013	1,148	1,148	64.460	60.884	2.091	23.345
2014	1,190	1,190	63.445	62.139	1.597	22.017
2015	1,198	1,198	59.683	69.362	1.252	15.860
2016	1,147	1,147	59.372	69.041	1.918	13.514
2017	1,297	1,297	63.685	68.655	1.619	11.951
2018	1,420	1,420	61.338	74.574	1.056	9.085
2019	1,562	1,562	63.124	80.060	0.448	6.530
2020	1,592	1,592	65.704	82.994	0.440	4.334
2021	2,283	2,283	59.571			
2022	2,196	2,196	46.585			
2023	370	370	52.703			
Pooled	17,583	17,583	60.803	66.663	1.012	11.153

Appendix B Table III: Financing Round Distribution by Industry

This table reports the number of startups and financing rounds by the indicated grouping, together with the share of the baseline financing-round sample represented by each grouping. Next Round is the share of observations followed by a subsequent financing round. Firm Failure, IPO, and M&A rates are calculated among observations with nonmissing outcome information.

Industry	Startups N	Financings N	Sample Share %	Next Round %	Firm Failure %	IPO %	M&A %
Software	7,166	7,166	40.755	63.871	61.397	0.237	13.508
Commercial Services	1,574	1,574	8.952	55.083	67.075	0.254	13.342
Pharmaceuticals and Biotechnology	1,099	1,099	6.250	65.332	50.324	9.918	7.552
Consumer Non-Durables	934	934	5.312	55.032	78.398	0.107	7.495
Healthcare Technology Systems	781	781	4.442	64.020	68.831	0.640	6.530
Healthcare Devices and Supplies	681	681	3.873	61.233	71.918	2.056	8.370
Media	640	640	3.640	51.875	72.589	0.156	14.531
Healthcare Services	637	637	3.623	59.498	79.574	0.628	4.710
Other Financial Services	630	630	3.583	60.635	68.182	0.476	6.984
Services (Non-Financial)	546	546	3.105	57.143	70.196	0	12.088
Commercial Products	355	355	2.019	55.493	78.528	0.563	5.352
Consumer Durables	308	308	1.752	60.390	79.058	0.325	9.740
Computer Hardware	292	292	1.661	66.096	64.925	1.370	11.986
Restaurants, Hotels and Leisure	280	280	1.592	52.500	76.970	0.357	11.071
Apparel and Accessories	225	225	1.280	52.889	83.721	0	8.889
Retail	201	201	1.143	54.229	77.500	1.493	9.453
IT Services	200	200	1.137	64.500	52.381	0.500	21.500
Communications and Networking	139	139	0.791	66.906	63.768	0.719	15.108
Transportation	126	126	0.717	56.349	70.968	0	11.905
Capital Markets/Institutions	118	118	0.671	47.458	81.034	1.695	5.085
Insurance	114	114	0.648	77.193	63.636	3.509	4.386
Energy Equipment	79	79	0.449	50.633	88.571	0	5.063
Semiconductors	71	71	0.404	71.831	62.069	0	12.676
Energy Services	70	70	0.398	51.429	78.125	0	8.571
Agriculture	68	68	0.387	55.882	84.848	0	4.412
Chemicals and Gases	55	55	0.313	56.364	66.667	0	14.545
Other Business Products and Services	37	37	0.210	56.757	70.588	0	10.811
Commercial Transportation	34	34	0.193	61.765	61.538	0	5.882
Commercial Banks	32	32	0.182	50	50	0	18.750
Exploration, Production and Refining	28	28	0.159	60.714	84.615	3.571	3.571
Utilities	19	19	0.108	52.632	84.615	0	0
Containers and Packaging	14	14	0.080	42.857	100	0	0
Other Materials	11	11	0.063	54.545	80	0	9.091
Metals, Minerals and Mining	8	8	0.045	75	50	0	12.500
Textiles	5	5	0.028	40	100	0	0
Construction (Non-Wood)	2	2	0.011	50		0	0
Other Healthcare	2	2	0.011	50		0	0
Other Information Technology	2	2	0.011	100		0	0
Pooled	17,583	17,583	100	60.803	66.663	1.012	11.153
Tech (pooled)	12,679	12,679	72.109	62.678	63.004	1.254	11.917
SaaS (pooled, PB)	5,774	5,774	32.839	66.505	62.133	0.260	11.153
SaaS (pooled, FinBERT)	7,852	7,852	44.657	65.563	60.923	0.318	11.793
AI (pooled, PB)	3,418	3,418	19.439	66.062	70.640	0.234	5.734
AI (pooled, FinBERT)	2,406	2,406	13.684	65.794	64.557	0.540	7.689

Notes: Tech (pooled) includes Software, Commercial Services, Pharmaceuticals and Biotechnology, Healthcare Technology Systems, Healthcare Devices and Supplies, Healthcare Services, Computer Hardware, IT Services, Semiconductors, Other Business Products and Services, Other Information Technology, and Communications and Networking. PitchBook SaaS and AI indicators are based on the corresponding PitchBook Verticals tags; observations without reported Verticals remain unclassified. FinBERT indicators are based on the text classifier. Pooled-row sample shares use the full baseline financing-round sample as the denominator.

Appendix B Table IV: Cox Proportional Hazards Models

This table reports hazard ratios from Cox proportional hazards models for the timing of the next financing round and firm failure. Columns (1) and (2) model the hazard of securing the next financing round, and columns (3) and (4) model the hazard of firm failure. Within each outcome, the first specification controls for pre-round runway risk and startup age, and the second adds focal-round investor characteristics, pre-round founder characteristics, and pre-round employee characteristics. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. The next-round model uses a continuous financing clock with entry at the focal-round start time, while the firm-failure model follows time to failure or censoring. All models are stratified by industry-by-employment-size-quartile-by-fund-size-quartile, financing year, and state. Each observation is a financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Standard errors are clustered at the company level and reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	Next Round		Firm Failure	
	(1)	(2)	(3)	(4)
Hiring Burn Rate	2.531***	2.867***	0.450***	0.468***
	(0.292)	(0.422)	(0.069)	(0.087)
Runway Risk PreRound	0.566***	0.438***	1.657***	1.888***
	(0.058)	(0.059)	(0.176)	(0.257)
Ln(Founder Prior-Work Exp.)		1.070***		0.952
		(0.027)		(0.032)
Founder Prestige Score		1.218***		0.931
		(0.089)		(0.094)
Serial Founder		0.917*		1.002
		(0.045)		(0.072)
Ln(Employee Prior-Work Exp.)		1.013		0.886***
		(0.026)		(0.031)
Employee Prestige Score		1.313***		0.692***
		(0.108)		(0.079)
Employee Prior Big-Firm Share		1.032		1.125
		(0.064)		(0.097)
Industry Hotness		0.498***		0.854
		(0.071)		(0.202)
Ln(Experienced Investors)		1.276***		0.885**
		(0.051)		(0.051)
Ln(Investor Syndicate)		1.013		0.911*
		(0.038)		(0.050)
Ln(Startup Age)	0.762***	0.819***	1.464***	1.419***
	(0.021)	(0.029)	(0.054)	(0.072)
Industry × Emp Size Q × Fund Size Q strata	Y	Y	Y	Y
Year strata	Y	Y	Y	Y
State strata	Y	Y	Y	Y
Startup controls	Y	Y	Y	Y
Current investor controls	N	Y	N	Y
Founder controls	N	Y	N	Y
Employee controls	N	Y	N	Y
N	17,644	13,044	17,644	13,044
Failures	9,719	7,343	5,450	3,479

Appendix B Table V: Employee-Count versus Salary-spending Hiring

This table compares hiring measured by employee-count changes with hiring measured by salary-spending changes. Columns (1)-(6) report Next Round, Δ Ln(Valuation), Firm Failure, Δ Ln(Employment) over the three-year period following the HBR window, Ln(Products) over the same period, and Ln(Patents) over the same period, respectively. Employee-Count HBR measures the change in employee count per million dollars raised between six months pre- and post-financing. HBR measures the corresponding change in monthly salary spending. The first block jointly includes Employee-Count HBR and HBR in a separate regression for each outcome. The residualized HBR row reports the Frisch-Waugh-Lovell-equivalent coefficient obtained after residualizing HBR on Employee-Count HBR, the full set of controls, and the fixed effects within each outcome sample. All specifications control for pre-round runway risk, startup age, focal-round investor characteristics, pre-round founder characteristics, and pre-round employee characteristics. All specifications include industry-by-employment-size-quartile-by-fund-size-quartile fixed effects, financing-year fixed effects, and state fixed effects. Each observation is a financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Cells report coefficient estimates, one-standard-deviation effects in brackets, and standard errors in parentheses. Standard errors are clustered at the company level. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	Next Round	Δ Ln(Valuation)	Firm failure	Δ Ln(Employment)	Ln(Product)	Ln(Patent)
				$t+3$	$t+3$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
Hiring-intensity channels						
Employee-count HBR	-0.001	-0.001	-0.000	−0.005***	-0.001	0.002
	[−0.014]	[−0.013]	[−0.004]	[−0.092]	[−0.015]	[0.031]
	(0.001)	(0.003)	(0.001)	(0.001)	(0.001)	(0.003)
Salary-spending HBR	0.420***	1.343***	−0.187*	0.816***	0.261*	0.088
	[0.064]	[0.133]	[−0.031]	[0.124]	[0.031]	[0.011]
	(0.076)	(0.319)	(0.097)	(0.152)	(0.138)	(0.348)
Residualized salary-intensity check						
Residualized Salary-spending HBR	0.420***	1.343***	−0.187*	0.816***	0.261*	0.088
	[0.023]	[0.050]	[−0.011]	[0.045]	[0.011]	[0.003]
	(0.076)	(0.319)	(0.097)	(0.152)	(0.138)	(0.348)
Industry × Emp Size Q × Fund Size Q FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y
Baseline controls	Y	Y	Y	Y	Y	Y
Current investor controls	Y	Y	Y	Y	Y	Y
Founder controls	Y	Y	Y	Y	Y	Y
Employee controls	Y	Y	Y	Y	Y	Y
N	12,962	4,693	5,613	12,424	4,258	2,924
Adj. R^2	0.096	0.122	0.153	0.074	0.084	0.171

Appendix B Table VI: Hiring Burn Rate by Employee Function, Seniority, and Wage

This table decomposes salary-based HBR by employee function, employee seniority, and employee wage. Panel A reports HBR by employee function, Panel B reports HBR by employee seniority, and Panel C reports HBR by employee wage quartile. Each HBR component measures the change in monthly salary spending per million dollars raised for the corresponding employee group between six months pre- and post-financing. Columns (1)-(6) report Next Round, Δ Ln(Valuation), Firm Failure, Δ Ln(Employment) over the three-year period following the HBR window, Ln(Products) over the same period, and Ln(Patents) over the same period, respectively. All specifications control for pre-round runway risk, startup age, focal-round investor characteristics, pre-round founder characteristics, and pre-round employee characteristics. All specifications include industry-by-employment-size-quartile-by-fund-size-quartile fixed effects, financing-year fixed effects, and state fixed effects. For brevity, the corresponding control and fixed-effect indicator rows are omitted. Each observation is a financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Cells report coefficient estimates, one-standard-deviation effects in brackets, and standard errors in parentheses. Bracketed effects are the primary quantities for comparing magnitudes across employee groups. Standard errors are clustered at the company level. * p < 0.10; ** p < 0.05; *** p < 0.01.

	Next Round	Δ Ln(Valuation)	Firm failure	Δ Ln(Employment)	Ln(Product)	Ln(Patent)
				$t+3$	$t+3$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Function HBR						
Sales	0.352***	0.893***	−0.149*	0.109	0.151	0.409
	[0.026]	[0.042]	[−0.012]	[0.008]	[0.008]	[0.023]
	(0.074)	(0.321)	(0.086)	(0.140)	(0.151)	(0.354)
Product	0.900***	2.350***	−0.710***	0.346	0.758***	0.875**
	[0.040]	[0.076]	[−0.035]	[0.015]	[0.029]	[0.035]
	(0.117)	(0.436)	(0.141)	(0.235)	(0.226)	(0.444)
R&D	1.424***	2.784**	-1.092	2.736***	-0.554	3.350**
	[0.015]	[0.027]	[−0.010]	[0.028]	[−0.005]	[0.035]
	(0.457)	(1.365)	(0.693)	(0.883)	(0.691)	(1.341)
Operations	0.270*	1.575**	-0.126	0.882***	-0.228	0.038
	[0.009]	[0.036]	[−0.004]	[0.030]	[−0.006]	[0.001]
	(0.156)	(0.621)	(0.183)	(0.285)	(0.270)	(0.728)
Panel B: Seniority HBR						
Leadership	0.457***	1.258**	-0.051	0.300	-0.368	0.089
	[0.015]	[0.028]	[−0.002]	[0.010]	[−0.009]	[0.003]
	(0.152)	(0.575)	(0.181)	(0.292)	(0.242)	(0.569)
Management	0.430***	0.904**	−0.283***	1.015***	0.405**	0.039
	[0.027]	[0.037]	[−0.019]	[0.064]	[0.020]	[0.002]
	(0.086)	(0.371)	(0.103)	(0.159)	(0.164)	(0.315)
Senior	0.890***	3.503***	−0.695***	0.267	0.403	2.105***
	[0.027]	[0.079]	[−0.023]	[0.008]	[0.010]	[0.053]
	(0.174)	(0.713)	(0.213)	(0.346)	(0.317)	(0.799)
Junior	0.551***	1.073*	−0.379**	−0.996***	0.192	0.534
	[0.019]	[0.026]	[−0.014]	[−0.034]	[0.005]	[0.013]
	(0.159)	(0.622)	(0.187)	(0.314)	(0.304)	(0.647)
Panel C: Wage-Quartile HBR						
Employee wage Q4	0.477***	2.092***	−0.178**	1.087***	0.148	0.334
	[0.036]	[0.108]	[−0.015]	[0.082]	[0.009]	[0.023]
	(0.070)	(0.283)	(0.084)	(0.135)	(0.129)	(0.281)
Employee wage Q3	0.382***	0.094	−0.312**	−0.378*	0.309	0.671
	[0.019]	[0.003]	[−0.017]	[−0.019]	[0.011]	[0.025]
	(0.112)	(0.457)	(0.127)	(0.215)	(0.236)	(0.467)
Employee wage Q2	0.478***	2.307***	−0.581***	−0.855**	0.324	-0.057
	[0.014]	[0.043]	[−0.019]	[−0.025]	[0.007]	[−0.001]
	(0.185)	(0.768)	(0.206)	(0.356)	(0.372)	(0.834)
Employee wage Q1	0.553	1.491	-0.553	0.153	-1.036	0.769
	[0.007]	[0.012]	[−0.007]	[0.002]	[−0.010]	[0.007]
	(0.398)	(1.616)	(0.496)	(0.828)	(0.760)	(1.755)
N	12,962	4,693	5,613	12,424	4,258	2,924
Panel A Adj. R^2	0.101	0.126	0.156	0.074	0.086	0.173
Panel B Adj. R^2	0.101	0.126	0.155	0.075	0.086	0.172
Panel C Adj. R^2	0.099	0.127	0.154	0.076	0.084	0.170

Appendix B Table VII. DML Out-of-Sample Fit and Nested-Increment Tests

This table reports first-stage out-of-sample fit for the double machine learning (DML) procedure and tests the incremental predictive contribution of additional feature layers. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. Columns (1)-(3) report out-of-sample R-squared, or Efron pseudo-R-squared for binary outcomes, from company-purged rolling walk-forward folds. M1 uses startup characteristics, M2 adds pre-round founder characteristics, and M3 adds pre-round employee characteristics. Columns (4) and (5) report the incremental change in R-squared from M1 to M2 and from M2 to M3, respectively. The line below each increment reports the company-clustered Diebold-Mariano t-statistic in parentheses. The HBR row reports the first-stage nuisance fit for the treatment, and the remaining rows report the outcome nuisance fits. Each observation is a financing round. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	(1) M1 Startup	(2) M2 + Founder	(3) M3 + Employee	(4) ΔR^2 Founder (DM t) (M2 − M1)	(5) ΔR^2 Employee (DM t) (M3 − M2)
Hiring Burn Rate (treatment)	0.408	0.444	0.449	+0.036***	+0.005
				(3.86)	(1.28)
Next Round	0.084	0.095	0.101	+0.011***	+0.006***
				(4.58)	(4.11)
Δ Ln(Valuation)	0.034	0.053	0.054	+0.019***	+0.002
				(3.69)	(0.53)
Firm Failure	0.068	0.086	0.095	+0.018***	+0.009***
				(4.10)	(3.26)
Δ Ln(Employment)$_{t,t+3}$	0.122	0.133	0.147	+0.011***	+0.013***
				(4.22)	(5.62)
Ln(Product)$_{t,t+3}$	0.050	0.064	0.060	+0.015***	−0.004
				(3.25)	(-1.14)

Appendix B Table VIII. Out-of-Sample Performance of Machine Learning versus Linear Models

This table compares out-of-sample performance for predicting HBR across machine-learning models and a linear-model benchmark. HBR measures the change in monthly salary spending per million dollars raised between six months pre- and post-financing. Each observation is a financing round, and performance is evaluated using company-purged rolling walk-forward folds. Panels A-C use nested feature sets consisting of startup characteristics, startup plus pre-round founder characteristics, and startup plus pre-round founder and pre-round employee characteristics, respectively. Columns report out-of-sample R-squared, root mean squared error (RMSE), and the percentage RMSE differential relative to the linear model; negative differentials indicate lower RMSE than the benchmark. The Ensemble row reports the nonnegative-least-squares-weighted combination of the six member learners (XGBoost, Random Forest, LightGBM, Histogram-based Gradient Boosting, CatBoost, and Extra Trees) used in the DML first stage. Stars indicate company-clustered Diebold-Mariano tests of squared-error improvement over the linear model. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	R^2_{OS} (1)	RMSE (2)	RMSE Differential (%) (3)
Panel A: Hiring Burn Rate (Startup Variables)			
XGBoost	0.375	0.112	−7.446***
RF	0.411	0.108	−10.104***
LightGBM	0.354	0.114	−5.860***
HistGB	0.387	0.111	−8.300***
CatBoost	0.402	0.109	−9.486***
ExtraTrees	0.392	0.110	−8.667***
LASSO	0.282	0.120	−0.790***
LM	0.271	0.121	0.000
Ensemble (6-Learner NNLS)	0.408	0.109	−9.907***
Panel B: Hiring Burn Rate (Startup-Founder Variables)			
XGBoost	0.422	0.107	−10.884***
RF	0.433	0.106	−11.718***
LightGBM	0.411	0.108	−10.018***
HistGB	0.423	0.107	−10.933***
CatBoost	0.444	0.105	−12.538***
ExtraTrees	0.421	0.107	−10.819***
LASSO	0.293	0.119	−1.425***
LM	0.273	0.120	0.000
Ensemble (6-Learner NNLS)	0.444	0.105	−12.552***
Panel C: Hiring Burn Rate (Startup-Founder-Employee Variables)			
XGBoost	0.427	0.107	−11.947***
RF	0.434	0.106	−12.539***
LightGBM	0.414	0.108	−11.005***
HistGB	0.434	0.106	−12.468***
CatBoost	0.449	0.105	−13.710***
ExtraTrees	0.420	0.108	−11.462***
LASSO	0.293	0.119	−2.186***
LM	0.261	0.121	0.000
Ensemble (6-Learner NNLS)	0.449	0.105	−13.650***

Appendix B Table IX: Alternative Hiring Burn Rate Measures

This table reports ordinary least squares (OLS) estimates for alternative Hiring Burn Rate (HBR) measures. Salary-spending HBR measures the change in monthly salary spending per million dollars raised; employee-count HBR measures the change in employee count per million dollars raised. Columns (1)-(6) report next-round financing, valuation growth, firm failure, the change in ln employment over the three-year period following the corresponding HBR window, Ln(Product), and Ln(Patent) over the same period, respectively. All specifications include industry-by-employment-size-quartile-by-fund-size-quartile fixed effects, financing-year fixed effects, and state fixed effects. Each observation is a financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Standard errors are clustered at the company level and reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	Next Round	Δ Ln(Valuation)	Firm Failure	Δ Ln(Employment)	Ln(Product)	Ln(Patent)
				$t+3$	$t+3$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
Salary-Spending $HBR_{[M-3,M+3]}$	0.402***	1.565***	−0.309***	0.273***	0.155*	0.055
	(0.050)	(0.228)	(0.061)	(0.098)	(0.089)	(0.220)
N	14,750	5,424	6,391	14,588	4,979	3,410
Adj. R^2	0.090	0.118	0.159	0.078	0.088	0.154
Employee-Count $HBR_{[M-6,M+6]}$	0.002***	0.009***	−0.002***	0.001**	0.001	0.003*
	(0.000)	(0.001)	(0.000)	(0.001)	(0.001)	(0.002)
N	12,962	4,693	5,613	12,424	4,258	2,924
Adj. R^2	0.094	0.116	0.153	0.072	0.084	0.171
Employee-Count $HBR_{[M-3,M+3]}$	0.003***	0.012***	−0.002***	0.001	0.001	-0.000
	(0.000)	(0.002)	(0.000)	(0.001)	(0.001)	(0.002)
N	14,750	5,424	6,391	14,588	4,979	3,410
Adj. R^2	0.088	0.115	0.158	0.077	0.087	0.154
Industry/size/funding-size FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y
Baseline controls	Y	Y	Y	Y	Y	Y
Current investor controls	Y	Y	Y	Y	Y	Y
Founder controls	Y	Y	Y	Y	Y	Y
Employee controls	Y	Y	Y	Y	Y	Y

Appendix B Table X: Matched-Sample Hiring Burn Rate Estimates

This table reports ordinary least squares (OLS) estimates of the salary-spending Hiring Burn Rate (HBR) in a matched sample. Each focal startup's first Seed or Early Stage VC financing round is matched to the first financing round of up to five startups at the same stage. Candidate peers must be in the same industry, be financed within five years of the focal round, and have firm age between 50% and 150% of the focal startup's age. Pre-round employment and deal-size ranges begin at 50%–150% of the focal values and widen as needed, up to 250%. We further require business-description cosine similarity above the 80th percentile of the pooled same-industry pair distribution among U.S. venture-backed startups whose first financing is a Seed or Early Stage VC round between 2010 and 2026. Each focal startup uses the narrowest range containing at least five qualifying peers, or the widest range if none does. We then keep up to five peers, ranked by similarity. Focal and matched firms are stacked into matched sets. Columns (1)-(6) report next-round financing, valuation growth, firm failure, the change in ln employment over the three-year period following the HBR window, Ln(Product), and Ln(Patent) over the same period, respectively. All columns include matched-set, financing-year, and state fixed effects and the full baseline, current investor, founder, and employee control set. Each observation is a financing round. Continuous variables are winsorized at the 1st and 99th percentiles. Standard errors are clustered at the company level and reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. The bottom block reports balance diagnostics for the matching variables, computed across the 63,109 focal-matched pairs (15,231 unique focal firms, 14,578 unique matched firms).

Panel A: Matched-Sample Estimates

	Next Round	Δ Ln(Valuation)	Firm Failure	Δ Ln(Employment)	Ln(Product)	Ln(Patent)
				$t+3$	$t+3$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
Salary-Spending $HBR_{[M-6,M+6]}$	0.611***	1.531***	−0.345***	0.761***	0.434***	0.982***
	(0.059)	(0.224)	(0.071)	(0.110)	(0.147)	(0.282)
Matched-set FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y
Baseline controls	Y	Y	Y	Y	Y	Y
Current investor controls	Y	Y	Y	Y	Y	Y
Founder controls	Y	Y	Y	Y	Y	Y
Employee controls	Y	Y	Y	Y	Y	Y
N	59,180	19,829	21,290	57,091	16,680	9,542
Adj. R^2	0.113	0.132	0.181	0.127	0.131	0.122

Panel B: Balance Diagnostics for the Matching Variables

	Focal Mean	Matched Mean	Difference	t-statistic	Normalized Difference
Employee Size PreRound	8.47	7.71	0.761	33.82	0.066
Deal Size ($M)	4.00	3.71	0.296	18.14	0.031
Firm Age	1.91	1.61	0.308	78.93	0.150
Financing Year	2017.63	2017.57	0.058	5.30	0.017

Appendix B Table XI: Founder Human Capital

This table reports summary statistics for 28,607 founder-deal observations attached to startups in the final DML second-stage sample. Founders are identified from focal-company Founder, Co-founder, CEO, or President positions that begin within two years of company founding. Education records must begin by $M-7$; prior-career variables use information available by $M-7$. Selected categories are reported.

Variable	Obs	Share (%)
Panel A: Demographics		
Gender		
Female	28,607	17.223
Ethnicity		
White	28,602	72.089
Asian or Pacific Islander	28,602	15.943
Panel B: Education		
Undergraduate institution		
Ivy-8 Degree	13,415	7.201
Ivy-Plus Degree	13,415	5.963
International Top Undergrad	13,415	1.752
Graduate education		
Global Top MBA	20,932	6.015
Other MBA	20,932	9.029
Advanced Degree	20,932	13.338
Undergraduate major		
Business Major	12,144	38.192
STEM Major	12,144	58.333
Panel C: Prior Career		
Prior employer industry (any position)		
Technology	28,607	51.466
Consulting	28,607	14.080
Investor Fund	28,607	7.103
Investment Banking	28,607	7.159
Prior employer type (any position)		
Worked at VC-Backed Startups	28,607	36.662
— Worked at VC-Backed Startups with Successful Exit	28,607	16.265
— Worked at VC-Backed Startups with Observed Failure	28,607	6.824
Worked at Young Firms (Non-VC-Backed)	28,607	41.647
Worked at Public Firms	28,607	28.122
Worked at Private Firms (Emp $\geq$ 500, Non-VC-Backed)	28,607	62.376
Worked at Other Firms	28,607	69.885
Prior entrepreneurial experience		
Prior Founder Experience at VC-Backed Startups	28,607	10.487
— Successful Exit	28,607	4.258
— Observed Failure	28,607	1.790

Appendix B Table XII: Employee Human Capital

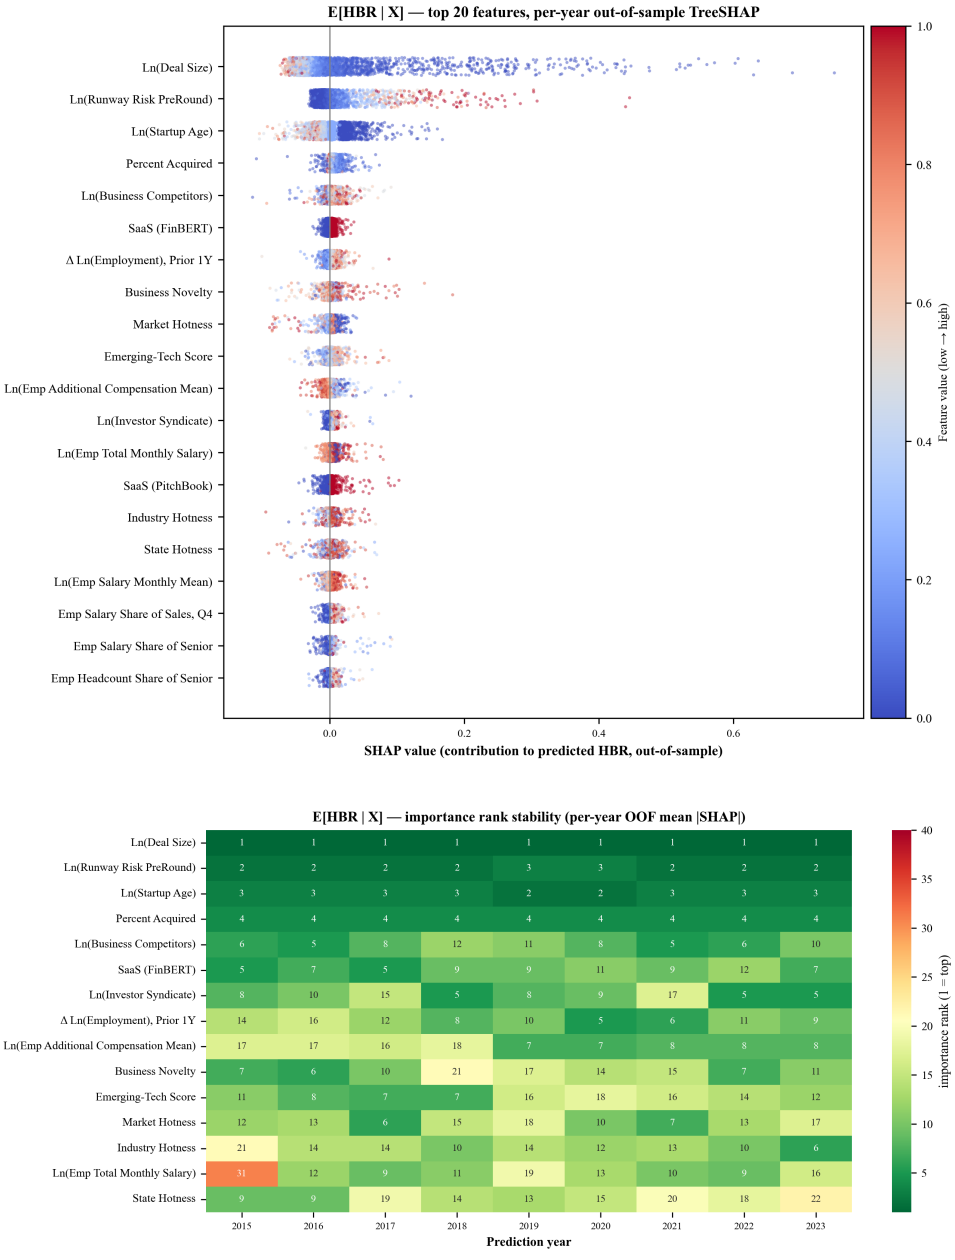
This table reports summary statistics for 131,844 employee-deal observations covering non-founder employees active at $M-7$ in startups in the final DML second-stage sample. Education records must begin by $M-7$; prior-career variables use information available by $M-7$. Selected categories are reported.

Variable	Obs	Share (%)
Panel A: Demographics		
Gender		
Female	131,844	36.801
Ethnicity		
White	131,811	65.056
Asian or Pacific Islander	131,811	17.842
Panel B: Education		
Undergraduate institution		
Ivy-8 Degree	58,875	2.402
Ivy-Plus Degree	58,875	1.841
International Top Undergrad	58,875	1.060
Graduate education		
Global Top MBA	94,618	1.549
Other MBA	94,618	5.515
Advanced Degree	94,618	4.707
Undergraduate major		
Business Major	47,743	41.225
STEM Major	47,743	52.785
Panel C: Prior Career		
Prior employer industry (any position)		
Technology	131,844	34.793
Consulting	131,844	8.549
Investor Fund	131,844	2.388
Investment Banking	131,844	2.989
Prior employer type (any position)		
Worked at VC-Backed Startups	131,844	19.284
— Worked at VC-Backed Startups with Successful Exit	131,844	7.140
— Worked at VC-Backed Startups with Observed Failure	131,844	3.699
Worked at Young Firms (Non-VC-Backed)	131,844	33.038
Worked at Public Firms	131,844	18.348
Worked at Private Firms (Emp $\geq$ 500, Non-VC-Backed)	131,844	48.729
Worked at Other Firms	131,844	64.074
Prior entrepreneurial experience		
Prior Founder Experience at VC-Backed Startups	131,844	1.298
— Successful Exit	131,844	0.389
— Observed Failure	131,844	0.264

Appendix B Figure I. HBR First-Stage Feature Importance

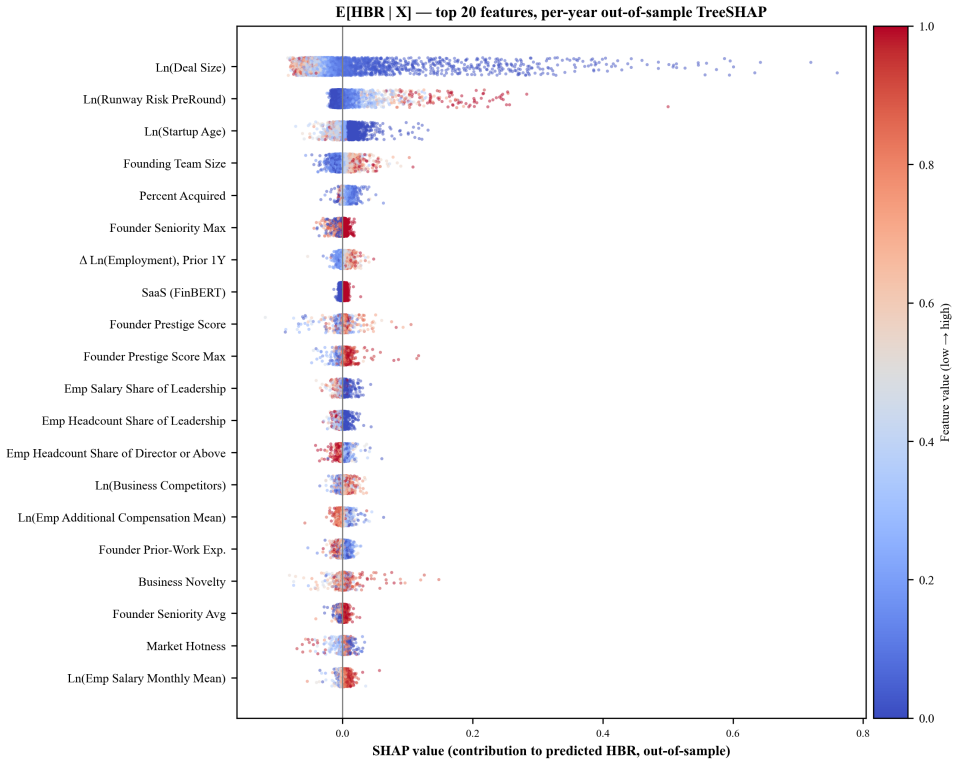
This figure reports first-stage HBR feature importance across three nested DML models using out-of-fold TreeSHAP: M1 includes 151 startup characteristics, M2 adds 66 founder characteristics, and M3 further adds 61 employee characteristics. The top panel reports the 20 features with the highest mean absolute SHAP values. Each point represents one financing round's SHAP value for a given feature; positive (negative) values indicate that the feature's value for that round raises (lowers) its predicted HBR, and color indicates the feature value (blue = low, red = high), so the color-position pattern—rather than the sign alone—reveals the direction of each feature's association. The bottom panel reports feature-importance ranks by prediction year for the 15 features with the best average rank, where 1 denotes the most important feature. For each prediction year, SHAP values are obtained from a model trained on the preceding five years, excluding all firms appearing in the prediction-year sample, following the walk-forward folds in Appendix B, Table VII. Spearman rank correlations over common features are 0.98 for M1–M2 and 0.98 for M2–M3, indicating that adding founder and employee characteristics does not materially change the relative importance of features already included in the simpler models.

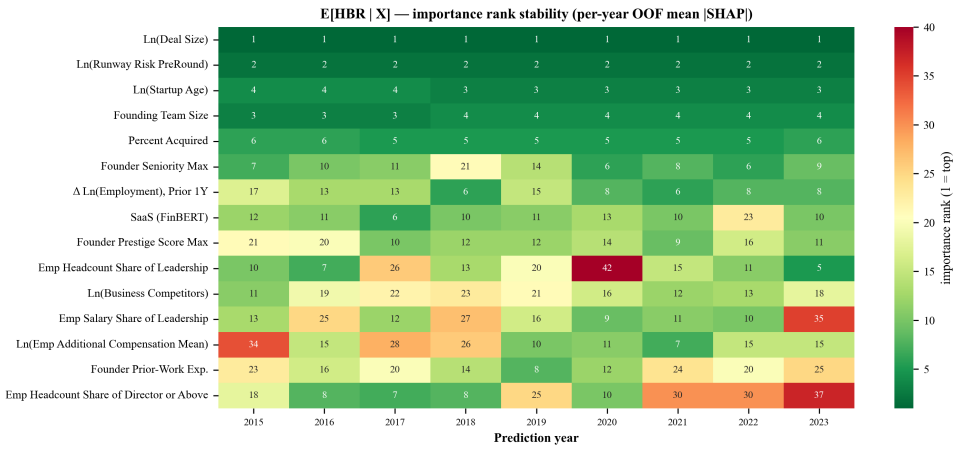
(a) M1 (Startup)



Appendix B Figure I (continued).

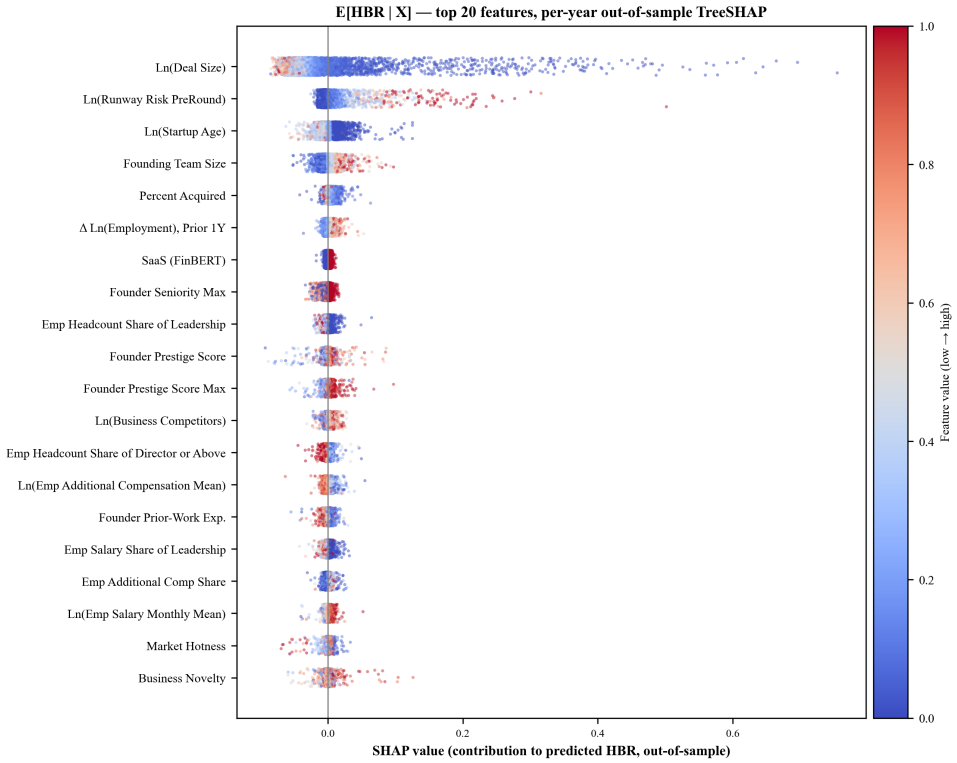
(b) M2 (+ Founder)

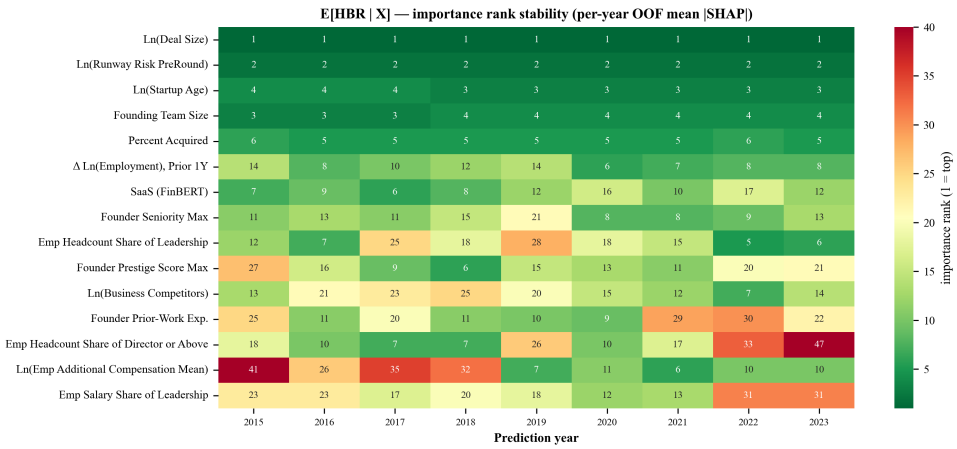




Appendix B Figure I (continued).

(c) M3 (+ Employee)





Appendix C: Variable Definition

Variable	Definition
Next Round	One if the startup receives follow-on VC financing within three years or, after its final observed VC round, has a successful exit or receives PE Growth/Expansion financing; zero otherwise.
Δ Ln(Valuation)	Log pre-money valuation of the subsequent financing round minus log post-money valuation of the focal round.
Firm Failure	One if the startup is classified as failed and zero if it has a successful exit.
Δ Ln(Employment)$_{t,t+3}$	Change in log employee count from M+7 to M+43.
Δ Ln(Employment)$_{t,t+5}$	Change in log employee count from M+7 to M+67.
Ln(Product)$_{t,t+3}$	Log number of observed product launches during the three years after the HBR window.
Ln(Product)$_{t,t+5}$	Log number of observed product launches during the five years after the HBR window.
Ln(Patent)$_{t,t+3}$	Log number of matched patent filings during the three years after the HBR window.
Ln(Patent)$_{t,t+5}$	Log number of matched patent filings during the five years after the HBR window.
Hiring Burn Rate	Change in monthly salary spending from M-6 to M+6 divided by the focal amount raised.
Hiring Burn Rate [M-3,M+3]	Change in monthly salary spending from M-3 to M+3 divided by the focal amount raised.
HBR [M-6,M+6], Employee-count	Change in employee count from M-6 to M+6 divided by the focal amount raised.
HBR [M-3,M+3], Employee-count	Change in employee count from M-3 to M+3 divided by the focal amount raised.
Runway Risk PreRound	Monthly salary spending at M-7 divided by the focal amount raised; a higher value indicates less salary-only coverage before financing.
Ln(Founder Prior-Work Exp.)	Log mean years of prior work experience among founders before the focal financing round.
Founder Prestige Score	Mean Revelio education-and-network prestige score among active founders at M-7. Revelio initializes the score using world-university rankings and updates it through worker-mobility networks.
Serial Founder	Share of founders with prior founder experience.
Ln(Employee Prior-Work Exp.)	Log mean years of prior work experience among active non-founder employees at M-7.
Employee Prestige Score	Mean Revelio education-and-network prestige score among active non-founder employees at M-7. Revelio initializes the score using world-university rankings and updates it through worker-mobility networks.
Employee Prior Big-Firm Share	Share of active non-founder employees at M-7 whose prior employer had at least 500 employees.
Industry Hotness	Log number of U.S. startups receiving a first Seed or Early Stage VC financing in the same industry during the preceding two quarters.
Ln(Experienced Investors)	Log number of focal-round investors with at least five prior VC investments in the same industry during the preceding two years.
Ln(Investor Syndicate)	Log number of investors in the focal-round syndicate.
Ln(Startup Age)	Log startup age at the focal financing round.
Employee Count PreRound	Number of active non-founder employees at M-7.
Amount Raised	Amount raised in the focal financing round, in millions of U.S. dollars.

Continued on next page

Variable	Definition
Percent Acquired	Share of startup ownership sold in the focal financing round.
SaaS (%)	Percentage of startups classified as SaaS by PitchBook.
AI (%)	Percentage of startups classified as AI by PitchBook.
Ln(Product), Prior 3Y	Log number of observed product launches during the three years before the HBR window.
Ln(Patent), Prior 3Y	Log number of matched patent filings during the three years before the HBR window.
Product Observed at Any Time	One if at least one matched product event is observed for the startup and zero otherwise.
Patent Observed at Any Time	One if at least one matched patent filing is observed for the startup and zero otherwise.
Sales	Employees in sales or marketing functions.
Product	Employees in engineering or product functions.
R&D	Employees in scientist or R&D functions.
Operations	Employees in administration, finance, or operations functions.
Leadership	Employees with Revelio seniority score 7.
Management	Employees with Revelio seniority scores 5–6.
Senior	Employees with Revelio seniority scores 3–4.
Junior	Employees with Revelio seniority scores 1–2.