# Venture Capital as Portfolios of Compound Options*

Sebastian Hillenbrand
Harvard Business School

Erik Stafford
Harvard Business School

July 2025

**Abstract**

A defining feature of startup financing is its staged structure: in each funding round, venture capital (VC) investors have the option to continue financing or to abandon a startup. By exercising the financing option, VC investors retain the option to eventually take the startup public, typically on the Nasdaq. We model startups as compound options on underlying firms that, upon successful exit, list on the Nasdaq. Our startup valuation model closely matches the IPO time series of startups, delivers closed-form expressions for startup betas over the lifecycle of these ventures, and yields a key insight: a portfolio of startups resembles a leveraged position in the Nasdaq. This replication strategy, implementable at low cost, tracks the performance of VC fund vintages with striking accuracy. Over the past two decades, VC has significantly underperformed this benchmark.

*We thank Taha Choukhmane, Joshua Coval, Peter Cziraki, Philipp Illeditsch, Tong Liu, Caio Machado, and Lawrence Schmidt as well as participants at several seminars and conferences for helpful comments and discussions. We thank Jennifer Beauregard for help with the data collection. Harvard Business School's Division of Research provided financial research support for both authors. Author emails: shillenbrand@hbs.edu (Hillenbrand), estafford@hbs.edu (Stafford).

# 1 Introduction

Large institutional investors, such as university endowments, allocate a substantial and growing share of their portfolios to venture capital (VC). VC offers the potential to fund the next unicorn that eventually lists on a public exchange, generating outsized returns for founders and investors. At the same time, most startups fail, resulting in a total loss of capital. To manage this downside risk, VC funds provide financing to startups in stages. This staged financing structure provides VC investors with the option to continue funding promising startups and discontinue those that underperform.[1] While the option-like nature and the skewed payoff of VC investments are often recognized, their implications for the risks and returns of VC investments remain imperfectly understood.

This paper develops a tractable compound-option framework (Geske, 1979) for valuing startup firms. Our model captures the staged nature of VC financing, reproduces the empirical pattern of a few blockbuster IPOs and mergers alongside many liquidations, and endogenously links exit payoffs to aggregate market conditions. Despite its simplicity, the framework closely matches two broad stylized facts: the time-series trajectory of VC-backed IPOs over the past thirty years and the return profiles of successive VC fund vintages. Remarkably, despite the optionality embedded within individual startups, a diversified portfolio of these staged-financing option claims behaves indistinguishably from a low-turnover leveraged position in the stock market index that drives their underlying systematic risk.

Building on prior studies of selection bias in deal-level returns (Cochrane, 2005; Korteweg and Sorensen, 2010) and the value of staging optionality (Berk, Green, and Naik, 2004), our framework provides clear, quantitative answers to several central questions facing both academics and practitioners. In particular, we demonstrate how high idiosyncratic risk coexists with significant sensitivity to market swings once firms successfully list; how the implicit leverage embedded in successive financing rounds shapes a startup's evolving market beta; the degree of convexity inherent in returns to well-diversified VC investors; which systematic risk factors best summarize the VC payoff distribution; and, finally, how investors can construct and implement a low-cost replication strategy—and hence a natural benchmark—for evaluating VC performance.

We model individual startup firms as compound call options. We treat each funding round as an embedded real option. At the first decision point, the startup invests only if the expected value of its remaining options exceeds the cost; otherwise it is discontinued/liquidated. If it proceeds, a second option grants the choice between a final growth investment (pursuing an IPO) or an exit via acquisition,

[1] See (Gompers, 1995; Lerner, 1994; Cumming and MacIntosh, 2003) for discussions and evidence of the role of staging in venture capital; and (see Raiffa and Schlaifer, 1961) for foundational treatments of staging and abandonment via decision trees.

with payoffs priced by Geske's compound-option formula early on (before the first decision point) and by Black–Scholes once only a single call remains (after the first decision point).

To analyze the systematic risks of VC, we integrate a CAPM-style market factor into our option framework. We drive firm asset values with both idiosyncratic shocks and a market factor, so that exit payoffs naturally depend on realized aggregate returns.[2] An important implication of this framework is that a startup can be viewed as a levered claim on its underlying asset value, with leverage naturally arising from the staged funding structure. If the underlying asset value is exposed to systematic risk, this exposure is amplified through the leverage inherent in staged funding, magnifying the systematic risk in startups. When startups are pooled into diversified funds, that amplified systematic risk is prevalent at both the fund and allocator-portfolio levels as idiosyncratic risk is diversified away.

To bring our model to the data, we establish several empirical facts. Using comprehensive deal-level data spanning four decades, we document several empirical patterns in funding activity. First, exit outcomes are highly skewed: 65% of startups are liquidated, 27% are acquired, and 8% go public via IPO. Second, among IPOs, we find that 80% of VC-backed firms list on the Nasdaq exchange, confirming the high-growth, tech-orientation of VC-backed firms and motivating a Nasdaq-based market factor (see Ritter (1991) for early evidence on IPO listings). Third, because an IPO represents the final stage, turning a private startup into a publicly traded security, it reveals the risk characteristics of the firm's underlying asset. Conditional on IPO, we observe an average volatility of 90% and a Nasdaq-100 beta of 1.4—evidence of substantial systematic risk exposure in the very asset our option-pricing model treats as the underlying. Finally, quarterly IPO volumes track Nasdaq-100 returns far more closely than S&P 500 returns, confirming that both exit timing and payoffs are strongly procyclical with tech-sector performance. These empirical regularities validate our use of the realized market return as the state variable in the option-pricing framework and guide the calibration of idiosyncratic volatility and beta parameters.

These empirical findings directly inform our model. They imply that the appropriate underlying asset—should have high idiosyncratic volatility around 90% and a beta of 1.4 with respect to the Nasdaq-100. We set just two calibration targets—the unconditional probabilities of startup liquidation (65%) and IPO (8%)—and fix all model parameters thereafter. We then ask the model to deliver two out-of-sample predictions over a thirty-year time series: (i) the quarterly pattern of VC-backed IPOs, and (ii) the vintage-level return profile of institutional portfolios. Remarkably, without any return-based fitting, the constant-parameter model explains 39% of the variation in IPO frequencies and 92% of the return

[2]The key idea is to use the realized market return as the relevant state space for asset pricing, consistent with the Sharpe (1964) and Lintner (1965) CAPM. This follows (Coval, Jurek, and Stafford, 2009) who integrate the CAPM into the Merton (1974) credit model to study the pricing of bond portfolios and CDO tranches.

variation of VC fund vintages—demonstrating that matching only unconditional exit rates suffices to capture rich conditional dynamics.

The model implies stock-market betas of 2.2, 1.6, and 1.4 for early-, late-, and mezzanine-stage startups, respectively. Consistent with the findings of Berk, Green, and Naik (2004), systematic risk declines as startups progress through funding rounds and implicit leverage decreases. We then compare the stage-level model payoffs of a portfolio of startups with the payoff of a leveraged position in the Nasdaq-100 with the same stage-level beta. We find that the two payoffs are nearly identical with plausible parameter sets. Key to this finding is that the idiosyncratic risk of startups is high, a feature emphasized in prior deal-level studies (Cochrane, 2005; Korteweg and Sorensen, 2010). In other words, the option-like payoff structure that exists at the individual startup level vanishes once we construct a diversified portfolio of startups. This implies a powerful result: once the stage-level VC beta is identified, VC portfolio returns can be replicated using a low-turnover leveraged position in the Nasdaq-100 index.

The model suggests that an appropriate public market benchmark (Kaplan and Schoar, 2005) for evaluating VC investments is a leveraged position in the Nasdaq-100. We implement a vintage-level replication strategy that mirrors actual VC limited-partner cash flows—both contributions and distributions—using a levered Nasdaq-100 portfolio. The initial portfolio leverage is calibrated to early-stage startup betas and then gradually decreases as the portfolio companies mature. This strategy, which can be implemented at low cost using Nasdaq-100 futures, replicates the performance of VC vintages well, explaining 90% of the variation in VC vintage returns—which is only slightly below the 92% that our option pricing model explains. Notably, the strategy is able to keep pace with the extremely high VC returns realized during the tech boom in the late 1990s.

From the perspective of the leveraged Nasdaq benchmark, we uncover a stark performance divergence around the tech boom supported by two key findings. First, pre-2000 VC vintages match or slightly outperform the levered Nasdaq benchmark, while every vintage since 2000 underperforms the benchmark. Second, we use the replicating strategy to build a calendar-time market-level VC index and compare its performance to the Cambridge Associates (CA) VC index. We find that the CA index underperforms the Nasdaq benchmark over the last two decades.

Against the backdrop of VC underperformance, it is surprising that institutional investors like endowments have increased their allocations to this asset class. One possible consequence of this surge in capital is that fundraising may have outpaced the availability of high-quality investment opportunities, potentially resulting in capital being allocated to weaker deals at inflated valuations and leading to lower subsequent returns (Gompers and Lerner, 2000; Ewens and Rhodes-Kropf, 2015). An alternative

possibility for the underperformance of VC relative to our benchmark is that the riskiness of VC has declined over the past two decades, as fewer firms go public and more are acquired by "big tech," but we find little empirical support for this.

Our results show that investors—even those without direct VC exposure or access—can cheaply replicate the risk–return profile of a diversified VC portfolio by applying stage-specific betas to leverage a Nasdaq-100 position. More generally, this paper illustrates how a parsimonious structural model—one that captures only staging optionality and market-beta risk—can deliver practical benchmarks for private-equity investing.

**Related Literature**

**Option-Based Valuation** Our framework builds on the Black–Scholes–Merton tradition for pricing standard options (Black and Scholes, 1973; Merton, 1973, 1974) and Geske's compound-option formula (Geske, 1979). Subsequent work applied these tools to multi-stage investments (Majd and Pindyck, 1987; Pindyck, 1993; Schwartz and Moon, 2000; Berk, Green, and Naik, 2004), but did not integrate systematic risk factors directly into the compound-option structure. In contrast, we embed a CAPM-style market factor to generate joint predictions for exit timing and portfolio returns (Cremers, Driessen, and Maenhout, 2008; Coval, Jurek, and Stafford, 2009).

**Risk and Return in Venture Capital** A large body of empirical literature documents the skewed, high-variance nature of VC returns (Kaplan and Sensoy, 2015; Korteweg, 2019). Deal-level studies confront severe selection bias—successful exits are overrepresented—requiring structural corrections (Cochrane, 2005; Korteweg and Sorensen, 2010)[3]. Fund-level studies face infrequent NAV updates and fee distortions (Gompers and Lerner, 1997; Phalippou and Gottschalg, 2009; Woodward, 2009; Ewens, Jones, and Rhodes-Kropf, 2013; Brown, Ghysels, and Gredil, 2023). Our structural model endogenizes selection through market-driven exit options, aligning our high systematic-risk estimates with both strands of the literature (Cochrane, 2005; Korteweg and Sorensen, 2010; Ang, Chen, Goetzmann, and Phalippou, 2018; Korteweg and Nagel, 2024).

**Lifecycle Betas** Empirical evidence on how systematic risk evolves through financing rounds is mixed: some find rising betas (Korteweg and Sorensen, 2010); others find declining betas (Cochrane, 2005). By modeling successive funding as nested options, our framework delivers a structural explanation for declining betas—leverage falls as projects survive early stages—consistent with Berk, Green,

[3]Deal-level analyses often abstract from seniority differences across share classes (Gornall and Strebulaev, 2020) and complex fund–entrepreneur contracts (Ewens, Gorbenko, and Korteweg, 2022).

and Naik (2004). Our framework yields startup betas that—due to their closed-form solutions—are inherently intuitive and transparent.

**Performance Benchmarking** The public-market-equivalent (PME) approach (Kaplan and Schoar, 2005; Korteweg and Nagel, 2016) remains the industry standard for evaluating private equity, despite its reliance on broad market indices. Recent work has refined this by introducing factor-tilted benchmarks (Franzoni, Nowak, and Phalippou, 2012; Gupta and Van Nieuwerburgh, 2021). We contribute by showing (1) the underlying asset risk is tied to the Nasdaq-100, and (2) that the conditional payoffs to startup portfolios (e.g., VC risk) are well approximated by a leveraged Nasdaq-100 exposure, implying that a low-turnover levered tech-sector index is a natural benchmark for VC performance.

## 2 Empirical Characteristics of Startups

In this section, we document several empirical features of startups backed by venture capitalists. We first examine the staged nature of VC funding and the eventual "exit" outcomes of VC-backed startups. Since the most successful startups typically list on public exchanges, we then analyze the risk characteristics of newly public, formerly VC-backed firms. Finally, we show that periods of elevated IPO activity coincide with strong prior performance in public equity markets.

### 2.1 Funding Rounds and Exits of VC-Backed Startups

It is well-documented that venture capitalists provide funding to startups in stages. VC funds follow this approach for several reasons (Gompers, 1995): staged financing (1) strengthens the entrepreneur's incentives to exert effort and (2) mitigates the risk of backing projects and technologies with highly uncertain prospects. By providing capital in rounds, venture capitalists retain the option to discontinue funding if a startup fails to meet key milestones. To study this behavior in the data, we use the VentureXpert deal-level and exits datasets, which we access via the LSEG platform.

In the deal-level data, we restrict attention to startups based in the United States and focus on venture capital rounds, which we identify as rounds where the "round type" is classified as "Seed", "Early Stage", "Later Stage", or "Expansion", or where the "round security" is labeled "Venture Capital Equity Investment". We further limit the sample to startups whose first funding round occurred between 1981 and 2016, yielding 38,796 startups and 112,937 funding rounds. To align the deal-level data with our two-stage model, we aggregate any follow-on VC financings occurring within six months into a single round. This prevents counting purely administrative "inside rounds" or "backstop financings"

as distinct milestone stages.[4] When the round size is missing—which is the case for 12.73% of observations—we assume a value of $1 million, corresponding to the 17th percentile of observed funding round amounts. We merge the deal-level data with the exit data, which contains information on exit type (e.g., IPO, mergers, write-offs) as well as the enterprise valuations for mergers. We classify a "High-Value Merger" as a merger or acquisition (M&A) in which the enterprise valuation at exit is at least five times the amount raised in the first VC round. "Low-Value Mergers" include M&As that do not meet this threshold and M&As where the enterprise value is unobserved, as well as buybacks, reverse takeovers (RTOs), and secondary sales. Finally, we classify an exit as a "Write Off" if the exit is explicitly recorded as a "Write Off" in VentureXpert or remains unobserved. The latter applies to 24,162 (62%) startups.

**Table 1: Funding Rounds and Exits of VC-Backed Startups**

| Frequency | | Count | Fraction (%) |
|---|---|---|---|
| | IPO | 3230 | 8.33 |
| | High-value mergers | 1744 | 4.50 |
| | Low-value mergers | 8742 | 22.53 |
| | Write-Off | 25080 | 64.65 |
| Number of VC rounds | | Mean | Median |
| | All outcomes | 1.9 | 1 |
| | IPO | 3.3 | 3 |
| | High-value mergers | 3.4 | 3 |
| | Low-value mergers | 2.8 | 2 |
| | Write-Off | 1.3 | 1 |
| Time to exit (years) | | Mean | Median |
| | IPO | 5.0 | 4.2 |
| | High-value mergers | 6.4 | 5.5 |
| | Low-value mergers | 5.9 | 5.1 |
| Funding round amount (2020 dollars) | | Mean | Median |
| | First round | 12.3 | 4.5 |
| | Second round | 21.3 | 8.8 |
| | Third round | 36.7 | 12.5 |
| Time between rounds (years) | | Mean | Median |
| | All outcomes | 1.6 | 1.2 |

**Note:** This table shows exit outcomes: number of VC rounds, time between the first funding round and the exit, funding amounts per funding round and the time between rounds. The sample includes all startups that received their first VC funding round between 1981 and 2016, as reported by VentureXpert.

Table 1 summarizes the funding and exit patterns of VC-backed startups. Only a small share of

[4]Inside rounds—where only incumbent investors participate—are documented to comprise roughly 30% of follow-ons and often serve merely as backstops when outside capital is unavailable (Ewens, Rhodes-Kropf, and Strebulaev, 2015). Similarly, Broughman and Fried (2012) find that most inside rounds reflect such administrative top-ups rather than true new investments (Broughman and Fried, 2012).

startups achieve successful exits through IPOs (8.3%) or high-value mergers (4.5%), while the majority (64.7%) are written off. Consistent with the notion of staged financing, more successful outcomes are associated with more funding rounds: startups that go public or exit via high-value mergers receive a median of three VC rounds, compared to only one for startups that are written off. Time to exit is relatively similar across exit types, with average and median durations between four and six years. We also compute the size of funding rounds, converted to 2020 dollars using the CPI. The average second-round funding amount is 21.3 million—roughly 1.73 times the average first-round amount (12.3 million); for medians, the ratio is 8.8 / 4.5 $\approx$ 1.96. The average third-round amount is 36.7 million, implying a ratio of 1.72 relative to the second round; for medians, the ratio is 12.5 / 8.8 $\approx$ 1.42. These funding patterns suggest that funding rounds increase over time, consistent with startups maturing and uncertainty about the underlying technology gradually resolving. Finally, the average time between funding rounds is 1.6 years, with a median of 1.2 years, further illustrating the staged nature and pacing of VC investments.

## 2.2 Stock Exchange Listing of VC-backed Startups

Since an IPO is the ultimate "high-payoff" exit in our two-stage framework, we next document where and how often VC-backed firms list on public markets. We use both VentureXpert (as described above) and Jay Ritter's U.S. IPO data to ensure robustness to any changes in coverage over time in either dataset.[5] As shown in Table 4, the time-series of VC-backed IPO activity is highly consistent across the two sources.

**Table 2: IPO Listings by Exchange**

| Fraction listing on exchange | VentureXpert | Ritter |
|---|---|---|
| AMEX (NYSE American) | 1.2 % | 1.5 % |
| NYSE | 8.7 % | 10.2 % |
| Nasdaq | 79.8 % | 88.2 % |
| Other | 10.2 % | |
| Count | 3230 | 3376 |

**Note:** This table reports the stock exchanges on which VC-backed startups eventually list their shares. The VentureXpert data include both U.S. and international listings, whereas the Jay Ritter dataset focuses exclusively on U.S. listings.

Table 2 reports the stock exchanges on which successful VC-backed startups list their shares. The

[5]The Ritter data is further described in Ritter (2015). The data focuses on U.S. IPOs after excluding those with an offer price below $5.00 per share, unit offers, small best efforts offers, American Depositary Receipts (ADRs), closed-end funds, natural resource limited partnerships, special purpose acquisition companies (SPACs), real estate investment trusts (REITs), bank and S&L IPOs, and firms not listed on the Center for Research in Security Prices (CRSP) returns files within six months of the IPO, thus restricting the sample to NYSE-, Nasdaq-, and Amex- (now NYSE MKT) listed stocks. The primary data source is the Thomson Reuters (also known as SDC (Securities Data Company)) new issues database.

distributions are highly consistent across the VentureXpert and Jay Ritter datasets, with the main difference being that VentureXpert includes a small number of non-U.S. listings. The table shows that the vast majority of VC-backed IPOs occur on Nasdaq: 79.8% in VentureXpert and 88.2% in Ritter's data. In contrast, listings on the NYSE and AMEX are relatively uncommon, accounting for less than 10% and 2% of IPOs, respectively. This pattern is consistent with Nasdaq's historical focus on younger, high-growth firms, particularly in the technology and biotech sectors.

Next, we examine the return dynamics of VC-backed firms following their IPOs. We obtain daily stock returns from CRSP and estimate each firm's beta with respect to the Nasdaq-100 and the S&P 500, as well as total return volatility. Beta estimates include 20 lags of market index returns to capture delayed responses and potential autocorrelation. Daily volatilities are annualized by multiplying by $\sqrt{252}$. We track these risk measures over the first five years of trading to assess how the risk profile of newly public firms evolves over time.

Table 3 provides an overview of the estimates. The Nasdaq-100 beta of newly listed VC-backed stocks is estimated to be 1.40 with a standard error of 0.084. The median across all stocks in the sample is 1.29. When benchmarking the returns against the S&P 500, one expects the beta to be higher since the Nasdaq-100 has a beta of 1.20 with respect to the S&P 500. In line with this, the S&P 500 beta of newly listed VC-backed stocks is estimated to be 1.81 with a standard error of 0.145; the median S&P 500 beta is 1.57. The annualized volatility is estimated to be 0.91.

**Table 3: Beta and Volatility of VC-backed newly IPO'd firms**

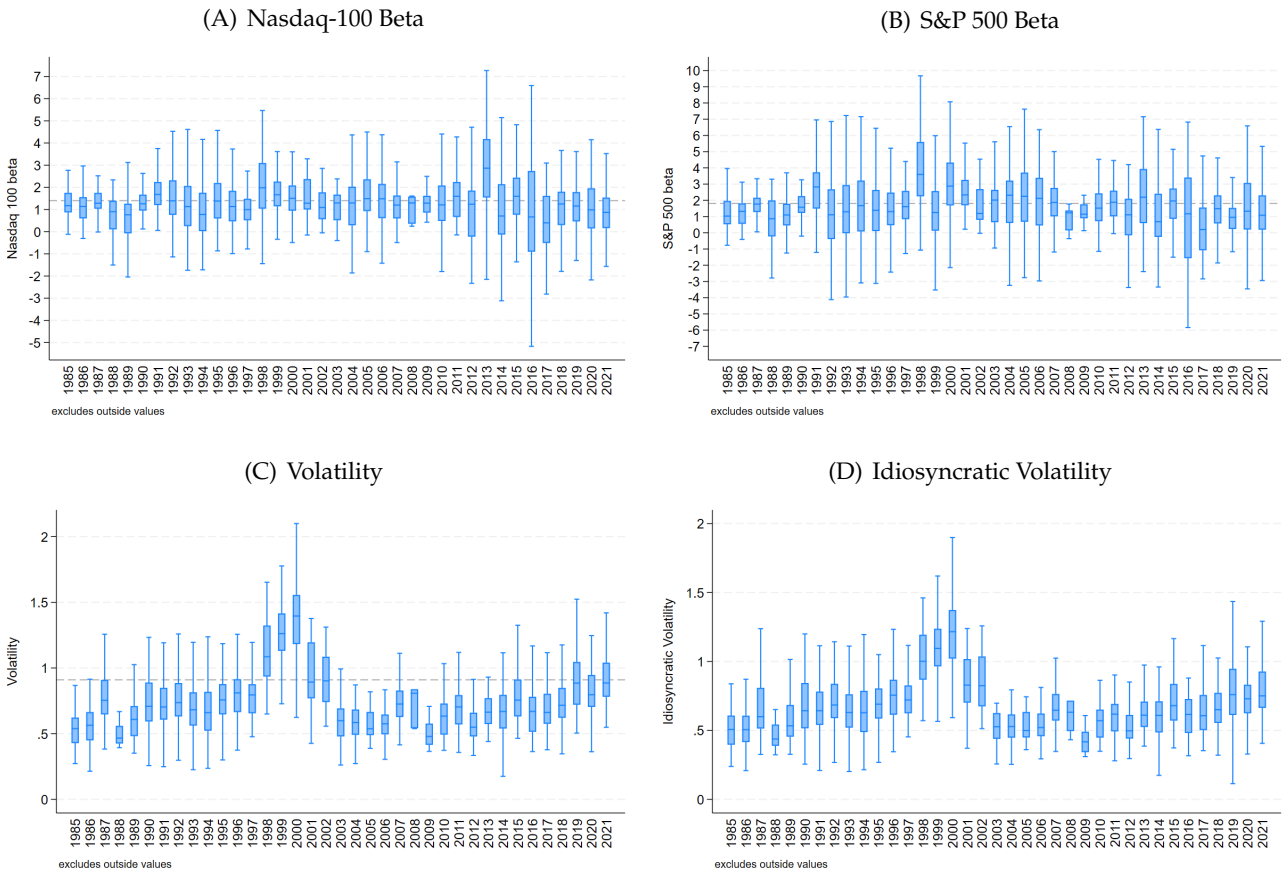
| | Variable | Mean | Standard error | p10 | p25 | p50 | p75 | p90 |
|---|---|---|---|---|---|---|---|---|
| First year after IPO (N = 3,325) | Nasdaq-100 Beta | 1.40 | 0.082 | -0.19 | 0.54 | 1.29 | 2.02 | 2.87 |
| | S&P 500 Beta | 1.81 | 0.145 | -0.67 | 0.43 | 1.57 | 2.82 | 4.28 |
| | Volatility | 0.91 | 0.008 | 0.49 | 0.62 | 0.77 | 1 | 1.33 |
| Second year after IPO (N=3,151) | Nasdaq-100 Beta | 1.30 | 0.062 | -0.19 | 0.47 | 1.17 | 1.92 | 2.7 |
| | S&P 500 Beta | 1.80 | 0.114 | -0.61 | 0.52 | 1.56 | 2.72 | 4.06 |
| | Volatility | 0.98 | 0.012 | 0.47 | 0.59 | 0.76 | 1.03 | 1.39 |
| Third year after IPO (N=2,760) | Nasdaq-100 Beta | 1.27 | 0.055 | -0.23 | 0.49 | 1.17 | 1.93 | 2.8 |
| | S&P 500 Beta | 1.78 | 0.098 | -0.48 | 0.58 | 1.56 | 2.72 | 4.05 |
| | Volatility | 1.01 | 0.020 | 0.45 | 0.58 | 0.76 | 1.02 | 1.34 |
| Fourth year after IPO (N=2,305) | Nasdaq-100 Beta | 1.25 | 0.055 | -0.2 | 0.45 | 1.16 | 1.88 | 2.77 |
| | S&P 500 Beta | 1.65 | 0.096 | -0.6 | 0.43 | 1.43 | 2.56 | 4.08 |
| | Volatility | 0.97 | 0.150 | 0.44 | 0.56 | 0.74 | 0.98 | 1.32 |
| Fifth year after IPO (N=2,024) | Nasdaq-100 Beta | 1.14 | 0.057 | -0.11 | 0.51 | 1.14 | 1.83 | 2.64 |
| | S&P 500 Beta | 1.59 | 0.095 | -0.39 | 0.63 | 1.56 | 2.66 | 4 |
| | Volatility | 0.96 | 0.190 | 0.44 | 0.56 | 0.73 | 0.99 | 1.32 |

**Note:** This table provides information on the return characteristics of newly-listed firms that were previously VC-backed. Daily return data from CRSP is used in the estimation.

When exploring the life cycle of the beta after the IPO, one finds that the beta is decreasing in subsequent years following the IPO. For instance, the Nasdaq-100 beta is 1.40 in the year following the IPO, decreasing to 1.30 in the second year. It then declines further to 1.27 and 1.25 in the third and fourth years, respectively, and reaches 1.14 by the fifth year after the IPO. A similar trend is observed for the S&P 500 beta, which decreases from 1.81 in the year following the IPO to 1.59 by the fifth year.

Figure 1 explores the time-series dimension by showing, for each calendar year, the distribution of return characteristics across all firms listed in that year during their first year of public trading. There is no obvious pattern in the stock betas (Panels A and B). In contrast, the volatility exhibits a more cyclical pattern (Panel C). The volatility of newly listed VC-backed stocks was the highest during the tech bubble, with a median value exceeding 1. The pattern for idiosyncratic volatility (Panel D) resembles the pattern of total volatility.

**Figure 1: Beta and Volatility of VC-backed newly IPO'd firms**




**Note:** This figure provides information on the return characteristics of newly-listed firms that were previously VC-backed, categorized by their IPO year. Daily return data from CRSP is used in the estimation.

### 2.3 The Time Series of VC-Backed IPOs

Next, we examine the time-series patterns of successful exits—IPOs and high-value mergers—and relate these to the number of startups that received venture funding. To mitigate concerns about coverage and selection bias—specifically, that successful outcomes are more likely to be observed in the data (Cochrane, 2005; Korteweg and Sorensen, 2010)—we draw on multiple data sources. The number of IPOs per quarter is obtained from both VentureXpert and Jay Ritter's dataset. Data on the number of startups receiving their first round of VC funding—i.e., startup vintages by year—are obtained from VentureXpert and Cambridge Associates. Because the Cambridge data report the number of funded startups only from 1996 onward, we approximate the earlier series by regressing the number of startups on the number of VC funds, which is available from Cambridge Associates beginning in 1981. The most limited data pertain to high-value mergers, for which we rely solely on VentureXpert. Notably, the dataset contains only a few mergers prior to 1992, suggesting limited or no coverage of mergers and acquisitions during the early portion of the sample.

Table 4 provides an overview of the data at an annual frequency. The number of startups receiving their first round of VC funding—i.e., startup vintages—ranged from 302 to 629 per year between 1981 and 1994. This number then increased steadily from 1995 to 2000, peaking at 3,521 in 2000. Following the burst of the tech bubble, startup formation declined sharply and has remained within a range of 900 to 1,700 firms per year since 2005. The number of IPOs exhibits a broadly similar pattern: VC-backed IPO activity was elevated between 1994 and 2000, with particularly high levels in 1996, 1999, and 2000, each exceeding 200 IPOs. Since the early 2000s, IPO activity has remained subdued, surpassing 100 offerings in only one year (2014). In contrast, the number of high-value mergers has shown a steady upward trend over time. This is consistent with prior evidence suggesting that many firms are better off remaining private due to the expansion and deregulation of private equity markets (Ewens and Farre-Mensa, 2020), as well as with the growing tendency of large public firms to acquire smaller companies in order to benefit from synergies in operations and innovation (Gao, Ritter, and Zhu, 2013; Bena and Li, 2014).

Using these data, we construct a quarterly measure of IPO probability defined as

$$\text{IPO Probability}_t = \frac{\#\text{IPOs}_t + \#\text{High-Value Mergers}_t}{\frac{1}{16}\sum_{\tau=t-23}^{t-8} \#\text{Startups}_\tau}, \tag{1}$$

where $\#\text{IPOs}_t$ and $\#\text{High-Value Mergers}_t$ denote the number of VC-backed IPOs and high-value mergers in quarter $t$, respectively, and $\#\text{Startups}_\tau$ is the number of VC-backed startups that received their first

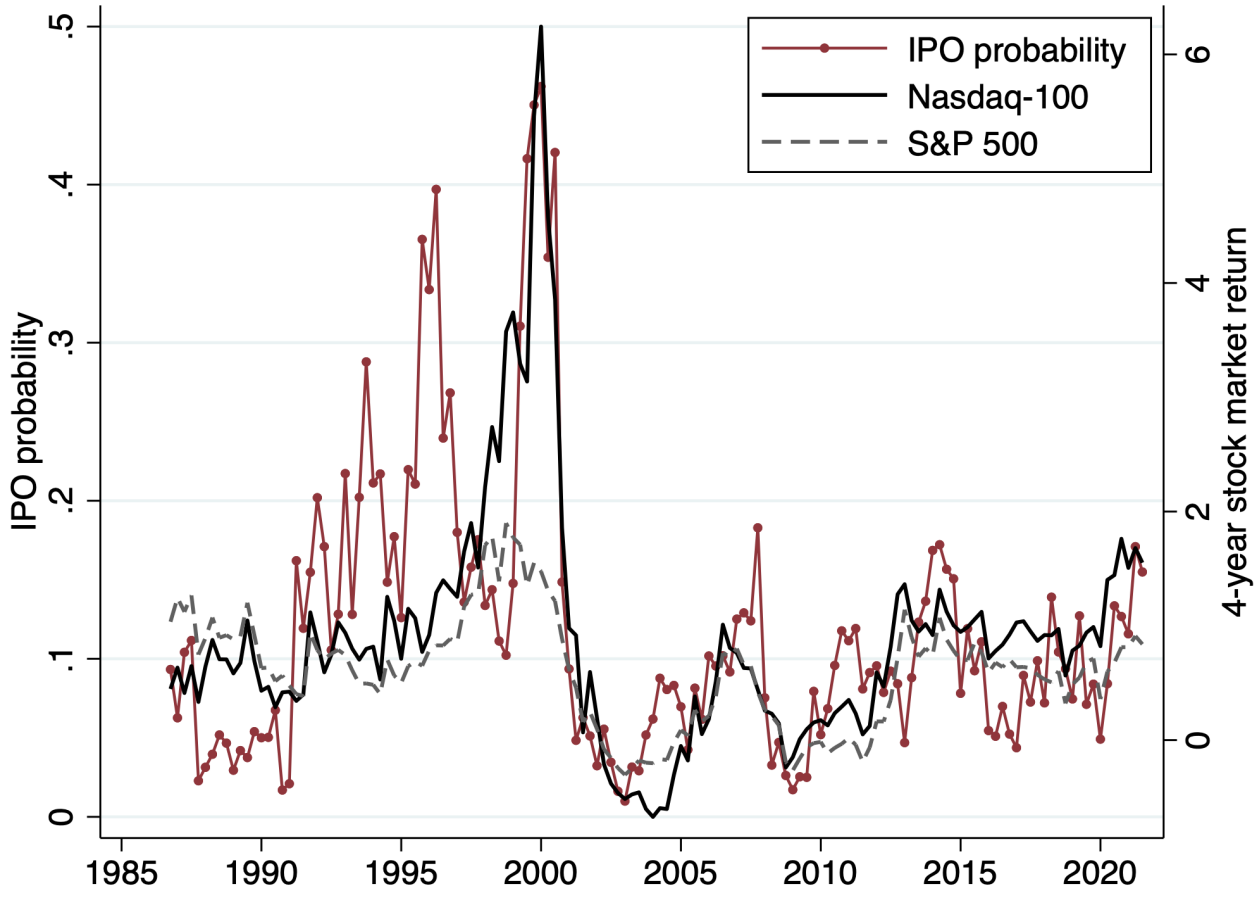
**Table 4: Time-Series of Startups and Exits**

| Year | #startups | #funds | #startups | #ipos | #ipos | #high-value mergers |
|---|---|---|---|---|---|---|
| Source | CA | CA | VX | Ritter | VX | VX |
| 1981 | | 9 | 468 | | | |
| 1982 | | 11 | 464 | | | |
| 1983 | | 28 | 639 | | 72 | |
| 1984 | | 32 | 584 | 42 | 36 | |
| 1985 | | 26 | 479 | 29 | 36 | |
| 1986 | | 30 | 538 | 78 | 81 | |
| 1987 | | 34 | 628 | 64 | 72 | 1 |
| 1988 | | 26 | 556 | 37 | 35 | 3 |
| 1989 | | 37 | 482 | 35 | 36 | 5 |
| 1990 | | 17 | 404 | 43 | 44 | 2 |
| 1991 | | 17 | 304 | 110 | 108 | 2 |
| 1992 | | 21 | 422 | 133 | 126 | 21 |
| 1993 | | 36 | 398 | 175 | 142 | 13 |
| 1994 | | 41 | 444 | 137 | 110 | 29 |
| 1995 | | 35 | 932 | 166 | 166 | 33 |
| 1996 | 1039 | 40 | 1203 | 239 | 244 | 45 |
| 1997 | 1046 | 70 | 1362 | 108 | 134 | 50 |
| 1998 | 1517 | 81 | 1847 | 61 | 84 | 52 |
| 1999 | 2336 | 110 | 2526 | 264 | 280 | 84 |
| 2000 | 3003 | 153 | 3521 | 249 | 233 | 134 |
| 2001 | 1481 | 54 | 1323 | 23 | 40 | 60 |
| 2002 | 1258 | 33 | 876 | 12 | 29 | 51 |
| 2003 | 1255 | 38 | 799 | 23 | 28 | 37 |
| 2004 | 1416 | 67 | 997 | 84 | 88 | 70 |
| 2005 | 1320 | 64 | 1104 | 52 | 54 | 76 |
| 2006 | 1508 | 80 | 1312 | 58 | 65 | 64 |
| 2007 | 1599 | 68 | 1461 | 83 | 90 | 86 |
| 2008 | 1418 | 66 | 1392 | 7 | 8 | 55 |
| 2009 | 908 | 23 | 907 | 14 | 13 | 40 |
| 2010 | 1206 | 50 | 1156 | 62 | 47 | 68 |
| 2011 | 1510 | 45 | 1433 | 53 | 50 | 97 |
| 2012 | 1368 | 56 | 1412 | 48 | 51 | 67 |
| 2013 | 1405 | 58 | 1535 | 79 | 81 | 54 |
| 2014 | 1522 | 81 | 1605 | 135 | 102 | 79 |
| 2015 | 1699 | 61 | 1771 | 79 | 73 | 46 |
| 2016 | 1351 | 68 | 1512 | 43 | 36 | 39 |
| 2017 | | | | 67 | 47 | 45 |
| 2018 | | | | 100 | 71 | 53 |
| 2019 | | | | 90 | 52 | 48 |
| 2020 | | | | 119 | 58 | 43 |
| 2021 | | | | 185 | 92 | 55 |

**Note:** This data provides an overview of the startup data.

round of funding in quarter $\tau$. The denominator reflects the average quarterly number of startups funded between $t-23$ and $t-8$, corresponding to a four-year window with a two-year lag. For example, the IPO probability in 2021Q4 is computed by dividing the number of IPOs and high-value mergers in 2021Q4 by the average number of startups funded per quarter from 2016Q1 through 2019Q4.

**Figure 2: IPO Probability vs. Stock Market Returns**




**Note:** This figure compares the quarterly IPO probability (as defined by eq. (1)) with the stock market return over the past 4-years.

Figure 6 plots the IPO (or "good outcome") probability over time. Prior to 1990, this probability remained below 10%, but it rose substantially during the 1990s, peaking near 50% during the tech bubble. Following the collapse of the bubble, the IPO probability fell below 5% and has fluctuated around 10% over the subsequent two decades.

The figure also compares the IPO probability to the cumulative returns of two major stock market indices—the S&P 500 (SPX) and the Nasdaq-100 (NDX)—measured over the preceding four years (i.e., from the end of quarter $t-16$ to quarter $t$). The figure highlights a strong positive correlation between public market returns and the likelihood of successful startup exits: when equity markets have per-

formed well, more VC-backed firms have either gone public or exited via high-value mergers. Gompers and Lerner (2004) document a similar pattern for the number of biotech IPOs and a biotech equity index.

Notably, the IPO probability has a stronger correlation with the NDX return (69%) than with the S&P 500 return (47%). This is consistent with the fact that most VC-backed startups list on Nasdaq (as shown in Table 2) and are therefore more exposed to the performance of the NDX index. This distinction is particularly salient during the tech bubble, when the performance of the NDX and SPX indices diverged sharply.

# 3 A Valuation Model of Startups and Venture Capital

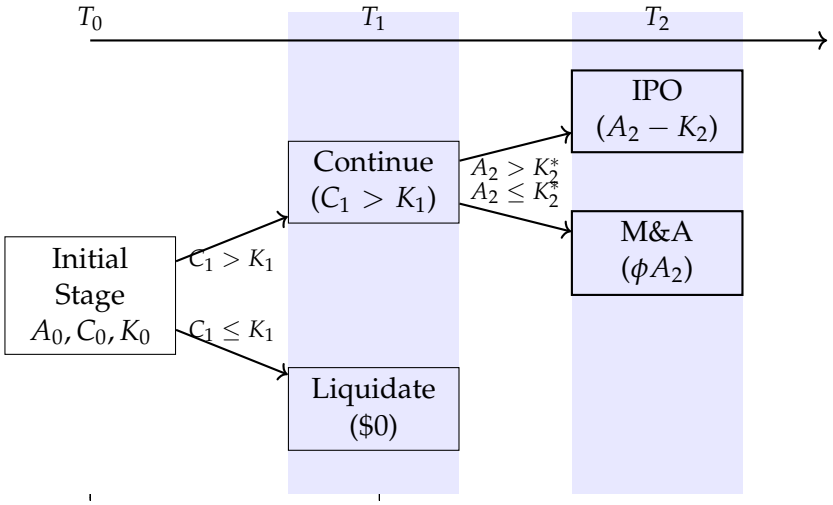
Startup ventures display three hallmark features. First, they face very high failure rates—roughly two-thirds of startups never return investors' capital—and extreme payoff skewness: most investments fail (zero payoff), while a small fraction deliver outsized gains, much like out-of-the-money call options (Gompers and Lerner, 2001; Kaplan and Strömberg, 2003; Berk, Green, and Naik, 2004; Cochrane, 2005). Second, to manage downside risk and align incentives, financing is divided into discrete, milestone-contingent tranches rather than provided all at once (Sahlman, 1990; Gompers, 1995; Wang and Zhou, 2002). Third, realized exit values and the timing of follow-on rounds co-move with aggregate market conditions—particularly Nasdaq index performance—so systematic market risk plays a key role in startup payoffs (Ritter, 1991). In this section, we embed these features into a parsimonious valuation framework that is readily taken to the data.

## 3.1 Staged-Financing Process

Figure 3 depicts the two-stage decision tree at the heart of our model. We model a startup as a series of staged investments $K_0, K_1, K_2$ across three development periods, combined with a claim on a terminal (post-money) payoff $A_2$. At each decision date $T_1$ and $T_2$, the investor chooses whether to continue by making the next required investment $K_t$, or to exit and realize the indicated payoff. We denote the startup continuation value before funding at $T_1$ as $C_1$ and its initial value as $C_0$. The initial investment, $K_0$, does not affect the initial value of the startup, but affects the founder's share, such that the founder's ownership share is $1 - \frac{K_0}{C_0}$.

Specifically, at $T_1$ (end of the "Early Stage"), the investor either **continues** by investing $K_1$ when the continuation value exceeds the required investment, $C_1 > K_1$, or **liquidates** for zero. If continued, then at $T_2$ (end of the "Late Stage"), the firm either **goes public** and realizes $A_2 - K_2$ if $A_2 > K_2^* = \frac{K_2}{1-\phi}$ or **exits via M&A** for a partial recovery of the underlying asset value, $\phi A_2$. To summarize, there are three exit possibilities: write-off, M&A, or IPO.

**Figure 3: Funding Stages and Payoffs over the Startup Lifecycle**




**Note:** This figure illustrates the startup model and its staged funding structure.

## 3.2 Startup Firms as Compound Options

We now map the decision tree into a nested call option on the underlying asset. The required investment amounts and their timing are pre-specified, making the optimal funding and continuation decisions dependent solely on the realized underlying asset value. The underlying asset pays no dividends and follows geometric Brownian motion. The distribution of the asset $A$ at time $t$ under the risk-neutral measure is given by

$$A_t = A_0 \cdot e^{\left(r_f - \frac{1}{2}\sigma_A^2\right) t + \sigma_A \sqrt{t}\, Z_A} \tag{2}$$

where $Z_A \sim \mathcal{N}(0,1)$ ($\mathcal{N}$ is the cumulative standard normal distribution), $r_f$ is the risk-free rate and $\sigma_A$ is the asset volatility.

### 3.2.1 Simplified example: Valuation with no recovery

To build intuition, we initially set the late-stage recovery parameter to zero ($\phi = 0$), so that the lower branch at $T_2$ pays zero and $K_2^* = K_2$. The late-stage payoff is then simply

$$\max(A_2 - K_2, 0),$$

a standard European call option that can be valued with the Black-Scholes option model (Black and Scholes (1973); Merton (1973)). Denote the time-$T_1$ value of the late stage payoff as $C_1$, then

$$C_1 = \text{Call}(A_1, K_2, \tau_2)$$

where $\text{Call}(A_1, K_2, \tau_2)$ denotes the Black-Scholes call option value where the underlying value is $A_1$, the strike price is $K_2$, and the option maturity is $\tau_2 = T_2 - T_1$.

At the decision date $T_1$, the firm continues only if the continuation value, $C_1$, exceeds the required investment amount, $K_1$, and otherwise it liquidates for zero. Hence, the early-stage payoff is

$$\max(C_1 - K_1, 0).$$

Standing at $T_0$, the startup can be viewed as a call option on the late-stage call option (there are two call decisions: one at $T_1$ and one at $T_2$). This call-on-call option can be valued with the Geske compound option valuation model (Geske (1979)).[6] Thus, the initial startup value is given

$$C_0 = \text{CompoundCall}(A_0, K_1, K_2, \tau_1, \tau_2)$$

where $\text{CompoundCall}(A_0, K_1, K_2, T_1, \tau_1, \tau_2)$ denotes the compound call option value where the underlying is $A_0$, the time to the first call date is $\tau_1 = T_1 - T_0$, the strike price at the first call date is $C_1$, the time between first and the second (final) call date is $\tau_2 = T_2 - T_1$, and the final strike price is $K_2$.

#### 3.2.2 Valuation with generalized recovery

We make one modification and allow for partial recovery on the lower late-stage branch at $T_2$, which will better match empirically observed late-stage exits. Specifically, if the late-stage investment at $T_2$ is not made, the firm is sold for $\phi \cdot A_2$, where $\phi \in [0, 1]$ is a recovery rate. This modification sets the late-stage payoff to be the greater of two positive values:

$$\max(A_2 - K_2, \phi \cdot A_2). \tag{3}$$

This means investment $K_2$ will be made if $A_2 - K_2 > \phi \cdot A_2$ and therefore if $A_2 > K_2^*$ where $K_2^* = \frac{K_2}{1-\phi}$. The terminal payoff in equation (3) can be replicated with a combination of the underlying asset and a standard call option with a modified strike price, $K_2^* = K_2/(1-\phi)$, which adjusts the strike to reflect the recovery condition.[7] Using the Black-Scholes call option formula, we can derive the startup value at

[6] Cassimon, Engelen, Thomassen, and Van Wouwe (2004) generalize the Geske (1979) two-stage compound option model to an N-stage model.

[7] This can be easily obtained from re-arranging: $\max(A_2 - K_2, \phi \cdot A_2) = \phi \cdot A_2 + \max((1-\phi)A_2 - K_2, 0) = \phi \cdot A_2 + (1-\phi) \cdot \max(A_2 - \frac{K_2}{1-\phi}, 0)$.

$T_1$ as

$$C_1 = \phi \cdot A_1 + (1-\phi) \cdot \underbrace{\left[A_1 \cdot \mathcal{N}(c_+) - K_2^* \cdot e^{-r_f \tau_2} \cdot \mathcal{N}(c_-)\right]}_{=\text{Call}(A_1, K_2^*, \tau_2)}$$

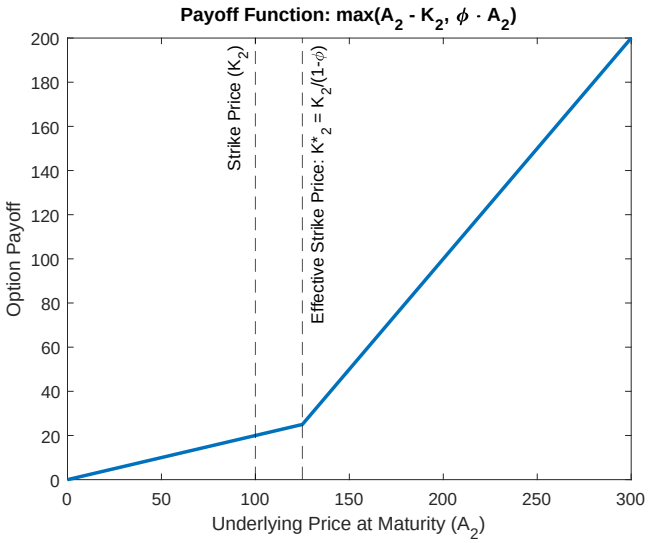
where

$$c_+ = \frac{\ln(A_1/K_2^*) + \left(r_f + \frac{1}{2}\sigma_A^2\right)\tau_2}{\sigma_A\sqrt{\tau_2}} \tag{4}$$

$$c_- = c_+ - \sigma_A\sqrt{\tau_2}.$$

Intuition for this replicating portfolio is developed in Figure 4, which plots the payoff function of an option with strike price $K_2 = 100$ and a recovery rate $\phi = 0.2$. The payoff is calculated as $\max(A_2 - K, \phi \cdot A_2)$, where $A_2$ represents the underlying asset value at maturity. The figure highlights two key points, including the strike price ($K_2$) and the effective strike price ($K_2^* = \frac{K_2}{1-\phi}$). A portfolio that owns $\phi$ shares of the underlying asset and $(1-\phi)$ call options on the underlying with strike price equal $K_2^*$ will perfectly replicate this payoff.

**Figure 4: Late-Stage Payoff with Partial Recovery**




This figure illustrates the payout profile of a call option with fractional recovery. The figure illustrates the stage-2 startup pre-IPO conditional payoff, where the terminal underlying asset value is $A_2$, the required investment to pursue an IPO (strike price) is $K_2 = 100$, and the fractional recovery rate is $\phi = 0.2$.

The early-stage payoff of the startup with a terminal recovery value (at $T_1$) is still

$$\max(C_1 - K_1, 0), \tag{5}$$

As before, the startup will only receive funding $K_1$ if its post-money value $C_1$ exceeds the funding cost,

i.e., $C_1 > K_1$. The initial startup continues to resemble a call option on the late-stage call option, but where the late-stage option now includes the recovery possibility. We next derive the initial startup value $C_0$ using a modified-version of the Geske (1979) compound option formula.

**Compound option pricing with recovery.** The Geske (1979) compound option formula is used to value an option on another option, such as a call-on-call. It works by accounting for the two stages of decision-making inherent in a compound option. At the first decision point $T_1$, the holder of the compound option decides whether to exercise it to obtain the embedded late-stage option. This occurs if the value of the late-stage option at $T_1$ exceeds the early-stage strike price. Denote the minimum $A_1$ that makes it optimal to exercise the stage-1 option as $A_1^*$. If this condition is satisfied, the holder pays $K_1$ to acquire the embedded late-stage option, and the compound option continues to the second stage. Otherwise, the compound option expires worthless. If the compound option is exercised at $T_1$, the final payoff depends on the late-stage option's value at its final maturity $T_2$. For our call-on-call, this occurs when the underlying asset price exceeds the embedded late-stage option's effective strike price $K_2^*$ (i.e., reflecting the recovery possibility).

The compound option's value at time $T_0$ reflects the risk-neutral probabilities of satisfying both conditions:

1. intermediate early-stage condition: $C_1 > K_1$ or, equivalently, $A_1 > A_1^*$,
2. final late-stage condition: $A_2 > K_2^*$.

The Geske model combines these probabilities into a two-dimensional framework. A formula for the modified compound option—a call option on a call option with a recovery value guarantee—can be assembled with the logic of the Geske model. Using the formula, the initial startup firm value $C_0$ can be

derived as

$$
\begin{aligned}
C_0 = \phi \cdot A_0 \cdot \mathcal{N}(a_+) + (1-\phi) \cdot A_0 \cdot \mathcal{N}_2(a_+, b_+, \sqrt{\tau_1/(\tau_2+\tau_1)}) \\
- (1-\phi) \cdot K_2^* \cdot e^{-r_f \cdot (\tau_2+\tau_1)} \cdot \mathcal{N}_2(a_-, b_-, \sqrt{\tau_1/(\tau_2+\tau_1)}) \\
- K_1 \cdot e^{-r_f \cdot \tau_1} \cdot \mathcal{N}(a_-),
\end{aligned}
$$

where

$$
\begin{aligned}
a_+ &= \frac{\ln(A_0/A_1^*) + \left(r_f + \frac{1}{2}\sigma_A^2\right)\tau_1}{\sigma_A\sqrt{\tau_1}} \\
a_- &= a_+ - \sigma_A\sqrt{\tau_1} \\
b_+ &= \frac{\ln(A_0/K_2^*) + \left(r_f + \frac{1}{2}\sigma_A^2\right)(\tau_2+\tau_1)}{\sigma_A\sqrt{\tau_2+\tau_1}} \\
b_- &= b_+ - \sigma_A\sqrt{\tau_2+\tau_1}.
\end{aligned} \tag{6}
$$

where: $\mathcal{N}_2(h,k,\rho)$ is the bivariate cumulative normal distribution function with $h$ and $k$ as upper integral limits and a correlation coefficient $\rho = \sqrt{\tau_1/(\tau_2+\tau_1)}$ reflecting the time to maturity of the initial call option relative to the time to maturity of the embedded call, where $\tau_1 = T_1 - T_0$ and $\tau_2 + \tau_1 = T_2 - T_0$.

The intuition behind equation 6 is as follows. $A_0 \cdot \mathcal{N}_2(a_+, b_+, \sqrt{\tau_1/(\tau_2+\tau_1)})$ is the expected terminal value of the underlying, given the joint probability of being in-the-money at $T_1$ and $T_2$; $\mathcal{N}_2(a_-, b_-, \sqrt{\tau_1/(\tau_2+\tau_1)})$ is the joint probability of being in-the-money at $T_1$ and $T_2$; $A_0 \cdot \mathcal{N}(a_+)$ is the expected value of the underlying asset at $T_1$, conditional on being in-the-money (i.e., $A_1 > A_1^*$); and $\mathcal{N}(a_-)$ is the probability of being in-the-money for the intermediate decision at $T_1$. The closed-form valuation formula derived from Geske's two-stage compound-option model explicitly captures the timing and interdependence of the two decision points, ensuring that the startup's initial value reflects the underlying asset's risk-neutral dynamics across both stages.

### 3.2.3 Exit Probabilities

Our model does not only allow us to value startup firms, but also provides closed-form solutions for the probabilities of various exit outcomes. The unconditional (Q-measure) probability of liquidation (i.e. exit at the end of stage 1) is equivalent to 1 - probability of the initial call option being in the money at the end of stage 1:

$$
\begin{aligned}
\Pr^{\mathrm{Q}}(\text{Continuation}) &= \mathcal{N}(a_-) \\
\Pr^{\mathrm{Q}}(\text{Liquidation}) &= 1 - \mathcal{N}(a_-).
\end{aligned} \tag{7}
$$

The unconditional (Q-measure) probability of IPO is the joint probability of being in-the-money at $T_1$ and $T_2$:

$$\Pr^{Q}(\text{IPO}) = \mathcal{N}_2(a_-, b_-, \sqrt{\tau_1/(\tau_2+\tau_1)}). \tag{8}$$

The M&A probability can be computed as 1 minus the sum of the unconditional liquidation and the unconditional IPO probability.

### 3.3 Integrating the CAPM into the Startup Valuation

As indicated by Figure 6, startup exit outcomes co-move with broad market performance, implying a non-diversifiable (systematic) risk component to each startup's total risk. When startups are pooled into VC funds and those funds become institutional portfolios, this market-driven risk becomes even more prevalent as idiosyncratic risk gets diversified away.

To analyze systematic risk, we integrate the compound option framework with a CAPM-style market factor. The key idea is to use the realized market return as the relevant state space for asset pricing, consistent with the Sharpe (1964) and Lintner (1965) CAPM. Terminal asset values are driven in part by a market factor, which allows for startup firm payoffs to be characterized as a function of the realized market return. This follows the spirit of structural-credit applications (e.g., Coval, Jurek, and Stafford (2009))[8], allowing us to capture how systematic shocks shape startup payoffs.

To discern systematic and startup-specific (idiosyncratic) factors, we start thinking about a large cross-section of startups and denote an individual startup by $i$. We assume that the risk-neutral distribution of the gross market return from time 0 to time $t$ is given by the following (where $M_t$ can be thought of as a total return index)

$$\frac{M_t}{M_0} = e^{\left(r_f - \frac{1}{2}\sigma_m^2\right)\cdot t + \sigma_m\sqrt{t}\,Z_m}, \tag{9}$$

where $\sigma_m$ is the volatility of the market return and $Z_m \sim \mathcal{N}(0,1)$ is a systematic market shock. The underlying asset value of startup i, $A_{i,t}$, is driven by a combination of the market shock and an idiosyncratic shock, $Z_{i,\epsilon} \sim \mathcal{N}(0,1)$, such that

$$A_{i,t} = A_{i,0}\cdot e^{\left(r_f - \frac{1}{2}\sigma_a^2\right)t + \beta_a\,\sigma_m\sqrt{t}\,Z_m + \sigma_\epsilon\sqrt{t}\,Z_{i,\epsilon}}, \tag{10}$$

where $\beta_A$ is the systematic market exposure (CAPM beta) of the asset, and $\sigma_\epsilon = \sqrt{\sigma_A^2 - \beta_A^2\sigma_m^2}$ is the idiosyncratic volatility of the underlying asset. Intuitively, one could think of this "change" as breaking

[8] Coval, Jurek, and Stafford (2009) embed the CAPM within Merton's (1974) credit model to study bond portfolios and CDO tranches. By mapping nonlinear systematic exposure into realized market returns, they enable no-arbitrage pricing via index option portfolios.

up the total shock $Z_{i,A}$ in equation (2) into a market- and a startup-specific component. Since this leaves the distribution of $A_{i,t}$ unchanged, the option pricing framework outlined above continues to hold.

From equation (9) it follows that $e^{\beta_A \sigma_m \sqrt{t} Z_m} = \left(\frac{M_t}{M_0}\right)^{\beta_A} \cdot e^{\left(-\beta_A r_f + \frac{1}{2}\beta_A \sigma_m^2\right) \cdot t}$, implying that the asset distribution conditional on the market return is

$$A_{i,t} = A_{i,0} \cdot \left(\frac{M_t}{M_0}\right)^{\beta_A} \cdot e^{\left((1-\beta_A) r_f - \frac{1}{2}\sigma_\epsilon^2\right) t + \sigma_\epsilon \sqrt{t} Z_{i,\epsilon}}. \tag{11}$$

Alternatively, one can also define the gross return of the market in excess of the risk-free rate as follows (since $M_t$ is compared to $M_0$ compounded at the risk-free rate)

$$\tilde{M}_t = \frac{M_t}{M_0 \cdot e^{r_f t}} \tag{12}$$

which means we can write

$$A_{i,t} = A_{i,0} \cdot \left(\tilde{M}_t\right)^{\beta_A} \cdot e^{\left(r_f - \frac{1}{2}\sigma_\epsilon^2\right) t + \sigma_\epsilon \sqrt{t} Z_{i,\epsilon}}. \tag{13}$$

Note that the only difference between equations (11) and (13) lies in how the risk-free return is incorporated into the asset value $A_{i,t}$. We will explore equation (13) when we analyze startup properties conditional on market excess returns below.

#### 3.3.1 Startup Betas

By embedding a systematic market factor into our option-pricing framework, we can directly derive the CAPM beta of startups. Specifically, startup beta is given by scaling the underlying asset's beta by the option's delta and the option-implied leverage (i.e., the asset-to-startup value ratio):

$$\beta_{i,t} = \beta_A \times \Delta_{i,t} \times \frac{A_{i,t}}{C_{i,t}}, \tag{14}$$

where the option delta $\Delta_{i,t} = \frac{\partial C_{i,t}}{\partial A_{i,t}}$ measures the sensitivity of the call price with respect to the underlying asset value. Solving for the option delta, one computes the **early-stage beta** as

$$\begin{aligned} \beta_{Startup,0} &= \beta_A \times \Delta_0 \times \frac{A_0}{C_0} \\ &\text{where} \quad \Delta_0 = \phi\, \mathcal{N}(a_+) + (1-\phi)\, \mathcal{N}_2\left(a_+, b_+, \sqrt{\tau_1/\tau_2}\right) \end{aligned} \tag{15}$$

Similarly, the late-stage beta of startup i can be computed as

$$\beta_{i,1} = \beta_A \times \Delta_{i,1} \times \frac{A_{i,1}}{C_{i,1}}$$
$$\text{where} \quad \Delta_{i,1} = \phi + (1-\phi)\mathcal{N}(c_+)$$

This means the value-weighted **late-stage beta** is given by

$$\beta_{Startup,1} = \mathbb{E}[\beta_{i,1}] = \int_i \beta_{i,1} \cdot C_{i,1} \cdot \mathbb{I}_{C_{i,1}>K_1} di, \tag{16}$$

where the indicator $\mathbb{I}_{C_{i,1}>K_1}$ keeps track of surviving startups.

### 3.4 The Dependence of Exit Probabilities and VC Returns on Market Returns

Integrating a market factor into our option pricing framework allows the model to speak to the market-dependence of (i) various VC exit probabilities and (ii) venture capital returns. We illustrate this next. For simplicity, we start by focusing on stage-1 outcomes and returns (the remainder of the paper then primarily deals with all-stage outcomes and returns).

The dependence of VC outcomes on the market return is illustrated in Figure 5. The figure demonstrates the important role of systematic risk in the underlying assets for shaping exit rates and portfolio returns. The left panels show results with $\beta_A = 0$ and the right panels show results with $\beta_A = 1.4$. All other parameters are set to the calibrated set shown in Table 5 below. The analysis draws on simulations of 10,000 startup firms. The x-axis in all plots is the market excess return defined in equation (12)—which we call "Realized Moneyness."

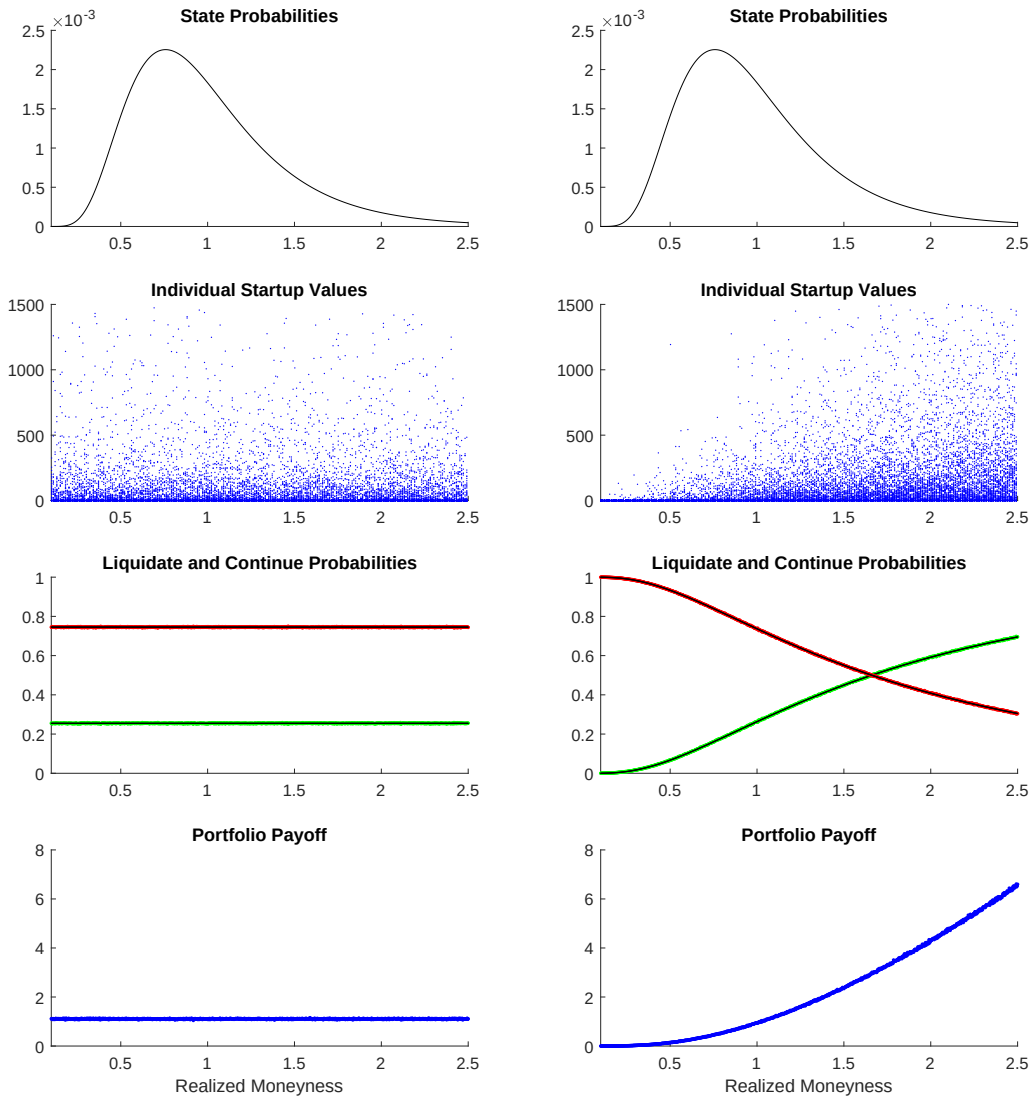
The *first row* displays the distribution of excess market returns. Of course, this is unaffected by the asset's beta.

The *second row* visualizes individual startup values $C_{i,1}$ at stage-1 after liquidations. For each realized market excess return value, we use equation (13) to simulate underlying asset values and then use the Black-Scholes option formula (4) to compute the startup value. If the value is below $K_1$, we set the value to 0. The key insight from these plots is that while startup values cluster near zero regardless of market conditions, the probability of high valuations increases with market returns when $\beta_A > 0$ (right panel).

The *third row* visualizes continuation and liquidation probabilities at stage-1 conditional on the excess market return at stage-1. Define the conditional cutoff (see the appendix for the derivation)

$$a_-(\tilde{M}_1) = \frac{\ln\left(A_0/A_1^*\right) + \left(r_f\,\tau_1 + \beta_A \cdot \ln(\tilde{M}_1)\right)}{\sigma_\epsilon\,\sqrt{\tau_1}},$$

**Figure 5: Stage-1 State-Contingent Model Properties**




**Note:** These plots illustrate the state-contingent properties of stage-1 model outcomes for $\beta_A = 0$ on the left-hand side, and $\beta_A = 1.4$ on the right-hand side. The graph is based on the simulation of 10,000 startups for each market excess return realization ("Realized Moneyness").

we can write the continuation and liquidation probabilities at stage-1 conditional on the excess market return at stage 1, $\tilde{M}_1$, as

$$\Pr\big(\text{Continuation} \mid \tilde{M}_1\big) = \mathcal{N}\big(a_-(\tilde{M}_1)\big), \tag{17}$$

$$\Pr\big(\text{Liquidation} \mid \tilde{M}_1\big) = 1 - \mathcal{N}\big(a_-(\tilde{M}_1)\big) = \mathcal{N}\big(-a_-(\tilde{M}_1)\big), \tag{18}$$

The conditional liquidation probabilities are computed both analytically, using the conditional formu-

las, and via simulation; the two approaches yield equivalent results except for "simulation noise". The simulated probability is computed as the fraction of the 10,000 startups where the continuation condition $C_{i,1} > K_1$ is satisfied. The plots reveal a key prediction: when startups have positive beta, their continuation probability increases strongly with market returns. For zero-beta startups, continuation probability is independent of market conditions. Thus, the sensitivity of VC outcomes to market returns depends on the underlying's systematic risk.

The *bottom row* analyzes the conditional expected portfolio returns for initial investors (founders and early-stage VCs) over stage-1. We compute the conditional expected return as the expected startup payoff by the initial investment (which is the initial startup value)

$$R(\tilde{M}_1) = \frac{\mathbb{E}[\max(C_{i,1} - K_1, 0) \mid \tilde{M}_1]}{C_0}. \tag{19}$$

The takeaway from the plots is that the VC portfolio payoff is strongly increasing in the market returns when the underlying asset has a positive beta (right panel). With zero beta, VC returns would exhibit no relationship to market returns (left panel).

# 4 Model Calibration and Evaluation

Our option-based model provides a simple framework to value startups. In this section, we (1) calibrate its parameters to match the unconditional rates of liquidation and IPO observed in the data, and (2) evaluate its ability to explain joint dynamics of startups' exit outcomes and equity market returns. Thus, the model is calibrated with *unconditional* moments, but evaluated with *conditional* moments.

## 4.1 Model Calibration

To calibrate the model, we infer several model parameters directly from the data and estimate the remaining ones by matching unconditional exit probabilities in the data. To match real-world exit outcomes, we need to compute model outcomes under the P-measure.

### 4.1.1 P-measure Exit Outcomes

Fortunately, in a lognormal setting, switching from the risk-neutral Q-measure to the physical P-measure is straightforward. This transformation—formally justified by Girsanov's theorem—amounts to replacing the risk-free drift $r_f$ in the underlying asset's process with its physical drift. To model the real-world, the drift of the market return becomes

$$r_f + \lambda_m$$

where $\lambda_m$ is the market risk premium. Similarly, the drift of the underlying asset's value is

$$r_f + \beta_A \cdot \lambda_m$$

where $\beta_A \cdot \lambda_m$ represents the asset's equilibrium risk premium under the CAPM, and $\beta_A$ is the asset's CAPM beta.

Thus, to convert the unconditional Q-measure liquidation and IPO probabilities given by equations (7) and (8) into objective P-measure probabilities, we replace the risk-free drift in the asset process with its physical drift in each of $a_-$ and $b_-$. The resulting probabilities, $\Pr^P\big(\text{Liquidation}\big)$ and $\Pr^P(\text{IPO})$ are then calibrated to match empirical exit rates.

#### 4.1.2 Calibration Results

To bring the model to the data, we rely on the data to the fullest extent possible. Several of the model parameters can be directly inferred from the data, such as the volatility and the beta of the underlying asset, see Table 3. We set the Nasdaq-100 risk premium to the ex-post realized excess return during our sample and the risk-free rate to the average 4-year Treasury rate. The strike ratios $K_2/K_1$ and the time-to-continuation decisions $\tau_1$ and $\tau_2$ are chosen in line with the VC characteristics documented in Table 1. This leaves two parameters to be calibrated: the stage-1 strike $K_1$ and the recovery rate $\phi$. We select the unconditional IPO rate of 12.8%[9] and the write-off rate of 64.6% (see Table 1) as calibration targets.

Table 5 shows the calibration results. It summarizes which parameters are inferred directly from data (asset value, volatility, beta, market risk premium, risk-free rate, stage durations, and $K_2/K_1$) and which are calibrated. We solve for the stage-1 strike $K_1$ and recovery rate $\phi$ such that the model's P-measure exit probabilities match the empirical targets. We find that $K_1 = 77$ and $\phi = 0.56$.

**Startup betas.** The calibrated model directly provides startup beta estimates, as outlined in equations (15) and (16). These stage-specific betas provide the foundation for the replicating levered-NDX strategy evaluated in Section 5. The calibration implies an early-stage beta of 2.2 and a late-stage beta of 1.6. Both exceed the underlying asset's CAPM beta of 1.4—exactly as theory predicts when a startup is viewed as a call-like derivative whose option-implied leverage amplifies the systematic market exposure of the underlying asset.

### 4.2 Model Evaluation: Startup Exits and Market Dynamics

Our structural compound-option model is calibrated to *unconditional* exit rates (liquidations and IPOs). We next evaluate the model based on the *conditional* IPO probability to see if our mechanism for market-

[9]The IPO probability measures the combined effects of "strong outcomes," including IPOs and high-value acquisitions.

**Table 5: Model Calibration**

| **Fixed Parameters** | Symbol | Value |
|---|---|---|
| Underlying value | $A_0$ | 100 |
| Duration first stage (yrs) | $\tau_1$ | 2 |
| Duration second stage (yrs) | $\tau_2$ | 2 |
| Vol of underlying | $\sigma_A$ | 0.90 |
| 4yr risk-free rate | $r_f$ | 0.05 |
| Ratio of funding amounts | $K_2/K_1$ | 1.73 |
| Beta of underlying | $\beta_A$ | 1.40 |
| Ex-post market risk premium | $\lambda_m$ | 0.13 |
| | | |
| **Moments** | Target | Calibrated |
| Unconditional IPO probability | 0.13 | 0.13 |
| Unconditional write-off probability | 0.65 | 0.65 |
| | | |
| **Calibrated Parameters** | Symbol | Value |
| Funding amount | $K1$ | 77 |
| Recovery parameter | $\phi$ | 0.56 |
| | | |
| **Implied values** | Symbol | Value |
| Startup value | $C_0$ | 42 |
| Early-stage NDX beta | $\beta_{Startup,0}$ | 2.23 |
| Late-stage NDX beta | $\beta_{Startup,1}$ | 1.55 |

**Note:** This table shows how the model parameters are calibrated.

dependent VC outcomes reproduces the empirical co-movement of IPO rates with public equity returns.

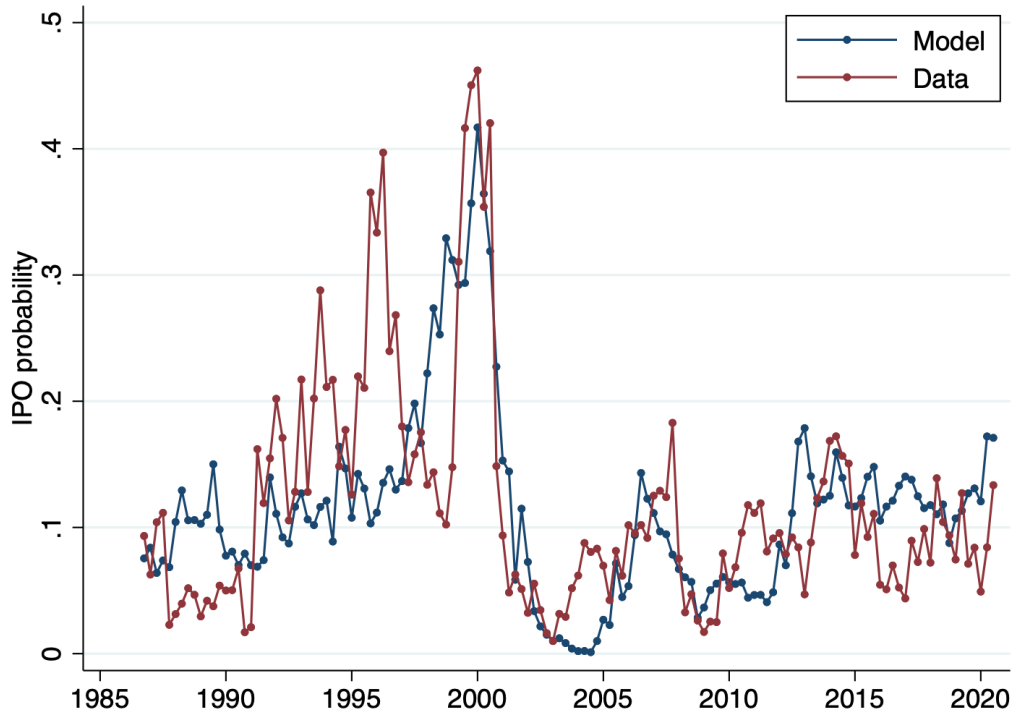
For each quarter $t$ from 1987 through 2024, we simulate the outcomes for 10,000 startups receiving their initial funding in quarter $t-16$, using realized equity market returns to compute asset values at each stage. We construct the startup's underlying asset value in two parts: a systematic market component and an idiosyncratic shock. Since this simulation is based entirely on realized returns, the resulting IPO probabilities are effectively computed under the P-measure.

To model the systematic component, we compound the daily returns of a portfolio that holds 140% in the Nasdaq-100 (i.e., borrowing 40%) over the two-year periods from $t-16$ to $t-8$ (for stage 1) and from $t-8$ to $t$ (for stage 2). This approach uses the actual empirical return distribution and avoids imposing a lognormal structure.[10] We then apply an idiosyncratic term by simulating $Z_{i,\epsilon} \sim \mathcal{N}(0,1)$ and scaling the systematic asset path by $e^{-\frac{1}{2}\sigma_\epsilon^2 t+\sigma_\epsilon\sqrt{t}Z_{i,\epsilon}}$.

This yields the stage-1 asset value $A_{i,1}$ for each startup $i$. We apply the stage-1 exercise rule $A_{i,1} > A_1^*$ to determine which startups continue to stage 2. For survivors, we simulate the stage-2 asset value $A_{i,2}$ using the same approach over the next 2-year window. We then apply the stage-2 threshold $A_{i,2} > K_2^*$; the fraction of startups satisfying both conditions defines the model-implied IPO probability in quarter $t$.

[10]The daily return of the systematic component is given by $r_f + 1.4 \cdot (r_{\text{NDX}} - r_f)$, compounded over each 2-year period. This corresponds to the term $\left(\frac{M_t}{M_0}\right)^{\beta_A} \cdot e^{(1-\beta_A)r_f t}$ from equation (11), but using actual realization of returns.

**Figure 6: Time series of IPO probabilities: Model vs. Data**




**Note:** This figure compares the model-implied quarterly IPO probabilities with the empirical IPO probabilities constructed following equation (1).

Figure 6 plots both the empirical and model-implied quarterly IPO probabilities. The model closely tracks the high run-up in IPO rates during the late-1990s and the subsequently lower rates post–2001. The model underpredicts the IPO spikes in the early 1990s when equity markets were strong, but growth firm valuations were less extreme than during the tech valuation peak. Overall, the close alignment between empirical and model-implied IPO rates demonstrates that our simple CAPM-driven compound-option mechanism captures the key co-movement between public market returns and VC exit outcomes.

To quantify the model's predicted sensitivity of startup exit rates to public markets, we estimate quarterly regressions of "good outcome" rates (IPOs + high-value acquisitions) on the prior four-year return of the Nasdaq-100 index:

$$\text{IPO_Rate}_t = b_0 + b_1 \cdot R^{\text{NDX}}_{t-4,t} + \varepsilon_t, \tag{20}$$

where $R^{\text{NDX}}_{t-4,t}$ denotes the log return of the Nasdaq-100 from four years before quarter $t$ through quarter $t$. We run this regression on (i) the empirical IPO series, (ii) the model-implied IPO series, and (iii) the empirical rates on the model rates, as well as in two subsamples (pre-2000 and post-2000).

Table 6 reports the slope coefficients and $R^2$ values. Empirical IPO rates are strongly linked to recent

NDX returns, with a slope coefficient of 0.06 (t-statistic = 5.8) and $R^2$ of 0.47. The model, calibrated only to the unconditional IPO rate, has a slope coefficient of 0.07 (t-statistic = 22.9) and $R^2$ of 0.93. A regression of empirical IPO rates on the model IPO rates produces a slope coefficient of 0.74 (t-statistic = 4.9), which is not reliably different from 1, and $R^2$ of 0.39. Finally, the last two specification show regressions of empirical IPO rates on 4-year NDX returns during the pre- and post-2000 samples, showing that these sensitivities are similar in both samples.

**Table 6: IPO Probability vs. Market Returns**

| | Quarterly IPO probability | | | | |
|---|---|---|---|---|---|
| | (1) Data | (2) Model | (3) Data | (4) Data | (5) Data |
| 4-year Nasdaq-100 return | 0.0598*** | 0.0711*** | | 0.0501** | 0.0607*** |
| | (0.0103) | (0.00311) | | (0.0229) | (0.00398) |
| Model-implied IPO probability | | | 0.735*** | | |
| | | | (0.150) | | |
| $R^2$ | 0.474 | 0.933 | 0.388 | 0.236 | 0.728 |
| N | 140 | 140 | 140 | 53 | 87 |
| Sample | | | | Pre-2000 | Post-2000 |

**Note:** Each column reports estimates of $\beta$ and $R^2$ from regressions of quarterly IPO (good outcome) rates on the preceding four-year Nasdaq-100 return, as in equation (20). Columns (1)–(2) use the empirically-observed and model-implied IPO probabilities, respectively; (3) regresses empirical on model-implied; (4)–(5) present the empirical sensitivity in pre-2000 and post-2000 subsamples.

We view these results as a major success for the basic design of the structural model. The compound option component allows for the calculation of various exit probabilities, and the market factor component introduces the ability to calculate these exit probabilities conditional on market return realizations, which match the empirical patterns remarkably well despite only being calibrated to the unconditional moments.

# 5 Venture Capital Returns

In this section, we deploy our structural compound-option framework to examine real-world VC performance. We begin by highlighting a surprising result: the model implies that the returns from startup investing can be closely replicated by taking a levered position in the Nasdaq-100. Building on this insight, we apply the levered-NDX strategy to assess the performance of actual VC funds along two dimensions. First, by evaluating vintage-level returns; second, by evaluating aggregate index-level returns.

### 5.1 Levered Nasdaq-100 Replicating Strategy

At first glance, using our model as a benchmark for VC may appear infeasible: compound options on individual Nasdaq stocks do not trade, and constructing a dynamic option-based replication strategy would be costly and complex. We therefore seek a practical alternative. Specifically, we compare model-implied payoffs to those of a levered Nasdaq-100 position. Such a levered strategy can be implemented at low cost by buying NDX futures or Nasdaq-100 ETF shares on margin.

The compound-option model implies that at each funding stage $k$ a startup claim has a beta $\beta_{Startup,k}$ with respect to the Nasdaq-100 index. We replicate this with a levered NDX strategy with the same beta. Specifically, from time $T_0$ to $T_1$, we hold $\beta_{Startup,0}$ units in a Nasdaq position. Startups surviving the early-stage then exhibit a beta of $\beta_{Startup,1}$, which we match by holding $\beta_{Startup,1}$ units in NDX. For example, in our calibrated model, an early-stage beta of 2.2 is matched by holding 220% of portfolio value in NDX (i.e., borrowing 120% at the risk-free rate). The late-stage beta of 1.6 corresponds to a 160% allocation to NDX (borrowing 60%).

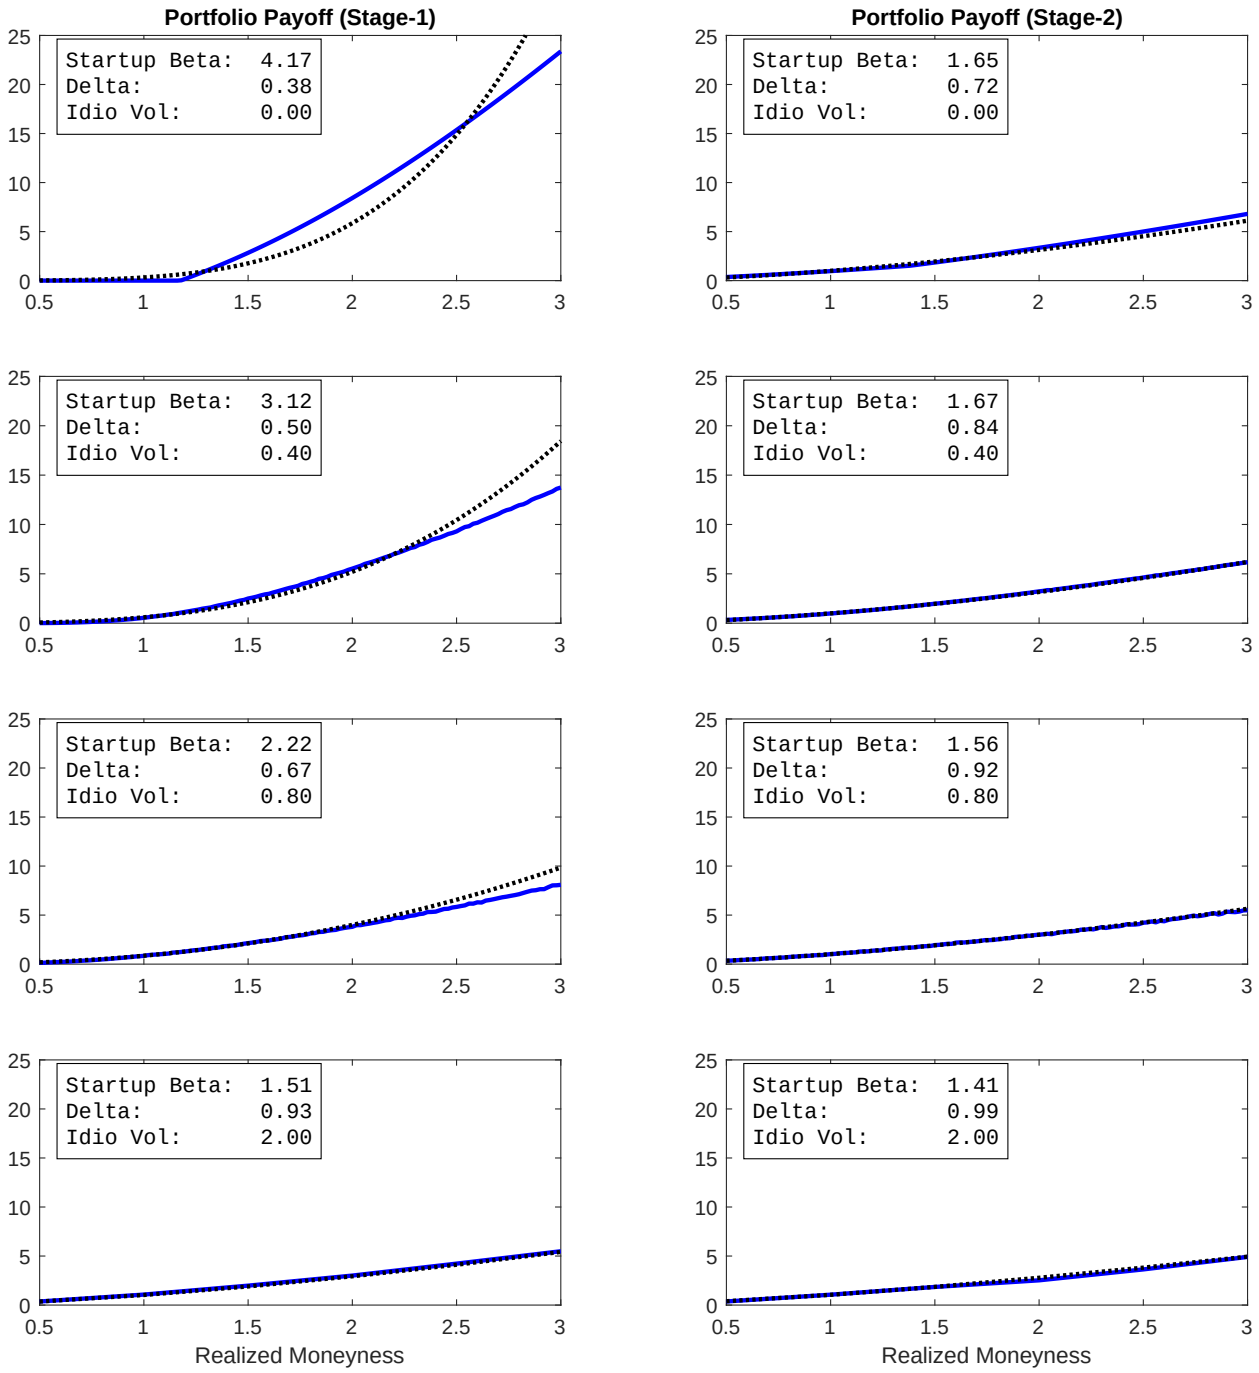
Figure 7 compares the conditional compound-option-model payoffs (solid lines) to this beta-matched levered-index strategy (dashed lines) for the first (left panel) and the second stage (right panel). The top to bottom panels show results for four levels of idiosyncratic volatility $\sigma_\epsilon$. In the top panel, with no idiosyncratic volatility (i.e., all startups share the same outcome), the model-implied payoff diverges significantly from the levered-index strategy. As idiosyncratic risk increases (moving down the rows), two things change. First, the convex "kink" at the stage-1 exercise boundary, in the dimension of market return, smooths out, dampening the stage-1 option's intrinsic convexity. Second, the effective portfolio beta declines, additionally reducing the levered index's convexity. Consequently, when $\sigma_\epsilon$ is large (in the range estimated for recently public VC-backed firms, e.g., $\sigma_\epsilon = 0.8$), the beta-matched levered-index strategy's payoffs match the model payoffs with high accuracy for both stages. Key to this finding is that the idiosyncratic risk of startups is high, a feature emphasized in prior deal-level studies (Cochrane, 2005; Korteweg and Sorensen, 2010).

The right-hand panels show late-stage payoffs assuming an early-stage moneyness $\tilde{M}_1 = 1.5$ (the market has done well during the early stage).[11] The recovery feature lowers model convexity (see Figure 4), and the late-stage beta is smaller, so the levered-NDX strategy matches the model payoff for all levels of idiosyncratic volatility.

We conclude that the model's conditional payoff profile can be closely replicated using a levered Nasdaq-100 strategy. By relying only on highly liquid futures or ETFs and model-implied beta ex-

[11] We find that the replication also works almost perfectly for different levels of first-stage moneyness (unreported).

**Figure 7: Approximating Model Payoffs with Levered Nasdaq-100 Exposure**




**Note:** This figure compares conditional model payoffs (blue lines) to the conditional payoffs of a levered Nasdaq-100 strategy (dotted lines). The leverage is chosen such that model-implied startup beta is matched. The left panels show stage-1 payoffs, and the right panels display stage-2 payoffs, conditional on a stage-1 moneyness of 1.5. The rows vary in terms of idiosyncratic risk, holding other calibrated parameters constant except for asset volatility, which remains a function of idiosyncratic volatility, $\sigma_A = \sqrt{\beta_A^2 \, \sigma_m^2 + \sigma_\epsilon^2}$.

posures, this approach offers a practical, low-cost implementation of systematic VC risk—without the need to trade long-dated options.

## 5.2 Vintage Returns

To evaluate how well our structural model and its simpler levered-NDX proxy capture real-world VC performance, we compare their vintage-level IRRs to those reported by two industry benchmarks: pooled vintage IRRs from Cambridge Associates (CA) and Preqin. Vintage-level internal rate of return (IRR) are a central metric of VC performance. To compute vintage-level pooled IRRs, the net cash flows (distribution minus contributions) of all funds are pooled together; the vintage IRR is computed as the discount rate that sets the net present value of these cash flows to zero.

Crucially, replicating VC returns in a public-market equivalent requires knowing both the appropriate market index and the systematic exposure (or, equivalently, the leverage). This cannot be directly inferred from private-fund returns, but is delivered by our structural model.

### 5.2.1 From Conditional Payoffs to Vintage-Level IRRs

To compare our model and the levered-NDX strategy to vintage-level VC IRRs, we convert the model's state-contingent payoffs into a vintage-level cash flows faced by limited partners (LPs) in a VC fund vintage. In practice, LPs make a series of capital contributions (negative cash flows) and receive distributions (positive cash flows) over time. We match this vintage-level cash flow profile with the following procedure.

**Constructing vintage-level cash flows & IRRs.** For each quarterly capital call in the historical vintage (at times $t_j$), we imagine investing that dollar into our replicating strategy:

1. At each contribution date $t_j$, for every \$1 of actual vintage capital called, we invest \$1 into the stage-1 replicator (either the calibrated model or the levered-NDX strategy) for a horizon of $\tau_1 = 2$ years.
    - Model: every \$1 is allocated to a diversified set of individual early-stage startups. The underlying asset of each startup is a daily-rebalanced NDX position with 1.4 beta plus idiosyncratic risk (consistent with our simulation assumptions).
    - Levered-NDX: every \$1 is used to enter a position with \$2.2 NDX-exposure.
2. At $t_j + \tau_1$, roll all remaining value into the stage-2 replicator with horizon $\tau_2 = 2$ years.
    - Model: Surviving startups become late-stage startups. The underlying asset of each startup is a daily-rebalanced NDX position with 1.4 beta plus idiosyncratic risk.

- Levered-NDX: the portfolio value is used to establish a position with 1.6 × Nasdaq-100 exposure.

3. At $t_j + \tau_1 + \tau_2$, roll any remaining value into a pure equity index position (beta equal to $\beta_A$) and hold until a maximum maturity of 6 years or until the final distribution date for that vintage.

   - Model: Startups become mature firms, replicated using a daily-rebalanced NDX position with 1.4 beta plus idiosyncratic risk.
   - Levered-NDX: Use the remaining value to establish a position with 1.4× Nasdaq-100 exposure.

At each historical distribution date, we exit positions and withdraw the cash needed to match the vintage's realized distribution. If the replicating strategy's balance is insufficient, the cashflow stream terminates at that point. We follow a first-in, first-out rule, so positions established earlier are exited first. Any surplus remaining at the final distribution is paid out as the terminal cash flow.

To get the cash flows of fund vintages, we use Preqin cash flow data and sum all distributions and contributions across all funds of a given vintage. We also count any remaining valuation at each fund's last observation as a distribution. These data are available for vintages starting in 1992. The vintage IRRs are highly consistent with the vintage IRR reported by CA, see Appendix Table IA1.

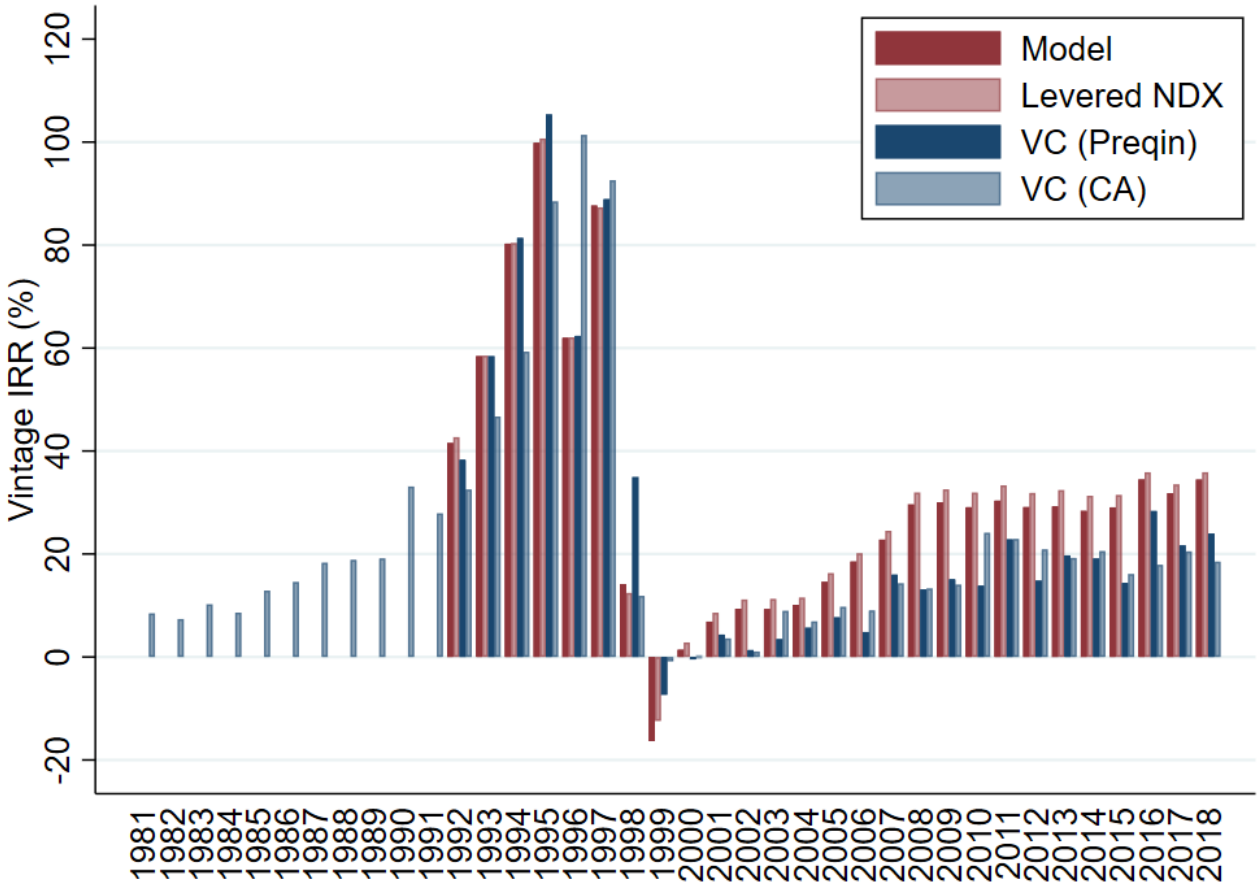
#### 5.2.2 Comparing VC Vintage IRRs to Model and Levered-NDX Benchmarks

Figure 8 presents a bar chart of IRRs by vintage cohort for all four series (model, levered-NDX, Cambridge, Preqin), while Table 7(A) reports detailed summary statistics and Table 7(B) shows the regression results of VC vintage IRRs on model-implied and levered-NDX vintage IRRs. Appendix Table IA1 reports the full set of IRRs for each historical vintage.

Over the main 1992–2018 sample, both the compound-option model and the levered-NDX replicator deliver nearly identical distributions of vintage IRRs, in terms of both means and volatilities (see Table 7). A regression of model-implied IRRs on the levered-NDX IRRs confirms the close alignment: the estimated slope is 1.02 (SE = 0.01), statistically indistinguishable from 1, with an intercept of $-2.3\%$, indicating the levered-NDX proxy has a slightly higher average return. Thus, the levered-NDX strategy achieves a close replication of the returns generated by our structural compound–option startup model.

Moreover, Table 7 compares actual VC IRRs to model-implied and levered-NDX IRRs. We find that the model-implied IRRs explain 92% of the variation in Preqin VC IRRs and 82% of variation in the CA VC IRRs. The slope coefficient is also statistically indistinguishable from 1. This is a remarkable success: our compound-option pricing model—calibrated only to unconditional exit outcomes—is able

**Figure 8: VC Vintage Returns (1981 - 2018)**




**Note:** This figures compares the vintage internal rate of returns (IRRs) of the model, the levered-NDX strategy, and Venture Capital returns from Preqin and Cambridge Associates. The levered NDX strategy is designed to replicate the contribution and distribution pattern of the VC vintages according to Preqin.

to explain the return variation across historical venture capital vintages with striking accuracy. The levered-NDX strategy does equally well, explaining 90% of the variation in Preqin VC IRRs and 83% of variation in the CA VC IRRs.

Notably, both Preqin and Cambridge Associates report very similar IRRs—whether for the full 1992–2018 period or when split into pre-2000 and post-2000 sub-samples. The returns of pre-2000 vintages are similar for actual VC funds and our replication strategies: the mean IRRs are 54% for Cambridge Associates, 58% for Preqin, 54% for our model, and 54% for the levered-NDX strategy. In contrast, the later vintages (2000–2018) show a pronounced divergence: Cambridge Associates and Preqin both average around 13–14%, whereas the model and levered-NDX benchmarks average approximately 23–25%. Thus, these VC vintages have substantially underperformed the benchmarks.

Despite this shift in mean returns, the standard deviations of the actual and model-based vintage

**Table 7: VC Vintage Returns**

(A) Summary Stats

| Sample | Statistic | Model | Levered NDX | VC (Preqin) | VC (CA) | VC (CA since 1981) |
|---|---|---|---|---|---|---|
| All vintages | Mean IRR (%) | 31.72 | 33.34 | 26.38 | 25.71 | 22.99 |
| | Standard deviation (%) | 26.38 | 25.81 | 28.71 | 27.96 | 24.22 |
| Pre-2000 vintages | Mean IRR (%) | 53.46 | 53.97 | 57.80 | 54.01 | 32.19 |
| | Standard deviation (%) | 39.10 | 38.44 | 35.77 | 38.27 | 31.19 |
| Post-2000 vintages | Mean IRR (%) | 22.57 | 24.66 | 13.15 | 13.79 | 13.79 |
| | Standard deviation (%) | 10.60 | 10.89 | 8.37 | 7.35 | 7.35 |

(B) Levered NDX vs. Venture Capital

| | IRR Model | IRR VC (Preqin) | | IRR VC (CA) | |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| IRR Levered NDX | 1.021*** | | 1.056*** | 0.980*** | |
| | (0.0103) | | (0.0993) | (0.113) | |
| IRR Model | | 1.042*** | | | 0.963*** |
| | | (0.0891) | | | (0.105) |
| Constant | −2.337*** | −6.667 | −8.818* | −6.980 | −4.848 |
| | (0.491) | (4.203) | (4.662) | (5.050) | (4.488) |
| $R^2$ | 0.998 | 0.916 | 0.900 | 0.818 | 0.826 |
| N | 27 | 27 | 27 | 27 | 27 |

**Note:** This table compares the vintage-level internal rate of returns (IRRs) of the model, the levered NDX strategy, and Venture Capital funds as reported by Preqin and Cambridge Associates (CA).

IRRs remain very similar across both sub-periods, underscoring that our replication captures the dispersion of IRRs across vintages. Having shown that the model and levered-NDX strategy match vintage-level VC performance, we now turn to market-level VC indices to evaluate aggregate returns and risk.

### 5.3 Index Returns

A widely used approach to assess VC performance is through market-level venture capital indices, such as the Cambridge Associates VC Index. We use the levered-NDX replication strategy to construct a comparable aggregate VC return series.

We obtain the quarterly cash flows of contributions and distributions from Cambridge Associates (1981–2024). Mirroring the vintage-level replication, each observed contribution is invested into a levered NDX position at 2.2× for the first two years, at 1.6× for the next two years, and finally at 1.4× for the subsequent four years. Each position is exited as cash flows are distributed to LPs or after a maximum of 6 years. As before, positions established earlier are exited first (first-in first-out). For each quarter $t$, we thus update the replicating strategy's market value by (i) subtracting any distributions paid out, (ii) adding any new contributions, and (iii) marking the remaining position to market using the realized return on the levered-NDX portfolio. From this series, we compute both the overall IRR (from the full

cash flow time series) and periodic net returns (percentage change in portfolio value with the difference between distributions and contributions treated as dividends), which enable standard time-series risk analysis.

Table 8 summarizes the results for the full sample (1981–2024) and two sub-periods (1981–2003, 2004–2024). In the early sample, the levered-NDX replicator achieves an IRR of 23%, compared to 19% for the Cambridge Associates VC Index. In the later sample, the levered-NDX's IRR of 19% exceeds the CA Index's 12%, mirroring the vintage-level divergence documented in Section 5.2.2. The levered-NDX strategy exhibits significantly deeper drawdowns than the Cambridge Index: –96% vs. –70% during the tech-bubble collapse, and –66% vs. –24% during the 2008 crisis.[12]

We also estimate standard factor regressions of quarterly excess returns on market excess returns (against the Nasdaq-100 and S&P). For the levered-NDX replicator, the estimated Nasdaq-100 beta is 1.74 (t-stat 15.2) and the S&P 500 beta is 2.43 (t-stat 18.7). Both are slightly above the model's average all-stage beta since there are only early-stage ventures active at the start of our exercise.. In contrast, regressions of the Cambridge VC Index include seven lags of the market return to account for valuation staleness and recover lower systematic betas (NDX beta 1.48; SPX beta 1.77). These findings likely reflect both (i) the well-identified systematic risk in our model-based VC replicating portfolio, and (ii) the potential impact of private-fund smoothing, which can understate market-driven volatility.

[12]These differences may reflect return smoothing arising from the staleness of reported startup valuations. We provide a discussion and analysis of smoothing mechanisms and their effects in Section 6.3.

**Table 8: VC Index Risk and Return**

| Sample | Stat | Levered NDX | CA VC Index |
|---|---|---|---|
| 1981Q1-2024Q2 | IRR (%) | 22.14 | 17.27 |
| | Geometric average (%) | 16.34 | 13.26 |
| | Arithmetic average (%) | 25.65 | 14.14 |
| | Standard deviation (%) | 45.22 | 19.13 |
| | Max drawdown (%) | -95.79 | -69.98 |
| | Nadaq-100 Alpha (%) | 0.00 (0.92) | -8.19 (2.67) |
| | Nadaq-100 Beta | 1.74 (0.03) | 1.48 (0.26) |
| | Nadaq-100 R-squared | 0.99 | 0.61 |
| | S&P 500 Alpha (%) | 0.69 (3.61) | -4.99 (2.27) |
| | S&P 500 Beta | 2.43 (0.14) | 1.77 (0.29) |
| | S&P 500 R-squared | 0.76 | 0.32 |
| 1981Q1-2003Q4 | IRR (%) | 22.90 | 19.04 |
| | Geometric average (%) | 12.97 | 14.19 |
| | Arithmetic average (%) | 26.86 | 15.79 |
| | Standard deviation (%) | 54.04 | 24.18 |
| | Max drawdown (%) | -95.79 | -69.98 |
| 2004Q1-2024Q2 | IRR (%) | 18.76 | 12.20 |
| | Geometric average (%) | 20.19 | 12.23 |
| | Arithmetic average (%) | 24.31 | 12.28 |
| | Standard deviation (%) | 33.12 | 11.07 |
| | Max drawdown (%) | -66.39 | -24.01 |

**Note:** This table reports the risk and return characteristics of the Cambridge Associates VC Index and our levered-NDX market-index replicator. We regress each index's quarterly excess returns on market excess returns—using only the contemporaneous return for the levered-NDX replicator and the contemporaneous return plus seven lags for the CA VC Index. Robust t-statistics are in parentheses.

# 6 Discussion

## 6.1 Key Lessons from the Structural Model

Our two-stage compound-option framework, calibrated solely to the unconditional probabilities of liquidation and IPO, succeeds in replicating both the time-series patterns of conditional exit rates and the vintage-level IRRs observed in the data. The model explains approximately 39% of the variation in quarterly IPO rates and 82% to 92% of the variation in vintage returns. This finding implies that the essential economics of VC-backed startups are captured by just two ingredients: the optionality embedded in staged financing, and exposure to systematic Nasdaq-100 index risk. By abstracting from dilution effects beyond fair-value common equity issuances, from matching between startups and VC investors (Sørensen, 2007), and from strategic timing discretion—which our IRR analysis captures via cashflow

matching—we isolate the pure value of the right to abandon and the market-beta channel.

Moreover, we show that a simple levered-NDX futures strategy—scaled to the model-implied betas to the Nasdaq-100 at each stage—closely approximates the model's conditional payoffs across a plausible range of idiosyncratic volatilities, offering a low-cost and transparent implementation of a VC replicating strategy.

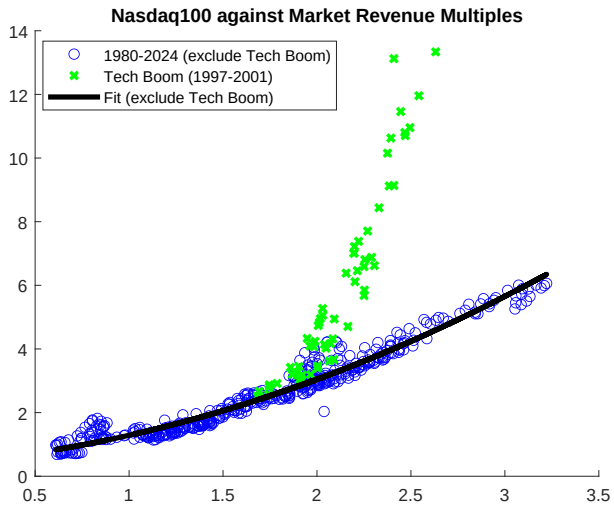
### 6.2 VC's Systematic Risk Exposure: Nasdaq-100 vs. S&P 500

Our analysis implies that the Nasdaq-100 index captures the relevant systematic market exposure of VC investments better than broader indices like the S&P 500. The underlying asset in our startup model is unobserved unless it becomes a viable publicly-traded firm. Ritter (1991) notes that most IPOs list on Nasdaq. This is also true for VC-backed startups: over 85% of VC-backed IPOs occur on Nasdaq (see Table 2). Consistent with this, we find that four-year Nasdaq-100 returns explain roughly 47% of the variation in quarterly IPO rates, compared to just 22% explained by the S&P 500 returns. We also regress quarterly Cambridge Associates VC Index returns on quarterly market returns (and 8 lags) and find an $R^2$ of 0.61 when using the Nasdaq-100 versus only 0.32 when using the S&P 500.

Nasdaq-100 exposure is distinct from broader US equity market exposure in this sample. First, in the post-2008 period, 2009-2024, the Nasdaq-100 exhibits positive risk-adjusted returns against the broader stock market (see Appendix Table IA2). This relative outperformance, sustained for nearly two decades, has meaningful consequences for the inferences from benchmarking exercises. Second, and more dramatically, the firms listed on the Nasdaq-100 exhibit convex valuations relative to average valuations. Figure 9 plots the value-weighted revenue multiples (enterprise value divided by revenues) for Nasdaq-listed firms against those for the broader U.S. stock market. The relationship is strongly positive and displays notable convexity, particularly pronounced during the tech boom period (1997-2001). This distinct nonlinearity underscores that Nasdaq exposure is fundamentally different from general market exposure, especially during market expansions characterized by rapid growth in technology and high-valuation sectors. Given that VC-backed startups predominantly list on Nasdaq, properly identifying and accounting for the Nasdaq-specific systematic risk exposure is essential when evaluating the investment performance and associated risks of VC investments.

Yet, most VC benchmarking studies nevertheless use the S&P 500 or a broad equity factor (such as the CRSP value-weighted returns) as the market proxy, potentially understating both the amount and nature of the systematic risk that LPs bear. Table 9 summarizes beta estimates reported in the VC risk and performance literature. Our stage-specific beta estimates provide novel insights into the evolution of market risk exposure through a startup's lifecycle, reflecting that the leverage on the underlying

**Figure 9: Stock Market Valuation: Nasdaq-100 vs. S&P 500**




This figures plots the monthly revenue multiples of stocks in the Nasdaq-100 and the US value-weight stock market over the period 1980-2024. The sub-period ranging from 1997-2001 is defined as the "tech boom" and is displayed as green 'x', while all other time periods are displayed as blue 'o'. The fitted values from a regression of Nasdaq-100 multiples on market multiples including the squared market multiple are plotted in black.

claim decreases as the startup survives the early stages, confirming the theoretical hypothesis of Berk, Green, and Naik (2004). The decreasing beta pattern is consistent with the estimates of Cochrane (2005), although our stage-level estimates seem quantitatively more realistic and contrast with the decreasing beta pattern of Korteweg and Sorensen (2010).

While few papers have estimated the life cycle beta of startups, many studies provide a general ("all-stage") beta estimate for VC funds. A key point to emphasize is that the systematic risk of VC investments in our study is derived from the systematic risk of newly-listed firms, which we estimate to be $\beta_{A,NDX} = 1.4$ and $\beta_{A,SPX} = 1.8$. If one accepts the option-like characteristic of venture capital, then the underlying asset betas serve as lower bounds for the venture capital beta, as the option-like structure inherently amplifies leverage of earlier-stage holdings. Under our calibration, the all-stage VC betas are $\beta_{VC,NDX} = 1.7$ and $\beta_{VC,SPX} = 2.3$, reflecting the leverage inherent in staged financing.

**Table 9: Market Betas of VC Investments**

| | | Stage 1 (Early) | Stage 2 (Late) | Mezzanine/IPO | All-stage Beta |
|---|---|---|---|---|---|
| NDX Beta | This paper | 2.2 | 1.6 | 1.4 | 1.7 |
| SPX Beta | This paper | 2.9 | 2.1 | 1.8 | 2.3 |
| | Cochrane (2005) | 1.1-0.9 | 0.7 | 0.5 | 1.9 |
| | Korteweg and Sorensen (2010) | 0.6–2.7 | 2.5 | 5.6 | 2.8 |
| | Gompers and Lerner (1997) | | | | 1.1–1.4 |
| | Peng (2001) | | | | 1.3–2.4 |
| | Woodward (2009) | | | | 2.2 |
| | Driessen, Lin, and Phalippou (2012) | | | | 2.7 |
| | Ewens, Jones, and Rhodes-Kropf (2013) | | | | 1.2 |
| | Ang, Chen, Goetzmann, and Phalippou (2018) | | | | 1.8 |
| | Brown, Ghysels, and Gredil (2023) | | | | 1.4–1.6 |
| | Korteweg and Nagel (2024) | | | | 2.4 |
| SPX Beta | Asset managers & Consultants | | | | 1.4 |

**Note:** This tables compares our market betas to other estimates in the literature.

### 6.3 Has Startup Risk Declined Post-2000?

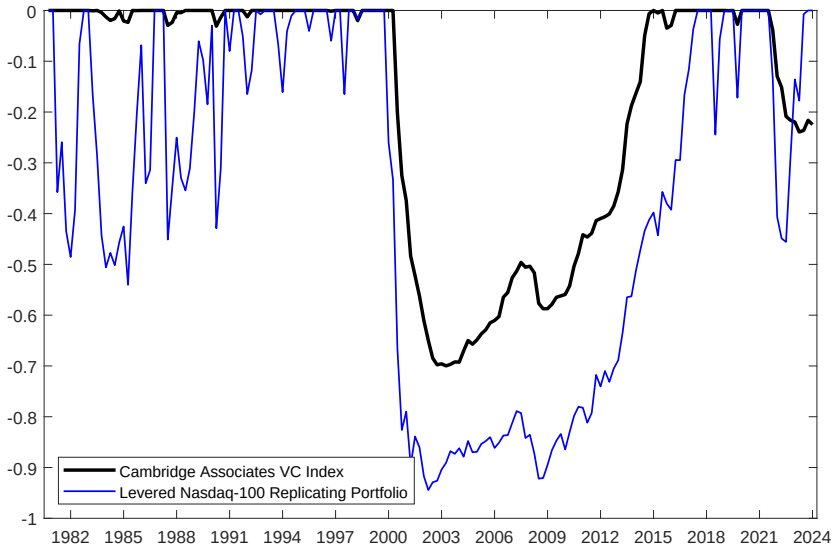
Since the "tech boom", venture-capital returns have underperformed relative to the returns of our compound-option benchmark and the returns of the levered NDX strategy. Vintage IRR data from Preqin and Cambridge Associates show that average VC vintage returns for cohorts between 2000 and 2016 lie in the 13–14% range, whereas our model and leveraged NDX strategy generates IRRs of roughly 23–25% for those same vintages (see Table 7). One possible interpretation is that the systematic risk profile of startups has come down over the past two decades, so that a model calibrated on unconditional exit probabilities over the full sample overstates the required risk premium over the last two decades.

Another observation that is challenging to interpret arises from the behavior of the Cambridge Index during market downturns. As shown in Figure 10, the index exhibits only three meaningful drawdowns over the past four decades: during the dot-com bust (2000-2002), the global financial crisis (2008-2009), and the recent venture correction beginning in 2022. In contrast, our well-marked replicating portfolio experiences substantial drawdowns in those three episodes as well as in numerous others, suggesting a more continuous and higher sensitivity to underlying risk factors. The muted decline in the Cambridge Index during the 2008 financial crisis is especially notable. While the levered-NDX replicating strategy declined by approximately 66%, the Cambridge-reported drawdown in VC returns was limited to around 24%. In one interpretation, this could signal that venture portfolios were less exposed to systematic risk than high-beta public tech equities. However, it may also reflect *artificial smoothing* in reported

returns.[13]

Smoothing might arise because the CA index is built off quarterly NAVs that incorporate stale or model-based valuations, rather than marking positions to market in real time. In benign environments, those NAVs tend to drift steadily upward as long as no valuation events (such as new funding rounds or exits) require reassessment. Meaningful declines only materialize when firms undergo down-round financings, impaired exits, or liquidations—at which point the drawdowns appear abruptly, often with a lag. Consistent with the smoothing interpretation, the CA index exhibits highly persistent return dynamics. Quarterly returns display large and statistically significant autocorrelations at the first three lags—approximately 0.6, 0.5, and 0.3, respectively—suggesting mechanically smoothed performance rather than timely price discovery. By contrast, our replicating portfolio shows no serial correlation, consistent with immediate valuation adjustments based on observable market signals.

**Figure 10: Venture Capital Drawdowns**




**Note:** This figures plots the quarterly drawdowns (through to peak levels of total return indices) for the Cambridge Associates Venture Capital Index (black) and the the levered Nasdaq-100 replicating portfolio (blue) over the period 1982-2024.

This return behavior underscores a fundamental challenge in private-firm risk assessment: without continuous market pricing, investors must infer both systematic and idiosyncratic risk from sparse, event-driven updates. Unlike public equities—whose daily valuations reflect the aggregated forward-looking views of market participants—venture-backed startups are effectively "invisible" between fi-

[13] A large body of literature documents return smoothing in hedge funds and private equity due to illiquidity and managerial discretion in NAV reporting. For hedge funds, see Asness, Krail, and Liew (2001); Getmansky, Lo, and Makarov (2004); Bollen and Pool (2008); Bollen and Pool (2009); Cassar and Gerakos (2011); and Cao, Chen, Liang, and Lo (2013). Jurek and Stafford (2015) show that smoothing in just two crisis months can obscure downside risk exposure in aggregate hedge fund indices. For private equity, see Jenkinson, Sousa, and Stucke (2019); Barber and Yasuda (2017); Chakraborty and Ewens (2018); Brown, Gredil, and Kaplan (2019); and Stafford (2022).

nancing events. Consequently, our compound-option framework provides one of the few transparent and economically grounded methods for revealing latent risk exposures, by linking observed exit dynamics and payoff structures to an underlying pricing model of risk.

Crucially, however, our model's calibration assumes that the startup underlying asset risk parameters remain constant over time, which translates into stage-specific betas being essentially stable over time. There are two empirical patterns that support the notion of constant underlying risk: First, the sensitivity of IPO exit rates to recent Nasdaq-100 returns is nearly identical before and after 2000. A regression of quarterly IPO probability on trailing 4-year market returns yields coefficients of 0.05 in the pre-2000 sample and 0.06 in the post-2000 sample (see Table 6). In both eras, "good-outcome" exits remain highly exposed to market performance. Second, post-IPO market betas also show no significant difference between the two subsamples. The average beta of firms going public between 1986 and 1999 is 1.43 and of those going public after 1999 is 1.39, indicating that conditional on an IPO exit, the asset's systematic risk has not drifted.

However, one change in exit patterns clearly stands out. High-value acquisitions now account for a much larger share of "good outcomes". Over our main 1992–2021 sample, 44% of successful exits are high-value M&A events; that fraction was just 22% in the early period and 53% after 2000 (see Table 4). This rise in high-value acquisition exits raises an important nuance. When startups are folded into large acquirers—as is often the case with technology buyouts—their underlying assets become part of publicly traded corporations that can exhibit lower systematic risk than standalone recent IPO firms.

We gauge the magnitude of this effect as follows. First, we identify the ten largest acquirers of VC-backed startups by deal count. We then estimate each acquirer's Nasdaq-100 beta and total volatility, as shown in Table 10. Next, we assume that startups are a levered claim on an underlying that is a 50/50 mixture of a representative standalone firm (with a beta of 1.4 and a total volatility of 0.90) and a representative acquirer (with a beta of 0.97 and a total volatility of 0.39). After recalibrating our model, the computed early-stage beta is 2.5 and the late-stage beta is 1.4. Thus, these beta estimates do not suggest that the systematic risk of VC has declined.

These analyses demonstrate that—even with a secular surge in high-value M&A exits—the model-implied systematic risk of VC-backed startups has remained remarkably stable. Taken together, the results reinforce the interpretation that VC returns since the tech bubble reflect not a structural reduction in startup risk, but rather underperformance relative to a transparent, risk-matched benchmark.

**Table 10: Beta and Volatility of 10 Most Frequent Startup Acquirers**

| Acquirer Name | Nasdaq-100 Beta | Volatility |
|---|---|---|
| Boston Scientific | 0.50 | 0.36 |
| Broadcom | 1.42 | 0.53 |
| Dell | 0.85 | 0.41 |
| Facebook | 1.08 | 0.37 |
| Google | 1.19 | 0.30 |
| IBM | 0.69 | 0.26 |
| Microsoft | 0.84 | 0.31 |
| Oracle | 0.71 | 0.38 |
| Salesforce | 1.26 | 0.43 |
| Yahoo | 1.11 | 0.51 |
| Average | 0.97 | 0.39 |

**Note:** This table provides beta and volatility estimates from 2000 to 2021 for the ten companies that acquired the largest number of VC-backed startups during the same period.

# 7 Conclusion

This paper develops and evaluates a parsimonious structural model that combines compound-option staging with CAPM-style systematic risk to provide a unified framework for understanding startup exits, payoff skewness, and VC portfolio returns. With only two calibration targets—unconditional liquidation and IPO probabilities—the model successfully replicates both (i) the time-series patterns in IPO activity and (ii) vintage-level VC returns, explaining 39% of the variation in quarterly IPO rates and 92% of the variation in vintage returns.

The model highlights the systematic risk exposure inherent in VC investing, particularly exposure to a tech-oriented equity index such as the Nasdaq-100. First, it links the well-documented "hot" and "cold" IPO cycles to underlying equity market conditions. Second, staged financing implies that startups can be viewed as levered claims on an underlying asset that—when successful—eventually lists on the public market, necessitating betas that exceed those of newly listed firms. Third, our results challenge the common practice of benchmarking VC performance against broad equity indices, given that most successful exits occur in technology-focused markets like the Nasdaq-100.

Most strikingly, our analysis implies that a diversified VC portfolio behaves similarly to a leveraged position in the Nasdaq-100. As such, a levered Nasdaq-100 investment offers a transparent, low-cost replicating strategy and a useful performance benchmark for limited partners. This replicator has meaningfully outperformed the average VC fund over the past two decades.

Future research could endogenize multi-stage investment timing, incorporate granular deal-level acquisition outcomes, and model dilution and GP incentive structures. These extensions would add nuance, but are unlikely to overturn the core insight: VC returns are largely explained by market-driven optionality and concentrated Nasdaq exposure.

## References


Ang, Andrew, Bingxu Chen, William N Goetzmann, and Ludovic Phalippou, 2018. Estimating private equity returns from limited partner cash flows. The Journal of Finance 73, 1751–1783.

Asness, Clifford S, Robert J Krail, and John M Liew, 2001. Do hedge funds hedge?. Journal of Portfolio Management 28, 6–19.

Barber, Brad M., and Ayako Yasuda, 2017. Interim fund performance and fundraising in private equity. Journal of Financial Economics 124, 172–194.

Bena, Jan, and Kai Li, 2014. Corporate innovations and mergers and acquisitions. The Journal of Finance 69, 1923–1960.

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 2004. Valuation and return dynamics of new ventures. Review of Financial Studies 17, 1–35.

Black, Fischer, and Myron Scholes, 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–654.

Bollen, Nicolas P.B., and Veronika K. Pool, 2008. Do hedge fund managers misreport returns? evidence from the pooled distribution. Journal of Finance 63, 2257–2288.

——— , 2009. Hedge fund risk dynamics: Implications for performance appraisal. Journal of Finance 64, 985–1035.

Broughman, Brian J., and Jesse M. Fried, 2012. Do vcs use inside rounds to dilute founders? some evidence from silicon valley. Journal of Corporate Finance 18, 1104–1120.

Brown, Gregory W, Eric Ghysels, and Oleg R Gredil, 2023. Nowcasting net asset values: The case of private equity. The Review of Financial Studies 36, 945–986.

Brown, Gregory W., Oleg R. Gredil, and Steven N. Kaplan, 2019. Do private equity funds manipulate returns?. Journal of Financial Economics 132, 267–287.

Cao, Charles, Yong Chen, Bing Liang, and Andrew W. Lo, 2013. Can hedge funds time market liquidity?. Journal of Financial Economics 109, 493–516.

Cassar, Gavin, and Joseph Gerakos, 2011. Hedge funds: Pricing controls and the smoothing of self-reported returns. Review of Financial Studies 24, 1699–1736.

Cassimon, Danny, Peter J. Engelen, Liesbeth Thomassen, and Martine Van Wouwe, 2004. Valuing new drug applications using n-fold compound options. Research Policy 33, 41–51.

Chakraborty, Itay, and Michael Ewens, 2018. Managing intangible capital. Review of Financial Studies 31, 1941–1970.

Cochrane, John H, 2005. The risk and return of venture capital. Journal of financial economics 75, 3–52.

Coval, Joshua D., Jakub W. Jurek, and Erik Stafford, 2009. The economics of structured finance. Journal of Economic Perspectives 23, 3–25.

Cremers, Martijn, Joost Driessen, and Pascal J. Maenhout, 2008. Explaining the level of credit spreads: Option-implied jump risk premia in a firm value model. Review of Financial Studies 21, 2209–2242.

Cumming, Douglas J., and John G. MacIntosh, 2003. A cross-country comparison of full and partial venture capital exits. Journal of Banking & Finance 27, 511–548.

Driessen, Joost, Tse-Chun Lin, and Ludovic Phalippou, 2012. A new method to estimate risk and return of nontraded assets from cash flows: The case of private equity funds. Journal of Financial and Quantitative Analysis 47, 511–535.

Ewens, Michael, and Joan Farre-Mensa, 2020. The deregulation of the private equity markets and the decline in ipos. The Review of Financial Studies 33, 5463–5509.

Ewens, Michael, Alexander Gorbenko, and Arthur Korteweg, 2022. Venture capital contracts. Journal of Financial Economics 143, 131–158.

Ewens, Michael, Charles M. Jones, and Matthew Rhodes-Kropf, 2013. The price of diversifiable risk in venture capital and private equity. Review of Financial Studies 26, 1854–1889.

Ewens, Michael, and Matthew Rhodes-Kropf, 2015. Is a vc partnership greater than the sum of its partners?. The Journal of Finance 70, 1081–1113.

——— , and Ilya A. Strebulaev, 2015. Inside rounds and venture capital returns. Discussion paper, SSRN Electronic Journal Preliminary draft available at SSRN.

Franzoni, Francesco, Eric Nowak, and Ludovic Phalippou, 2012. Private equity performance and liquidity risk. The Journal of Finance 67, 2341–2373.

Gao, Xiaohui, Jay R Ritter, and Zhongyan Zhu, 2013. Where have all the ipos gone?. Journal of Financial and Quantitative Analysis 48, 1663–1692.

Geske, Robert, 1979. The valuation of compound options. Journal of Financial Economics 7, 63–81.

Getmansky, Mila, Andrew W Lo, and Igor Makarov, 2004. An econometric model of serial correlation and illiquidity in hedge fund returns. Journal of Financial Economics 74, 529–609.

Gompers, Paul A., 1995. Optimal investment, monitoring, and the staging of venture capital. Journal of Finance 50, 1461–1489.

——— , and Josh Lerner, 1997. Risk and reward in private equity investments: The challenge of performance assessment. Journal of Private Equity 1, 5–12.

——— , 2000. Money chasing deals? the impact of fund inflows on private equity valuations. Journal of Financial Economics 55, 281–325.

——— , 2001. The venture capital revolution. Journal of Economic Perspectives 15, 145–168.

Gompers, Paul Alan, and Joshua Lerner, 2004. The venture capital cycle MIT press.

Gornall, Will, and Ilya A Strebulaev, 2020. Squaring venture capital valuations with reality. Journal of Financial Economics 135, 120–143.

Gupta, Arpit, and Stijn Van Nieuwerburgh, 2021. Valuing private equity investments strip by strip. The Journal of Finance 76, 3255–3307.

Jenkinson, Tim, Miguel Sousa, and Ruediger Stucke, 2019. Agency conflicts in asset management: Evidence from private equity. Journal of Financial Economics 132, 288–306.

Jurek, Jakub W., and Erik Stafford, 2015. The cost of capital for alternative investments. Journal of Finance 70, 2185–2226.

Kaplan, Steven N., and Antoinette Schoar, 2005. Private equity performance: Returns, persistence, and capital flows. Journal of Finance 60, 1791–1823.

Kaplan, Steven N, and Berk A Sensoy, 2015. Private equity performance: A survey. Annual Review of Financial Economics 7, 597–614.

Kaplan, Steven N., and Per Strömberg, 2003. Financial contracting theory meets the real world: An empirical analysis of venture capital contracts. Review of Economic Studies 70, 281–315.

Korteweg, Arthur, and Stefan Nagel, 2016. Risk-adjusting the returns to venture capital. The Journal of Finance 71, 1437–1470.

——— , 2024. Risk-adjusted returns of private equity funds: a new approach. The Review of Financial Studies p. hhae067.

Korteweg, Arthur, and Morten Sorensen, 2010. Risk and return characteristics of venture capital-backed entrepreneurial companies. The Review of Financial Studies 23, 3738–3772.

Korteweg, Arthur G, 2019. Risk adjustment in private equity returns. Annual Review of Financial Economics 11, 131–152.

Lerner, Josh, 1994. The syndication of venture capital investments. Financial Management 23, 16–27.

Lintner, John, 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. The Review of Economics and Statistics 47, 13–37.

Majd, Saman, and Robert S. Pindyck, 1987. Time to build, option value, and investment decisions. Journal of Financial Economics 18, 7–27.

Merton, Robert C., 1973. Theory of rational option pricing. Bell Journal of Economics and Management Science 4, 141–183.

——— , 1974. On the pricing of corporate debt: The risk structure of interest rates. Journal of Finance 29, 449–470.

Peng, Liang, 2001. Building a venture capital index. Working Paper.

Phalippou, Ludovic, and Oliver Gottschalg, 2009. The performance of private equity funds. Review of Financial Studies 22, 1747–1776.

Pindyck, Robert S., 1993. Investments of uncertain cost. Journal of Financial Economics 34, 53–76.

Raiffa, Howard, and Robert Schlaifer, 1961. Applied Statistical Decision Theory Harvard University Press.

Ritter, Jay R., 1991. The long-run performance of initial public offerings. Journal of Finance 46, 3–27.

Ritter, Jay R, 2015. Growth capital-backed ipos. Financial Review 50, 481–515.

Sahlman, William A., 1990. The structure and governance of venture-capital organizations. Journal of Financial Economics 27, 473–521.

Schwartz, Eduardo S., and Mark Moon, 2000. Rational pricing of internet companies. Financial Analysts Journal 56, 62–75.

Sharpe, William F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. The Journal of Finance 19, 425–442.

Sørensen, Morten, 2007. How smart is smart money? a two-sided matching model of venture capital. The Journal of Finance 62, 2725–2762.

Stafford, Erik, 2022. Replicating private equity with value investing, homemade leverage, and hold-to-maturity accounting. Journal of Finance 77, 2299–2343.

Wang, Jie, and Wenyu Zhou, 2002. Staged financing and endogenous soft information. Journal of Finance 57, 81–114.

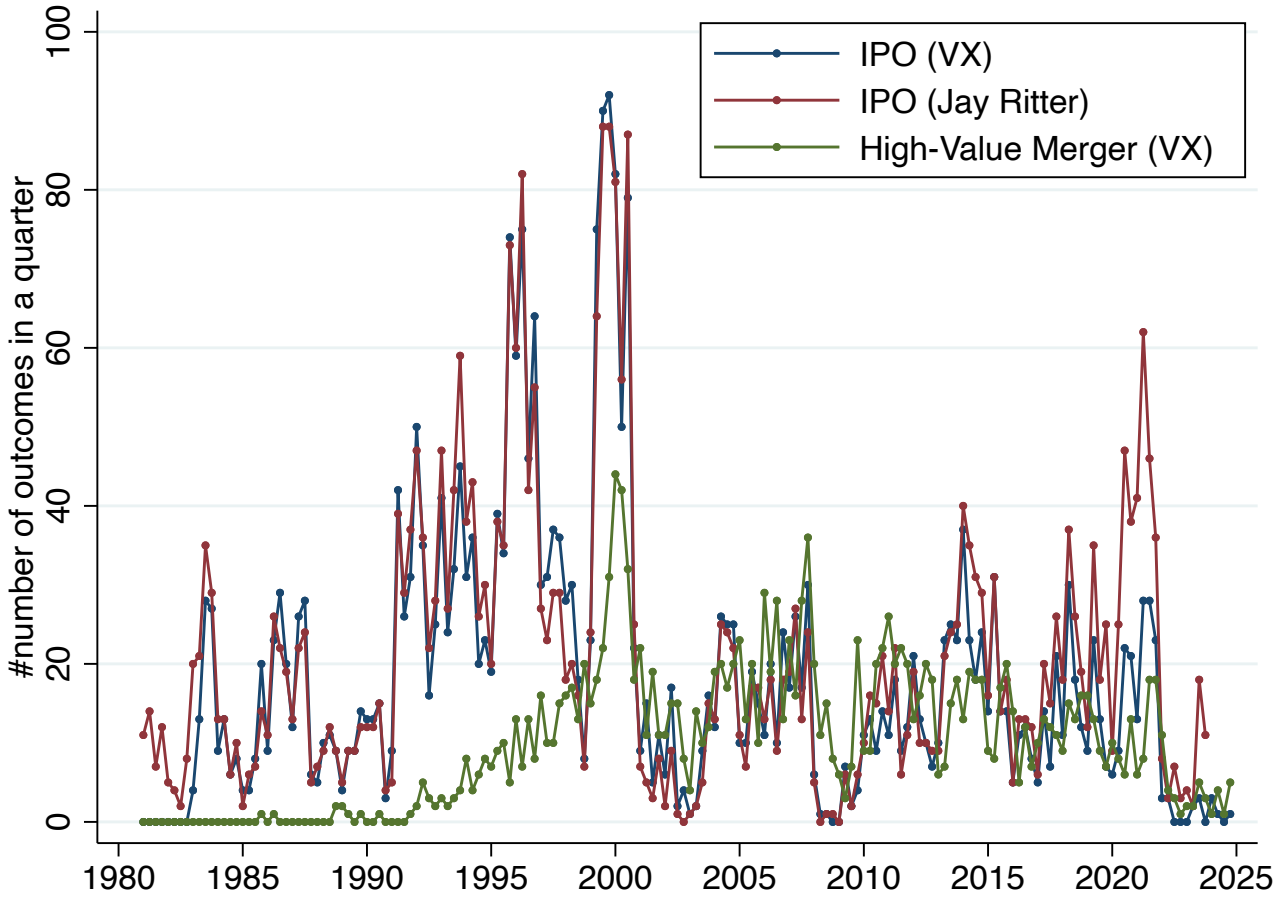
Woodward, Susan E., 2009. Measuring risk for venture capital and private equity portfolios. Discussion paper, Sand Hill Econometrics Available at `https://sandhillecon.com/pdf/MeasuringRiskForVentureAndBuyouts.pdf`.

Internet Appendix for:

# Venture Capital as Portfolios of Compound Options

# A Time-Series of Startup Exits

**Figure IA1: Time-series of Successful Startup Exits**




**Note:** This figure shows the time-series of quarterly startup exits. "VX" denotes the VentureXpert data.

# B VC Returns: Vintage By Vintage

**Table IA1: Vintage Returns**

| Vintage | Model | Levered NDX | VC (Preqin) | VC (CA) |
|---|---|---|---|---|
| 2018 | 34.5 | 35.9 | 24.0 | 18.5 |
| 2017 | 31.8 | 33.5 | 21.7 | 20.5 |
| 2016 | 34.5 | 35.9 | 28.3 | 17.9 |
| 2015 | 29.1 | 31.5 | 14.4 | 16.1 |
| 2014 | 28.4 | 31.3 | 19.1 | 20.6 |
| 2013 | 29.3 | 32.4 | 19.7 | 19.2 |
| 2012 | 29.1 | 31.9 | 14.8 | 20.9 |
| 2011 | 30.4 | 33.3 | 22.8 | 22.9 |
| 2010 | 29.1 | 31.9 | 13.9 | 24.1 |
| 2009 | 30.1 | 32.6 | 15.1 | 14.1 |
| 2008 | 29.6 | 31.9 | 13.1 | 13.3 |
| 2007 | 22.8 | 24.5 | 16.0 | 14.3 |
| 2006 | 18.5 | 20.1 | 4.8 | 9.1 |
| 2005 | 14.6 | 16.3 | 7.7 | 9.8 |
| 2004 | 10.1 | 11.5 | 5.7 | 6.9 |
| 2003 | 9.3 | 11.3 | 3.5 | 9.0 |
| 2002 | 9.4 | 11.2 | 1.3 | 1.1 |
| 2001 | 6.9 | 8.6 | 4.4 | 3.6 |
| 2000 | 1.4 | 2.8 | -0.5 | 0.0 |
| 1999 | -16.3 | -12.4 | -7.3 | -0.8 |
| 1998 | 14.2 | 12.4 | 34.9 | 11.9 |
| 1997 | 87.7 | 87.3 | 88.9 | 92.6 |
| 1996 | 62.0 | 62.1 | 62.4 | 101.4 |
| 1995 | 99.9 | 100.7 | 105.4 | 88.5 |
| 1994 | 80.3 | 80.4 | 81.4 | 59.3 |
| 1993 | 58.4 | 58.5 | 58.4 | 46.7 |
| 1992 | 41.6 | 42.7 | 38.3 | 32.5 |
| 1991 | . | . | . | 27.9 |
| 1990 | . | . | . | 33.1 |
| 1989 | . | . | . | 19.2 |
| 1988 | . | . | . | 18.9 |
| 1987 | . | . | . | 18.3 |
| 1986 | . | . | . | 14.6 |
| 1985 | . | . | . | 12.9 |
| 1984 | . | . | . | 8.6 |
| 1983 | . | . | . | 10.2 |
| 1982 | . | . | . | 7.4 |
| 1981 | . | . | . | 8.5 |

**Note:** This table reports the vintage internal rate of returns (IRRs) of the model, the levered NDX strategy, and Venture Capital vintages. The IRRs are reported in %.

# C Stock Market Returns

**Table IA2: Stock Market Returns**

| Sample | Stat | S&P 500 (SPX) | Nasdaq-100 (NDX) |
|---|---|---|---|
| 1981Q1-2024Q2 | Geometric average (%) | 11.61 | 13.68 |
| | Arithmetic average (%) | 12.45 | 16.33 |
| | Standard deviation (%) | 16.16 | 25.67 |
| | Max drawdown (%) | -45.76 | -81.04 |
| 1981Q1-2003Q4 | Geometric average (%) | 12.88 | 12.97 |
| | Arithmetic average (%) | 13.66 | 17.01 |
| | Standard deviation (%) | 16.54 | 30.42 |
| | Max drawdown (%) | -43.75 | -81.04 |
| 2004Q1-2024Q2 | Geometric average (%) | 10.20 | 14.48 |
| | Arithmetic average (%) | 11.10 | 15.57 |
| | Standard deviation (%) | 15.80 | 19.16 |
| | Max drawdown (%) | -45.76 | -41.70 |

**Note:** This table reports the return characteristics of stock market indices.

# D Conditional Continuation Probability

To derive the conditional continuation cutoff $a_-(\tilde{M}_1)$, start from the conditional asset process in equation (13):

$$A_{i,1} = A_0 \cdot \tilde{M}_1^{\beta_A} \cdot e^{\left(r_f - \frac{1}{2}\sigma_\epsilon^2\right)\tau_1 + \sigma_\epsilon \sqrt{\tau_1} Z_{i,\epsilon}}.$$

Taking logs and rearranging the continuation condition $A_{i,1} > A_1^*$, we obtain:

$$Z_{i,\epsilon} > \frac{-\left[\log(A_0/A_1^*) + \beta_A \log(\tilde{M}_1) + r_f \tau_1\right]}{\sigma_\epsilon \sqrt{\tau_1}}.$$

Thus, the standardized cutoff becomes:

$$a_-(\tilde{M}_1) = \frac{\log(A_0/A_1^*) + \beta_A \log(\tilde{M}_1) + r_f \tau_1}{\sigma_\epsilon \sqrt{\tau_1}},$$

yielding the conditional continuation probability:

$$\Pr\left(\text{Continuation} \mid \tilde{M}_1\right) = \mathcal{N}\left(a_-(\tilde{M}_1)\right).$$