Venture Capital as Portfolios of Compound Options*

Sebastian Hillenbrand
Harvard Business School

Erik Stafford
Harvard Business School

January 2025

Abstract

Startups are modeled as two-stage compound call options on an underlying asset with systematic risk exposure. The compound option structure captures the staged investment and funding decisions that determine whether a startup continues its development or liquidates at the end of the early stage. By integrating a CAPM-style market factor into the option framework, we characterize the nonlinear systematic risk exposure of startups as a function of realized market returns. We calibrate the model using only unconditional startup liquidation and IPO rates, then evaluate its predictions about exit rates and returns conditional on market performance. The calibrated model successfully matches two key empirical patterns: the time series of IPO rates, and the cross-section of VC vintage returns. The structural approach suggests feasible VC replication strategies using traded index options.

*We thank Joshua Coval for helpful comments and discussions. 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), yet the relationship between market conditions and VC investment outcomes remains imperfectly understood. While these investments offer the potential for significant returns, they also involve substantial risks, particularly due to their sensitivity to market conditions. Successful startup investments typically exit through IPOs or acquisitions during strong market conditions, generating high payoffs for founders and VC funds. However, many startup investments fail entirely or exit at modest valuations, especially during market downturns. This payoff profile, characterized by high skewness and extreme uncertainty, reflects both the option-like nature of startup investments and their systematic exposure to market conditions.

This paper develops and tests a structural model of startup firm valuation that integrates standard option pricing theory with systematic risk exposure through a market factor. Our framework builds on prior work exploring various aspects of VC investing, such as sample selection (Cochrane, 2005; Korteweg and Sorensen, 2010) and the optionality created by staged investments (Berk, Green, and Naik, 2004). Our model is designed to be tractable enough to empirically evaluate predictions about market-conditional startup exit rates and VC investment returns.

We model individual startup firms as compound call options, where each stage of development represents an option on the next stage's option. At each decision point, the firm incurs an investment cost to continue development and acquire the option on the subsequent stage. The firm proceeds only if the expected value of future options exceeds the required investment; otherwise, it exits and realizes the terminal value of the current stage.

To analyze the systematic risks of startups and portfolios, we integrate this 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. Underlying asset values are partially driven by a market factor, which allows startup firm payoffs to be characterized as a function of the realized market return.[1] An important implication of this framework is that an early-stage 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 early-stage startups.

[1]This approach follows the spirit of (Coval, Jurek, and Stafford, 2009) (CJS) who integrate the CAPM into the (Merton, 1974) credit model to study the pricing of bond portfolios and CDO tranches.

We calibrate the model parameters to match unconditional startup liquidation and IPO rates and then evaluate two key predictions. First, predicted IPO rates should track actual rates in the time series. Second, model-implied VC vintage returns should match actual returns. Despite being calibrated only to unconditional exit probabilities, the model explains approximately 23% to 42% of the variation in quarterly IPO rates and 55% to 74% of the variation in vintage returns. Notably, the calibration procedure does not make use of any return data other than unconditional market factor parameters and the underlying asset beta.

Our results provide several insights for understanding entrepreneurial finance. The model's success suggests that the compound option structure combined with market factor exposure captures key economic forces driving startup outcomes. Moreover, the framework provides a basis for analyzing practical applications in startup investing, including the design of VC replication strategies using traded index options.

Related Literature Our paper builds on the literature pricing standard options (Black and Scholes, 1973; Merton, 1973, 1974) and compound options (Geske, 1979). The option framework naturally extends to modeling growth options and staged investment decisions (Majd and Pindyck, 1987; Pindyck, 1993; Childs and Triantis, 1999; Berk, Green, and Naik, 1999; Schwartz and Moon, 2000). While earlier work applied this approach to startups [Berk, Green, and Naik (2004) and Cassimon, Engelen, Thomassen, and Van Wouwe (2004)], we develop a tractable framework that integrates systematic risk exposure following Coval, Jurek, and Stafford (2009) and Cremers, Driessen, and Maenhout (2008). This integration allows us to characterize how option-like payoffs interact with market risk to generate predictions about both exit decisions and returns.

A large literature examines VC risk and return characteristics; see Kaplan and Sensoy (2015) and Korteweg (2019) for reviews of this literature. Deal-level studies must address significant selection bias, as successful outcomes are more likely to be observed, e.g. Cochrane (2005) and Korteweg and Sorensen (2010).[2] Our structural model provides a framework for understanding this selection, as exit decisions depend explicitly on market conditions. Fund-level studies face different challenges due to infrequent valuations and complex fee structures (e.g., Gompers and Lerner, 1997; Phalippou and Gottschalg, 2009; Woodward, 2009; Ewens, Jones, and Rhodes-Kropf, 2013; Brown, Ghysels, and Gredil, 2023). Despite methodological differences, our finding of high systematic risk among startup firms aligns with estimates from both deal-level Cochrane (2005) and Korteweg and Sorensen (2010) and fund-level analyses

[2]Another set of challenges that the deal-level literature has side-stepped are the seniority differences across different share classes (Gornall and Strebulaev, 2020) and the risk-sharing features in contracts between VC funds and entrepreneurs (Ewens, Gorbenko, and Korteweg, 2022).

(e.g. Woodward, 2009; Driessen, Lin, and Phalippou, 2012; Ang, Chen, Goetzmann, and Phalippou, 2018; Korteweg and Nagel, 2024).

The evolution of systematic risk through startup lifecycles remains debated. Some studies find increasing betas (Korteweg and Sorensen, 2010), while others document decreases (Cochrane, 2005). Our model provides a structural foundation for declining betas, as the compound option structure implies decreasing leverage as startups survive early stages, consistent with Berk, Green, and Naik (2004). This pattern has implications for risk measurement, investment evaluation, and portfolio allocation.

The performance evaluation literature has developed various approaches including Public Market Equivalents (PME) (Kaplan and Schoar, 2005) and their extensions (e.g. Korteweg and Nagel (2016)). While many studies benchmark performance against equity indices (e.g. Woodward and Hall, 2004; Kaplan and Schoar, 2005; Korteweg and Nagel, 2016), recent work emphasizes broader risk factor exposure (e.g., Franzoni, Nowak, and Phalippou, 2012; Gupta and Van Nieuwerburgh, 2021). Our framework suggests specific option-based replication strategies that capture both the nonlinear payoffs and systematic risk exposure of VC investments.

2 A Valuation Model of Startups and Venture Capital

This section develops a state-contingent valuation methodology for startups and portfolios of startups (e.g. VC funds and allocations to VC funds). Startup firms are characterized by staged, milestone-driven investment decisions, substantial failure rates, and highly skewed return distributions where a small fraction of firms generate most industry returns through successful exits. These firms face extreme uncertainty, with volatile valuations driven by limited operating history, unproven business models, and rapidly evolving market conditions. Startup firms are commonly described as options. For example, Cochrane (2005) writes, "Venture capital investments are like options; they have a small chance of a huge payoff."

We model individual startup firms as compound call options, where each stage represents an option on the next stage's option. At each decision point, the firm must pay an investment cost (strike price) to continue development and acquire the option on the next stage. The firm continues only when the expected value of future options exceeds the required investment; otherwise, it exits and realizes the current stage's exit value.

To analyze systematic risk, we integrate this compound option framework with a CAPM-style market factor. This integration allows us to examine how market risk factors affect both individual startups and startup portfolios. Portfolio returns and risks are evaluated at each development stage, conditional on market return realizations, where a portfolio is simply a combination of startup firms.

A key insight from the option-based model is that an early-stage startup represents a levered claim on its underlying asset value. This leverage arises naturally from the staged funding structure, where each investment acts as an option on future stages. If the underlying asset value has exposure to systematic risk (beta), this exposure gets magnified through the leverage inherent in the staged funding structure. Thus, even when most of the startup's return variation is idiosyncratic, its systematic risk exposure can substantially exceed that of newly public firms - despite both representing equity claims on similar underlying assets.

The share of total risk that is systematic increases as startup firms are combined into larger portfolios like VC funds and portfolios of VC funds like the large allocations made by many college endowments. A key insight from the structural model of startup portfolios is that continuation (or equivalently, liquidation) probabilities and expected returns are smooth functions of the realized market return that can be neatly characterized. The model demonstrates that with sufficiently high beta in the underlying assets, expected VC returns are convex in market returns.

2.1 Startup Firms as Compound Options

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 $\tilde{V}_{T_2}$. This terminal payoff represents the underlying asset value on which the startup's structure serves as a derivative claim. The investment amounts and their timing are pre-specified, making the optimal funding and continuation decisions dependent solely on the realized underlying asset values. The underlying asset pays no dividends.

The investment process proceeds in three stages. The founder and early-stage investor group must invest K_0 to start the firm in order to proceed to stage 1 of the firm's development. At the end of stage 1, it must be decided whether to invest K_1 so that the firm can proceed to stage 2 or the firm is liquidated for zero. At the end of stage 2, it must be decided whether to invest K_2 to proceed to the final stage where the firm has a chance for a "high success" outcome like an IPO, or whether to forgo this last investment and exit via a mezzanine VC investors, respectively.

The startup value is a pre-money (i.e., before fund raising) object, which differs from the present value of the underlying asset defined as a post-money object. A natural comparison is the startup value to the present value of the underlying asset minus the present value of the series of investments, where the gap between these values measures the value of being able to stage the funding/continuation decisions.

The startup firm's value resembles a compound option, specifically a call option on a call option. If the first-stage call option is not exercised, the second (embedded) option expires worthless. (Geske,

1979) develops a pricing method for compound options and a formula for a call-on-call option. In this section, the Geske model is modified to accommodate a modified stage 2 payoff that allows for partial recovery in the event that the startup does not achieve a "high success" outcome, to better match the outcomes of actual late stage startups where zero value exits are relatively rare.

2.2 Modified Geske (1979) Compound Option Formula

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.

Stage 1: Intermediate Decision

At the first decision point (T_1), the holder of the compound option decides whether to exercise it. This occurs if the value of the underlying option or portfolio at T_1 exceeds the compound option's strike price:

$$V(T_1) > K_1.$$

If this condition is satisfied, the holder pays K_1 to acquire the underlying asset or option, and the compound option continues to the second stage. Otherwise, the compound option expires worthless.

Stage 2: Final Payoff

If the compound option is exercised at T_1, the final payoff depends on the embedded option's value at its maturity (T_2). For a call-on-call, this occurs when the underlying asset price exceeds the embedded option's strike price K_2.

Joint Behavior

The compound option's value at time $t = 0$ reflects the risk-neutral probabilities of satisfying both conditions:

- The intermediate condition: $V(T_1) > K_1$,
- The final condition: the embedded option is in-the-money at T_2 when $V(T_2) > K_2$.

Formula

The Geske model combines these probabilities into a two-dimensional framework. The compound option price is given by:

$$C_{\text{compound}} = S \cdot N_2(a_+, b_+, \sqrt{\tau_1/\tau_2}) - K_2 \cdot e^{-r \cdot T_2} \cdot N_2(a_-, b_-, \sqrt{\tau_1/\tau_2}) - K_1 \cdot e^{-r \cdot T_1} \cdot N(a_-),$$

where:

- $N_2(h,k,\rho)$ is the bivariate cumulative normal distribution function with h and k as upper integral limits and a correlation coefficient, ρ,
- $\rho = \sqrt{\tau_1/\tau_2}$ 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 $\tau_2 = T_2 - T_0$,
- $N(h)$ is the cumulative standard normal distribution function,
- $a_+ = \frac{ln(S/S^*)+(r+0.5\sigma_V^2)\tau_1}{\sigma_V\sqrt{\tau_1}}$,
- $a_- = a_+ - \sigma_V\sqrt{\tau_1}$,
- $b_+ = \frac{ln(S/K_2)+(r+0.5\sigma_V^2)\tau_2}{\sigma_V\sqrt{\tau_2}}$,
- $b_- = b_+ - \sigma_V\sqrt{\tau_2}$,
- S^* is the minimum underlying asset value at which it is optimal to exercise the option at T_1.

In this model, $N_2(a_+, b_+, \sqrt{\tau_1/\tau_2})$ is the expected terminal value of the underlying, given the joint probability of being in-the-money at T_1 and T_2; $N_2(a_-, b_-, \sqrt{\tau_1/\tau_2})$ is the joint probability of being in-the-money at T_1 and T_2; and $N(a_-)$ is the probability of being in-the-money for the intermediate decision at T_1. This approach accounts for the timing and interdependence of the two decision points, ensuring the compound option value reflects the underlying risk-neutral dynamics.

Adding a Recovery Value to the Final Payoff

Many startups liquidate for zero, as captured in the intermediate decision at the end of stage 1. Conditional on advancing to a later stage, a relatively small share of VC-backed startups exit via IPO, with many being sold to operating companies, potentially for a lower value than the cumulative paid in capital. We modify the Geske compound option formula to accommodate this final stage payoff structure.

At the end of the first stage, K_1 is invested or the startup is liquidated for zero, as before. At the end of the second stage, the startup can choose to invest K_2 and proceed to IPO. Otherwise, if the time 2 investment is not made the firm is sold for $\phi \cdot S(T_2)$, where $\phi \in [0,1]$ is a recovery rate. This modification sets the stage 2 payoff to be the greater of two components: the standard call option payoff or a proportion of the underlying asset's value. For a compound call-on-call with recovery, the embedded option's value at T_2 is adjusted to:

$$\max(S_{T_2} - K_2, \phi \cdot S_{T_2}).$$

This recovery structure impacts the replicating portfolio's value at the intermediate decision point T_1. The payoff can be replicated with a combination of the underlying and a standard call option with a modified strike price. The portfolio value becomes:

$$V(T_1) = \phi \cdot S_{T_1} + (1 - \phi) \cdot C(S_{T_1}, K_2', T_2 - T_1),$$

where $K_2' = K_2/(1 - \phi)$ adjusts the strike to reflect the recovery condition.

Figure 1 illustrates the payoff function of an option with strike price $K = 100$ and a recovery rate $\phi = 0.2$. The payoff is calculated as $\max(S_T - K, \phi \cdot S_T)$, where S_T represents the stock price at maturity. The figure highlights two key points, including the strike price (K) and the effective strike price ($K' = \frac{K}{1-\phi}$).

Figure 1: Payout Profile

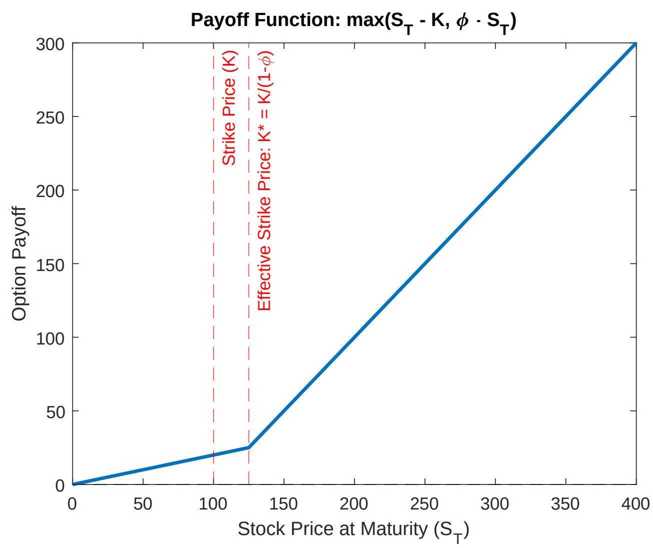

Note: This figure illustrates the payout profile of a call option with recovery.

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. The value of the startup firm is:

$$
\begin{aligned}
C_0 = \phi \cdot S \cdot N(a_+) + (1-\phi) \cdot S \cdot N_2(a_+, b_+, \sqrt{\tau_1/\tau_2}) \\
- (1-\phi) \cdot K_2' \cdot e^{-r \cdot T2} \cdot N_2(a_-, b_-, \sqrt{\tau_1/\tau_2}) \\
- K_1 \cdot e^{-r \cdot T_1} \cdot N(a_-),
\end{aligned}
\tag{1}
$$

where $N(a_+)$ is the expected value of the underlying asset at T_1, conditional on being in-the-money (i.e. $S_{T_1} > S^*$). This framework maintains the flexibility of the Geske model, while accounting for partial recovery in the final stage.

Unconditional Model Properties

There are several model properties of interest besides the startup firm value, some of which can be calculated analytically from the startup valuation formula in equation 1. The **unconditional 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:

$$
Pr(Liquidation_{T_1}) = 1 - N(a_-). \tag{2}
$$

The **unconditional probability of IPO** is the joint probability of being in-the-money at T_1 and T_2:

$$
Pr(IPO) = N_2(a_-, b_-, \sqrt{\tau_1/\tau_2}). \tag{3}
$$

Both of these risk-neutral probabilities can be converted to objective (P-measure) probabilities by adding an equilibrium risk premium to the drift of the underlying asset, which we do when calibrating the model to startup exit rate data.

An additional model-based calculation is the **value of staged investment**. The startup value represents the present value of the new venture, given the option to exit without making future investments. We define the value of staged investment to be the difference between the startup value in equation 1 and the net present value of the underlying post-money asset value along with mandatory future investments:

$$
\text{Value of Staged Investment} = C_0 - (S_0 - K_1 \cdot e^{-r \cdot \tau_1} - K_2 \cdot e^{-r \cdot \tau_2}). \tag{4}
$$

2.3 Integrating the CAPM into the Startup Valuation

To analyze systematic risk, we integrate this 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 approach follows the spirit of Coval, Jurek, and Stafford (2009) (CJS) who integrate the CAPM into the Merton (1974) credit model to study the pricing of bond portfolios and CDO tranches. CJS use the CAPM-integrated structural model to map the nonlinear systematic risk exposure of corporate credit into the realized market return, which can then be priced via no arbitrage and index option portfolios.

The underlying asset for a startup is a post-money object with returns driven by a combination of idiosyncratic, $\tilde{Z}_{i,\epsilon}$, and systematic shocks, $\tilde{Z}_m$. Conditional on the realization of the market factor, idiosyncratic returns are assumed to be Gaussian (i.e. $\tilde{Z}_{i,\epsilon}$ is normally distributed). The terminal asset value, conditional on a specific market realization is:

$$\tilde{A}_{i,T}(m_\tau) = A_{i,t}e^{\mu_A \tau + \beta_A m_\tau + \sigma_\epsilon \sqrt{\tau}\tilde{Z}_{i,\epsilon}}, \tag{5}$$

where $\mu_A = r - \sigma_A^2/2$, $\tau = T - t$, $m(\tau)$ is the log market return, β_A is the asset CAPM beta, and σ_ϵ is the idiosyncratic asset volatility. Equation 5 is agnostic about the market return distribution, but it is useful to assume the log market return is normally distributed when making comparisons to known results of the Black-Scholes-Merton option model. With the log market return assumption, the value of the initial startup firm is not altered by the CAPM factor, holding asset volatility constant, as the market factor is simply decomposing total risk into systematic and idiosyncratic components. Likewise, unconditional startup firm exit rates and return distributions are not affected by the market factor. The market factor allows us to characterize expected payoffs, returns, and exit rates conditional on market return realizations.

The startup beta can be calculated from the formula for an option beta, as a function of the underlying asset beta and model outputs,

$$\beta_{Startup,t} = \beta_A \cdot \Delta_t \cdot \frac{A_t}{C_t} \tag{6a}$$

$$\Delta_t = \phi \cdot N(a_+) + (1-\phi) \cdot N_2(a_+, b_+, \sqrt{\tau_1/\tau_2}). \tag{6b}$$

2.4 Illustrating the Model

To set ideas, consider the model parameter assumptions listed in Table 1. The initial underlying asset value is $A_{i,T_0} = 100$, and represents the present value of the terminal payoff distribution ignoring all required investments. The two strike prices of $K_1 = 30$ and $K_2 = 100$ are required to obtain the terminal distribution, and clearly have a present value that exceeds the initial underlying asset, such that this is a negative NPV investment if these investments cannot be abandoned.

However, the option to abandon creates substantial value. The initial startup value is $C_0 = 53.7$, reflecting the value created by the staged investment structure. The startup's initial beta of 2.6 significantly exceeds the underlying asset beta of 1.8, demonstrating the leverage effect created by staged investments. Under the physical measure P, the model implies a 53.4% probability of liquidation at the end of stage 1 and a 19.9% probability of achieving an IPO.

Table 1: Summary of Model Inputs and Outputs

Model Inputs	Value	Model Outputs	Value
Underlying value	100.0	Startup Value	53.7
K1	30.0	Startup Beta	2.6
K2	100.0	Prob Liquidate at 1	58.5
Recovery	20.0	Prob Liquidate (P)	53.4
K2_prime	125.0	Prob IPO at 2	15.4
T1	2.0	Prob IPO at 2 (P)	19.9
T2	4.0	Critical Value, A*	64.5
Market Vol	14.8		
Underlying Vol	90.0		
Underlying Beta	1.8		
Riskfree Rate	5.0		
Market Risk Premium	5.0		
Asset Risk Premium	9.0		

Note: This figure summarizes the model inputs and outputs.

Table 2 illustrates the compound option model using three identical startups under the assumed parameters. Initially, all firms share the same underlying asset values, required investments (strike prices), and therefore the same startup values. By time 1, the firms' underlying asset values have diverged significantly, leading to different option values and continuation decisions. Firm 1 has a post-money continuation value of 20.8, which is below the required $K_1 = 30$ investment, so liquidation is optimal. Firms 2 and 3 have continuation values exceeding K_1, making further investment optimal. At time 2, firms face a choice between selling or investing $K_2 = 100$ to pursue an IPO, with the decision governed by an effective strike price of $K_2^* = 125$. Firm 2 has an underlying asset value of 120, making an asset sale optimal, yielding 20% of the terminal value without requiring the K_2 investment, generating a payoff of 24. Firm 3 has a higher asset value, so investing $K_2 = 100$ to complete an IPO is optimal, generating a net payoff of 275.

Table 2: Model Example

	T0	T1	T2
Strike Prices			
Firm 1	0.0	30.0	100.0
Firm 2	0.0	30.0	100.0
Firm 3	0.0	30.0	100.0
Underlying Asset Values			
Firm 1	100.0	50.0	120.0
Firm 2	100.0	70.0	120.0
Firm 3	100.0	190.0	375.0
Post-Money Startup Values			
Firm 1	53.7	20.8	0.0
Firm 2	53.7	33.7	24.0
Firm 3	53.7	130.4	375.0
Pre-Money Startup Values			
Firm 1	53.7	0.0	0.0
Firm 2	53.7	3.7	24.0
Firm 3	53.7	100.4	275.0

Note: This tables illustrate how the model works with three example paths of how the underlying asset and the startup value may evolve.

Consider an early-stage venture capital investor who initially owns identical positions in all three startups. This investor maintains their original ownership stakes through both development stages, but does not provide the capital for either the time 1 (K_1) or time 2 (K_2) investments. The investor's returns evolve across two stages. The stage 1 return is calculated as the sum of the pre-money startup values at time 1, divided by the sum of the three initial startup values. At time 1, one firm liquidates with zero value, while two firms continue after raising K_1 each from new investors, for a total capital raise of $2K_1$. The VC's stage 2 return is then calculated as the sum of the pre-money startup values at time 2, divided by the sum of the post-money startup values at time 1. This return calculation reflects that the time 1 investment capital comes from new investors rather than the early-stage VC. In this example, the early-stage VC initially invests 161 in the three startups, which are worth only 104 at the end of stage 1, realizing a stage 1 return of -35%. Because the early-stage VC does not provide the time 1 investment capital, their ownership stake is reduced in each startup. We assume that the new capital provider's share of the startup is K_1/Post-Money Value. In the second stage, the early-stage VC earns a 62% return, providing a total return multiple of just 1.05x over the full vintage life.

The startup valuation formula can be applied in both unconditional and state-contingent frameworks. In the state-contingent framework, we decompose asset return variance into systematic (market-related) and idiosyncratic (firm-specific) components. Importantly, this decomposition preserves the total startup value when the total asset volatility remains constant and log market returns are normally distributed (as in the BSM and Geske models). This equivalence means we can use the unconditional valuation formula (equation 1) alongside state-contingent analyses.

Stage 1 Model Properties Conditional on Market Returns

The state-contingent properties are most easily illustrated through the first stage outcomes, shown in Figure 2. The x-axis in all plots is "moneyness", defined as the realized market index at the end of stage 1 divided by the market index at the beginning of stage 1 (i.e. the gross return). The rows of this figure present the:

1. distribution of market returns ('moneyness');
2. stage 1 conditional payoffs across different market states;
3. stage 1 conditional probability of continuation vs. exit;
4. stage 1 conditional expected gross return for initial investors.

Figure 2: Model Illustration of Early-stage

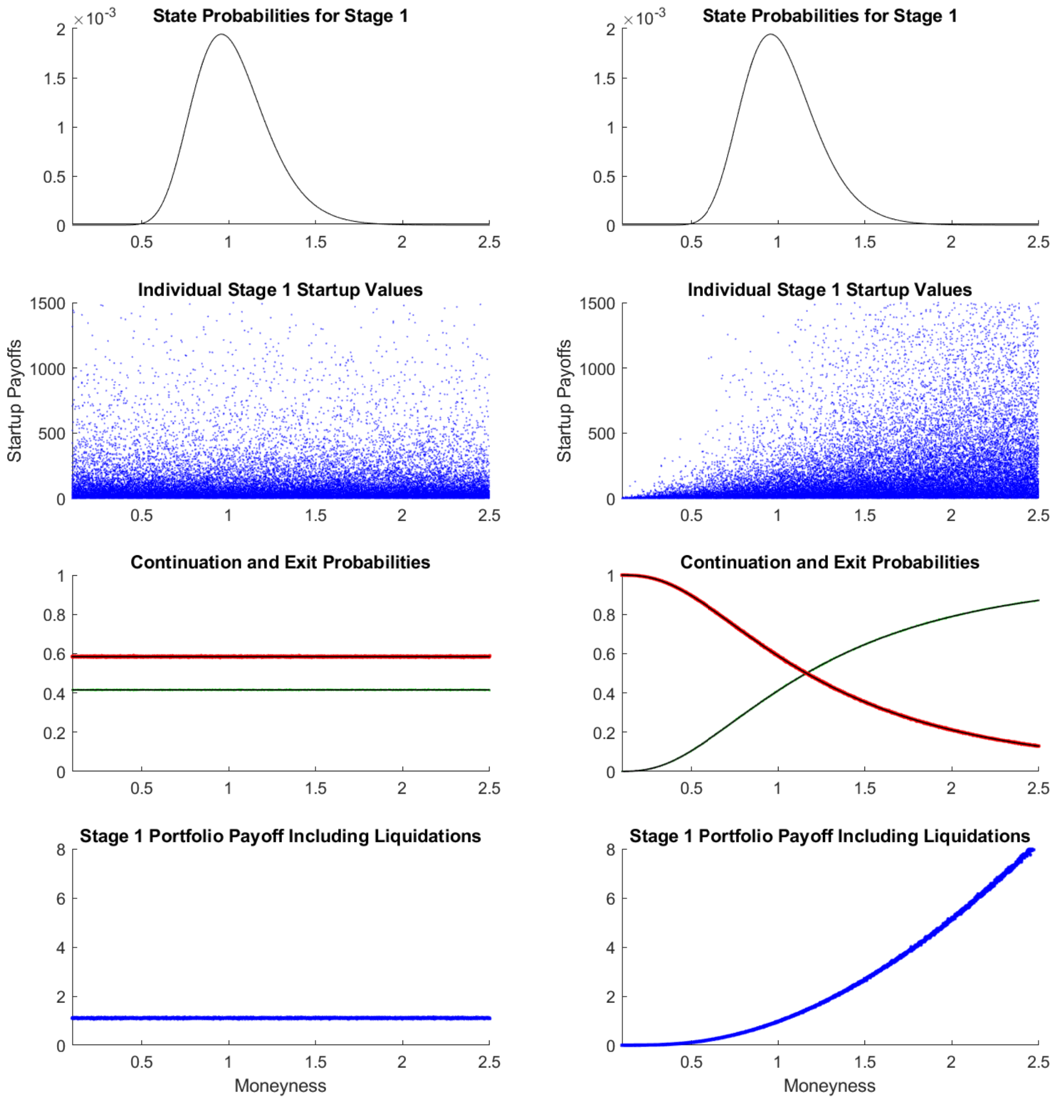

Note: This figures illustrates the state contingency of the first-stage outcomes. The underlying asset has a zero beta on the left-hand side and a positive beta on the right-hand side.

For each analysis, we compare two scenarios: startups with zero systematic risk ($\beta_A = 0$) and high systematic risk ($\beta_A = 1.8$). This comparison reveals how systematic risk affects startup and portfolio outcomes across different market realizations.

The top row shows the probability distribution of market returns (moneyness), defined as the ratio of the market index at the end of stage 1 to its initial value. This distribution $f(m)$ assumes normally distributed log market returns, using the capital market parameters from Table 1. Since this is purely a market distribution, it is identical across both columns ($\beta_A = 0$ and $\beta_A = 1.8$). The remaining analysis draws on simulations of 10,000 startup firms.

The second row visualizes potential startup values at the end of stage 1. For each market return level ('moneyness'), we use equation (5) to simulate underlying asset values. Each asset value determines an embedded option value - specifically, a one-period call option with time to maturity: $(T_2 - T_1)$, strike price: K_2, and the recovery value guarantee (see Appendix for details of this calculation). We denote these simulated values as $C_i^*(m_{T_1})$. The dots in the plots represent these values before considering whether continuation is optimal (i.e., before the stage 1 funding decision). 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 examines continuation (and liquidation) probabilities. A startup continues at stage 1 if its embedded option value exceeds the required investment: $C_i^*(m_{T_1}) > K_1$. For each market return level, we calculate a simulated continuation probability as $Pr_{sim}(Liquidation(m_{T_1}))$. The simulated conditional continuation probability is the fraction of the 10,000 startups where this condition is satisfied, with the remaining share representing exit/liquidation. The plot shows that the conditional continuation probability can be strongly increasing in the realized market return when there is systematic risk in the underlying asset. The unconditional simulated liquidation probability,

$$Pr_{sim}(Continuation) = \int_0^\infty Pr_{sim}(Continuation(m_{T_1}))f(m_{T_1})\, dm_{T_1}, \tag{7}$$

matches what is calculated with equation 2. The conditional continuation probability can also be calculated analytically:

$$Pr(Continuation(m_{\tau_1})) = N(\eta(m_{\tau_1})) \tag{8a}$$

$$\eta(m_{\tau_1}) = \frac{ln(A_{T_0}/A^*_{\tau_1}) + (\mu_A - \frac{1}{2}\sigma_A^2)\tau_1 + \beta_A ln(m_{\tau_1})}{\sigma_\epsilon \sqrt{\tau_1}}, \tag{8b}$$

where A^* corresponds to the critical underlying asset value, S^*, where continuation is optimal in the unconditional startup valuation model in equation (1). This is useful for verifying the accuracy of the simulation. The plots reveal a key prediction: when startups have positive beta, their continuation probability increases strongly with market returns. This market-dependent survival pattern is testable using startup investment data. For zero-beta startups, continuation probability is independent of market conditions.

The bottom row analyzes expected returns for initial investors (founders and early-stage VCs) over stage 1. The analysis proceeds in two steps. First, we calculate the startup's value after the continuation decision: $C_i(m_{T_1}) = \max[0, C_i^*(m_{T_1}) - K_1]$. Then, we compute the conditional expected return for each market state:

$$\tilde{R}(m_{T_1}) = \frac{\mathbb{E}_i[\tilde{C}_{i,T_1}(m_{T_1})]}{C_{T_0}} \tag{9}$$

where: $\mathbb{E}_i$ averages across startups at a fixed market return m_{T_1}.

Similarly, equation 9 can be used to calculate portfolio returns by summing across startups. Equation (9 is important for determining how to evaluate and replicate the systematic risks of **early stage** startups and allocations to **early stage** venture capital funds. For startups with sufficiently positive beta, returns are convex in market returns. A neat feature of the state-contingent model is that the nonlinear portfolio payoff can be practically approximated with an index option replicating portfolio.

Stage 2 Model Properties Conditional on Market Returns

Characterizing the stage 2 properties is more complex due to two factors: the path dependence inherent in compound options and the factor structure in the underlying asset return. Figure 3 illustrates key stage 2 model properties, comparing outcomes under two stage 1 market scenarios: poor market returns ($m_{T_1} = 0.7$) and strongly positive returns ($m_{T_1} = 1.5$). The figure presents:

1. distribution of startup values at the beginning of stage 2
2. stage 2 conditional portfolio return as a multiple of the initial time T_1 investment
3. stage 2 conditional portfolio return as a multiple of the beginning of stage 2 value
4. stage 2 conditional IPO probability

The top row reveals how market conditions affect continuation rates and startup values. When the market is up 50%, 65% of startups continue to stage 2, compared to only 23% when the market is down 30%. Moreover, the value of continuing startups is 6.5 times higher under strong market performance versus poor performance. This continuation decision is determined by comparing each startup's value to the required stage 2 investment K_1 (shown by the vertical line in the plots).

The middle rows examine portfolio returns from two perspectives: scaled by the initial time T_0 portfolio value (row 2) and by the post-continuation portfolio value at the beginning of stage 2 (row 3). These two scalings reveal (1) stage 1 market returns significantly impact cumulative performance, and (2) stage 2 conditional returns show similar patterns when normalized by stage 1 values.

The bottom row shows that IPO probabilities, conditional on reaching stage 2, are substantially higher than the ex ante probabilities. These probabilities maintain a strong positive relationship with stage 2 market returns, providing a testable implication of the model.

Figure 3: Model Illustration of Late-stage

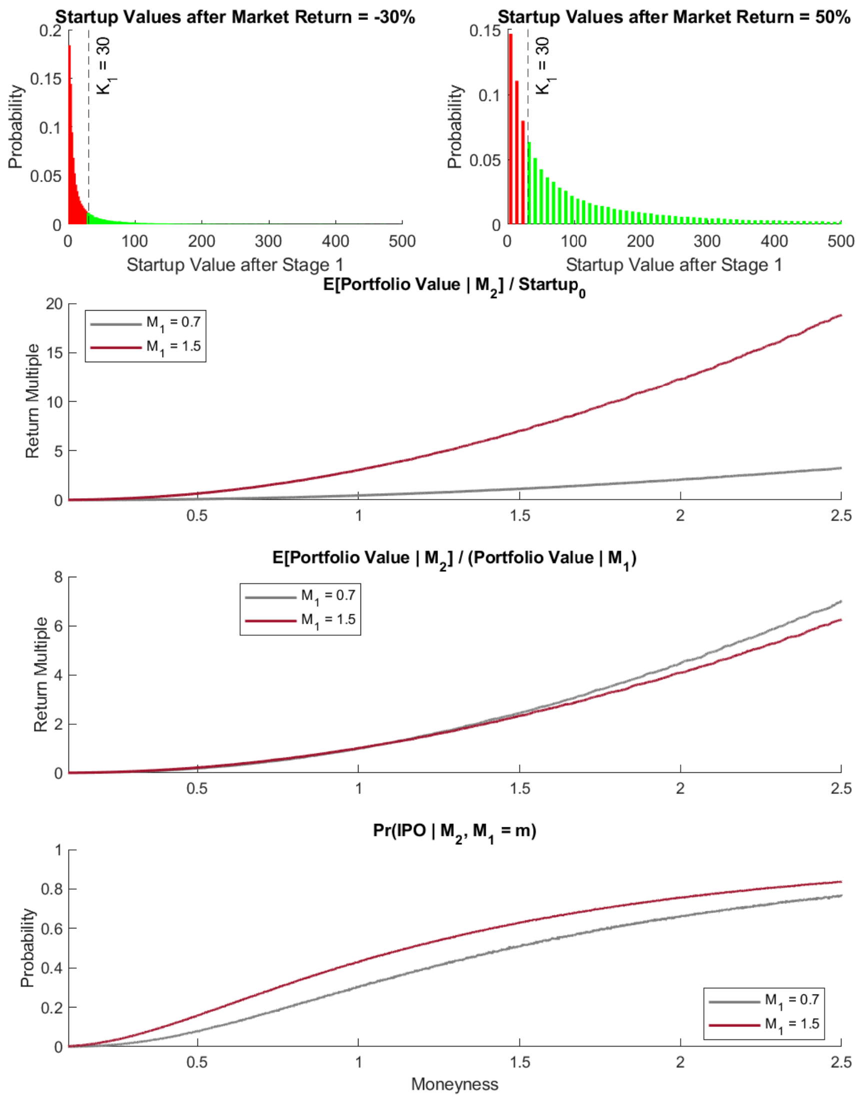

Note: The figures illustrates key model properties of the second stage, comparing outcomes under two scenarios for the first stage: poor market returns ($m_{T_1} = 0.7$) and strongly positive returns ($m_{T_1} = 1.5$).

3 Data Description

This section describes the data used to calibrate and test our model. The calibration relies on unconditional startup exit probabilities. We then test the model's predictions about conditional IPO rates using a time series of VC-backed IPOs, and evaluate predicted returns using VC fund vintage performance data. Our analysis covers VC investments from 1985 to 2016, tracking outcomes through 2021, to allow funds enough time to exit their investments.

We use two market indices, S&P 500 (SPX) and Nasdaq 100 (NDX), with total return indices obtained from Bloomberg. While SPX comprises the 500 largest U.S.-listed companies across exchanges, NDX consists of the 100 largest non-financial Nasdaq-listed firms, providing greater technology exposure. The behavior of stocks in these two indices is different, as shown by 4. The valuation of Nasdaq stocks is convex in the valuation of S&P 500 stocks. This may reflect Nasdaq's heavier weight on younger firms with growth opportunities. These embedded growth opportunities could introduce convexity, akin to how option-like structures shapes VC returns. The risk-free rate is the 4-year constant-maturity Treasury yield, obtained from the website of the Federal Reserve.

3.1 Venture Capital

Fund and portfolio company data from Cambridge Associates. Cambridge Associates (CA) provides quarterly reports containing the Cambridge Associates US Venture Capital Index and selected benchmark statistics. CA obtains the fund and investment-level performance information from the quarterly unaudited and annual audited financial statements produced by the fund managers (GPs) for their Limited Partners (LPs). These documents are provided to Cambridge Associates by the GPs themselves.[3] CA acknowledges the potential survivorship bias in their data but asserts that it is small.[4]

An advantage of the CA data is its provision of an extended time series with consistent inclusion criteria. CA includes "institutional quality" funds that tend to meet the following criteria: closed-end funds, commingled funds that invest third-party capital (we exclude firms that invest off of their balance sheet, such as a bank's principal investing group or a corporate's venture capital arm), and fund vehicles. The time-series consistency of CA is particularly noteworthy since the coverage of many VC data sets often changes substantially over time, making it hard to obtain a consistent time series. An extended

[3]Each report contains information on fund vintages over the past roughly three decades. We use several reports to assemble our final data set.

[4]CA states that "over the last ten years, the number of fund managers that stopped reporting to Cambridge Associates before liquidation represented an average of 0.8% (per year) of the total number of funds in the database during the respective year, and an average of 0.6% (per year) as a percentage of total NAV in the database during that respective year." Additionally, they write, "The performance of the small number of funds that have stopped reporting has been spread amongst all quartiles and has not been concentrated consistently in the poorer performing quartiles."

Figure 4: Stock Market Valuation: Nasdaq vs. S&P 500

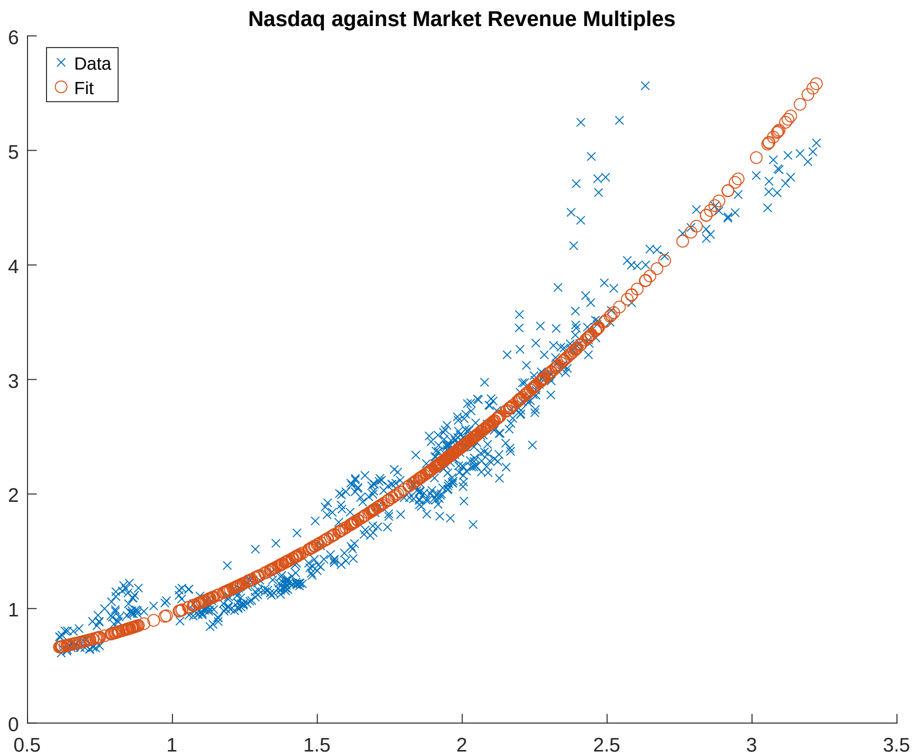

Note: This figures compares the revenue multiples of stocks in the Nasdaq and the S&P 500 indices.

time series is useful when investigating the systematic risk of VC investments.

From the CA reports, we collect the number of funds started in the year, i.e., funds that are included in a "fund vintage", and the performance of these funds as measured by the internal rate of return (IRR) and the total value to paid-in capital (TVPI). The performance measures are computed based on the pooled cash flows of all funds in each vintage, net of fees, expenses, and carried interest. From these two performance measures, we estimate the average holding period by $log(TVPI)/log(1 + IRR)$.

Figure 5 compares the vintage IRRs with the horizon-matched annualized returns of the S&P 500 (Panel A) and the Nasdaq 100 (Panel B). The figure highlights two key points. First, VC IRRs are highly related to the returns of stock market indices. The fitted curve, obtained from regressing the VC IRRs on market returns and squared market returns, exhibits a slope far above one across most of the figure. For example, the 1997 VC vintage achieved an annualized return of more than 110%, while the Nasdaq 100 and the S&P 500 returned around 50% and 30%, respectively. Thus, the figure is consistent with prior studies finding that the beta of VC is well-above 1 or even above 2 (e.g., Cochrane, 2005; Korteweg and Sorensen, 2010; Korteweg and Nagel, 2024). Second, the figure highlights convexity in VC returns. This is driven by relatively good performance in both the up and down markets around the 2000-era tech boom and crash. As the market return increases, the pooled IRR grows at an accelerating rate. Several vintages (1995, 1996, and 1997) prior to the dot-com bubble have achieved extremely high returns, with the following vintages realizing pooled IRRs were near zero while the market experienced losses. Overall, the figure illustrates a strong, nonlinear relationship between VC and equity market returns.

Figure 5: Venture Capital vs. Stock Market Returns

(A) VC vintages vs Nasdaq-100

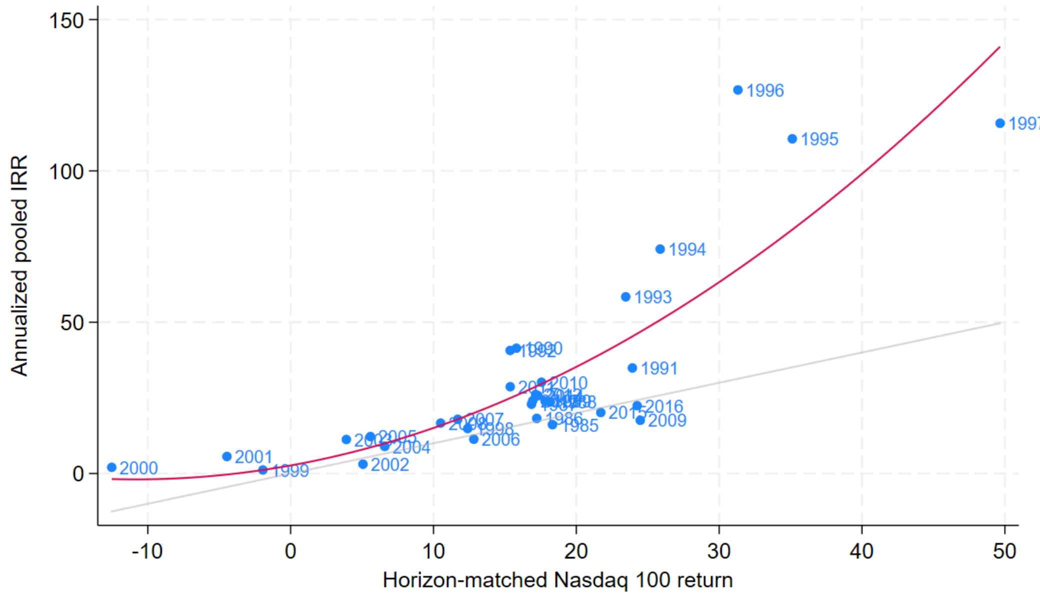

(B) VC vintages vs S&P 500

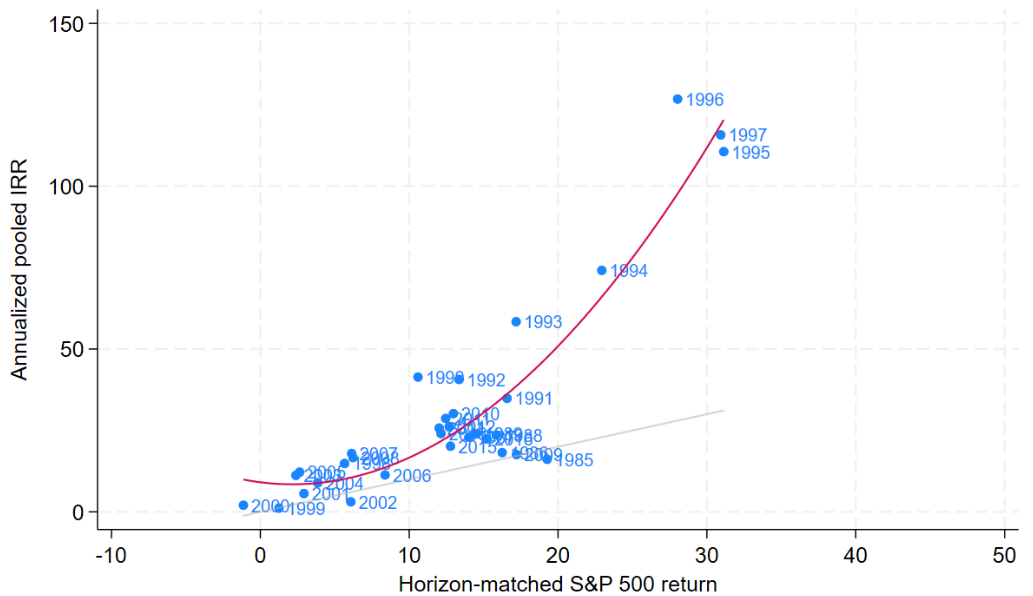

Note: This figure compares the VC vintage returns (IRR) with the returns of the Nasdaq 100 (Panel A) and the S&P 500 (Panel B). The returns of the equity indices are calculated over the holding period of each VC Vintage.

In addition to the fund vintage statistics, CA also reports the number of startup companies that receive their initial investment each year. Since reports contain only the recent history of company funding, we are only able to obtain the company vintage data going back to 1996. We extend this series back in time by multiplying the number of funds launched in each year prior to 1996 by 20, which is the average ratio of portfolio firms and funds between 1996 and 2016.

Number of VC-backed IPOs. A list of U.S. IPOs is obtained from Jay Ritter's website.[5] The data contains the firm name, the IPO date, and an indicator of whether the firm going public was previously backed by Venture Capital. We focus on VC-backed IPOs in the following. The number of VC-backed IPOs closely matches what prior studies report.

Deal-level data from VentureXpert and Pitchbook. Deal-level VC data is obtained from VentureXpert. From the deal-level data, we compute the number of startups that receive initial VC funding each year. The time series of portfolio firms is used to validate the CA data. We find that the two data sets match closely.

Additional deal-level data is obtained from Pitchbook. Pitchbook offers detailed deal-level information on the ultimate exit of each startup. However, its coverage changes substantially over time, becoming more comprehensive after 2000. The following two data filters are imposed. First, only investments in US-headquartered firms are considered. Second, to ensure consistency with other data sets that focus on "institutional quality funds" (CA), we restrict attention to startups that receive at least one million in VC funding. Of these 22,789 startups that received at least one million VC funding between 2001 and 2016, the exits probabilities are as follows: IPO (4.87%), M&A (34.17%), buyout (18.47%), bankruptcy (7.41%), out of business (14.52%), unobserved (10.44%), or still active (10.13%). To compute the unconditional failure probability, we assume that almost all active firms can be considered a failure since they have not been exited after seven years. Accordingly, we estimate the unconditional failure probability at 40% as the approximate sum of the probabilities of the bankruptcy, out-of-business, unobserved, and active outcomes. The unconditional failure probability is one two key empirical moments that we use to calibrate our model.

[5]This 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 Securities Data Company) new issues database.

Table 3: Data Overview

	Interest Rates and Equity Returns			Venture Capital			
Year	4-year risk-free rate	Nasdaq 100	S&P 500	# VC funds (CA)	# VC firms (CA)	# VC firms (VX)	# VC-backed IPOs
1985	9.9	19.3	31.7	26		439	29
1986	7.2	6.9	18.7	30			78
1987	7.8	10.5	5.3	34			64
1988	8.4	13.5	16.6	26		515	37
1989	8.5	26.2	31.7	37		450	35
1990	8.3	-10.4	-3.2	17		355	43
1991	7.1	65.0	30.4	17		259	110
1992	5.7	8.9	7.6	21		390	133
1993	4.8	11.8	10.1	36		350	175
1994	6.5	1.8	1.3	41		421	137
1995	6.3	43.1	37.6	35		882	166
1996	6.1	42.8	22.9	40	1039	1132	239
1997	6.2	20.8	33.3	70	1046	1288	108
1998	5.1	85.5	28.6	81	1517	1751	61
1999	5.5	102.1	21.0	110	2336	2409	264
2000	6.2	-36.8	-9.1	153	3003	3385	249
2001	4.3	-32.6	-11.9	54	1481	724	23
2002	3.5	-37.5	-22.1	33	1258	847	12
2003	2.5	49.4	28.7	38	1255	775	23
2004	3.1	10.7	10.9	67	1416	972	84
2005	4.0	1.9	4.9	64	1320	1038	52
2006	4.8	7.2	15.8	80	1508		58
2007	4.4	19.2	5.6	68	1599		83
2008	2.5	-41.6	-37.0	66	1418		7
2009	1.8	54.6	26.4	23	908	839	14
2010	1.5	20.1	15.1	50	1206	1106	62
2011	1.1	3.6	2.1	45	1510	1333	53
2012	0.6	18.3	16.0	56	1368	1360	48
2013	0.9	36.9	32.4	58	1405	1472	79
2014	1.3	19.4	13.7	81	1522	1511	135
2015	1.3	9.7	1.4	61	1699	1651	79
2016	1.2	7.3	11.9	68	1351	1411	43
2017	1.7	33.0	21.8	101	1526	1609	67
2018	2.7	0.0	-4.4	111	1860	1768	100
2019	1.9	39.5	31.5	104	1929	2037	90
2020	0.5	48.9	18.4	113	2111		119
2021	0.7	27.5	28.7	174	3341		185
Mean	4.2	19.1	13.4	62	1613	1149	90
Stdev	2.7	30.9	16.7	37	559	699	66

Note: This data provides an overview of the data. The Nasdaq Composite index is used instead of the Nasdaq 100 prior to the launch of the Nasdaq 100 index on January 21, 1985.

The second empirical moment for the model calibration is the unconditional IPO probability. To obtain an estimate for the unconditional IPO probability, we divide the number of VC-backed IPOs by the number of startups receiving VC funding. This results in an unconditional IPO probability of 7.2%. Overall, these numbers might appear different from previous studies or even the estimate in the Pitchbook data, but the differences are explained by the use of different sample periods. When focusing on the Pitchbook sample period from 2001 onwards, we obtain an IPO probability of 5.7% quite similar to the IPO probability in the Pitchbook data.

To obtain a time series for the IPO probability, we divide the number of VC-backed IPOs in a month by the number of startups receiving VC funding in an average month over the four years t-2 to t-5. For example, to estimate the IPO probability in January 2021, we divide the number of VC-backed IPOs in January 2021 by 48 times the number of portfolio firms that received funding in the years 2016, 2017, 2018, and 2019.

3.2 Return Characteristics of Newly-listed VC-backed Firms

Since many successful VC-backed firms eventually end up trading on public stock exchanges, we investigate the return dynamics of these stocks. To do so, daily returns of stocks listed on the NYSE, NYSE American (formerly Amex), and Nasdaq are obtained from CRSP. The data is then merged with Jay Ritter's IPO data to obtain the stock returns of newly listed VC-backed firms.

Using daily return data, the Nasdaq 100 beta, the S&P 500 beta, and the total volatility of these stocks are estimated. When estimating the beta, 20 lags of market index returns are included. To obtain annual volatilities, daily volatilities are multiplied by $\sqrt{252}$. We then track the return characteristics over the first five years of trading following an IPO.

Table 4 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 4: 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 6 explores the time series dimension by showing the distribution of return characteristics across all stocks listed in each calendar year. There is no obvious pattern in the stock betas (Panel 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 6: Beta and Volatility of VC-backed newly IPO'd firms

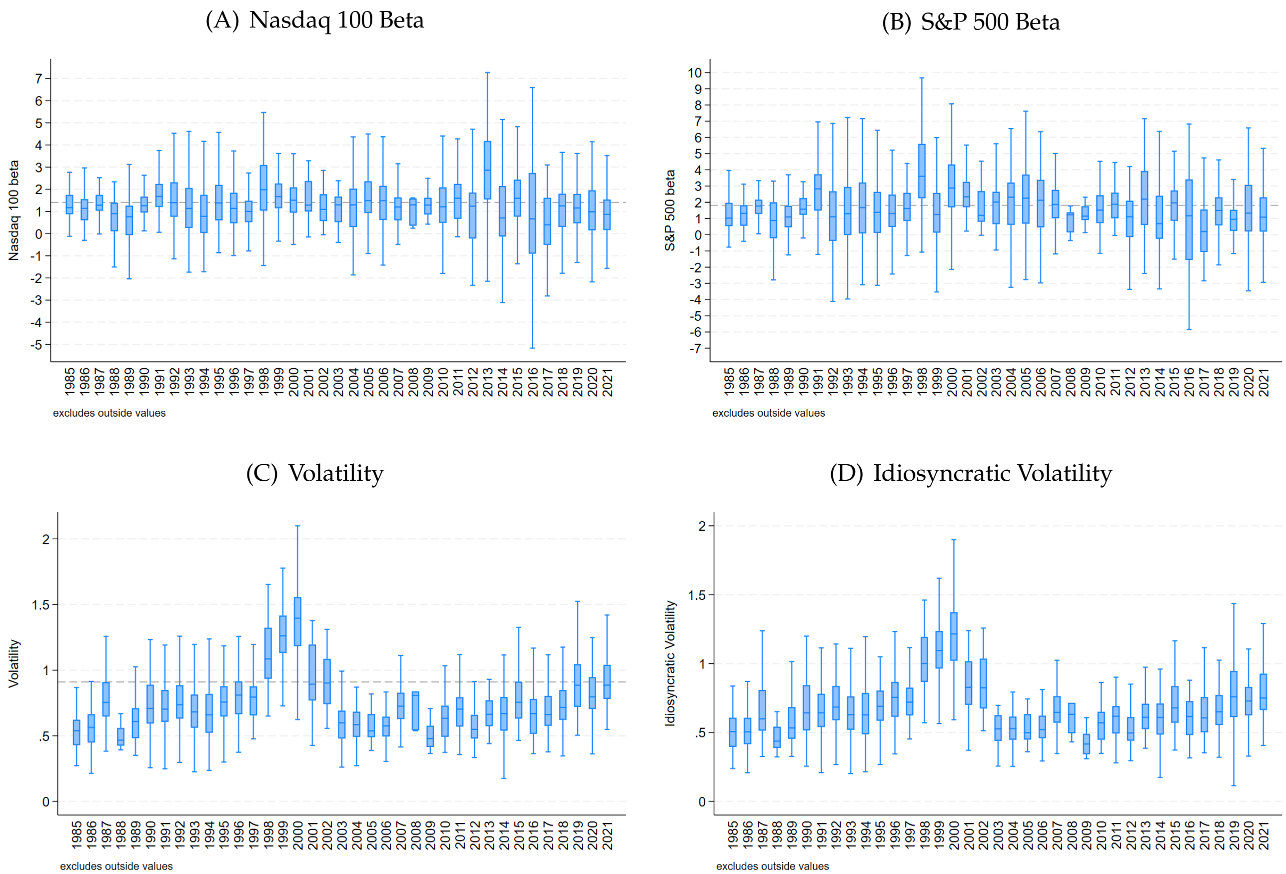

Note: This figure provides information on the return characteristics of newly-listed firms that were previously VC-backed, categorized by their IPO year. Daily return data from CRSP is used in the estimation.

4 Model Calibration and Evaluation

The structural option-based model allows for startup firm valuations (and returns), as well as portfolios of startup firms, to be characterized as functions of realized market returns. Additionally, the model calculates various exit probabilities like liquidation and IPO both unconditionally and conditional on market returns. In this section, we describe how we calibrate the model parameters to match the unconditional startup liquidation and IPO rates and then evaluate how well the model explains IPO and market return dynamics. A second validation exercise involves calculating model-implied VC vintage returns and comparing these with VC vintage returns reported by Cambridge Associates. The calibrated model does not make use of any return data other than unconditional market factor parameters and the underlying asset beta.

4.1 Calibration

Table 5 summarizes how the model parameters are determined. The unconditional IPO probability of 7.2% and the unconditional liquidation probability of 40%, described in Section 3, are the targets to be matched in the calibration exercise. The underlying asset value is 100 and simply operates to normalize startup values. The underlying asset return volatility is 90%, which is the estimated return volatility from recent US VC-backed IPO firms, as described in Section 3. The high return volatility is consistent with other research that finds high startup return volatilities (e.g., Cochrane, 2005; Korteweg and Nagel, 2016). We consider two market indices, the S&P 500 (SPX) and the Nasdaq 100 (NDX). We use each of these indices to measure the beta of the underlying asset, which we proxy as a recent VC-backed IPO firm, as reported earlier. The estimated underlying asset beta against the SPX is 1.8 and is 1.4 against the NDX. We assume a market risk premium for the SPX of 5% and 6.4% for NDX, which generate the same asset risk premium ($\beta_{A,m} \cdot MRP_m$) of 9%. The market risk premia are necessary to convert the risk-neutral exit probabilities to P-measure (or objective) exit probabilities. The 3.6% riskfree rate is the average 4-year yield on US Treasury securities over the sample. The duration of each stage is assumed to be two years, so that $T_1 = 2$ and $T_2 = 4$. These maturities are chosen to match the maturities of long-dated index options that can feasibly be traded in markets during our sample.

Table 5: Model Calibration

Empirical Moments		
Target Pr(IPO)	0.072	
Target Pr(EXIT)	0.400	
Fixed Parameters		
Underlying value	100.0	
T1	2.000	
T2	4.000	
Underlying Vol	0.900	
4yr Riskfree Rate	0.036	
	SPX	NDX
Market Risk Premium	0.050	0.064
Underlying Beta	1.800	1.400
Asset Risk Premium	0.090	0.090
Calibrated Parameters		
Calibrated K1	26.3	26.3
Calibrated K2	160.5	160.5
Calibrated Recovery	0.628	0.628
Model Implied Values		
Startup value	58.1	58.1
Initial startup beta	2.43	1.89

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

This leaves three parameters to be calibrated – the two strike prices and the stage 2 recovery rate. The time 2 strike price, K_1, is especially important for determining the stage 1 exit rate. The time 2 strike price, K_2, and the recovery rate, ϕ, together are especially important for determining the IPO rate (and equivalently the acquisition rate) at the end of stage 2. We use the formulas for the unconditional startup liquidation and IPO rates in equations 2 and 3 adjusted to include the asset risk premium in the underlying asset drift term. Since the asset risk premium and the asset volatility are the same for both considered indices, there is just one set of startup-specific calibrated parameters. We find that $K_1 = 26.3$, $K_2 = 160.5.3$, and $\phi = 0.628$.

With these parameters, the initial value of a startup is 58.1 using both indices. Using equation 6a, the startup initially has a beta against the SPX of 2.43 and a beta against the NDX of 1.89. The high startup betas relative to the underlying asset beta is an expected feature of the model viewing the startup structure as a call-like derivative on the underlying asset.

4.2 Startup Exit and Market Dynamics

The structural framework makes predictions about model-implied exit probabilities conditional on market returns. Our first analysis compares the time series of empirical startup IPO rates described in Section 3 with a time series of model-implied IPO rates, conditional on recent realized market returns.

Each quarter from 1990 through 2024, we predict the conditional IPO rate for a portfolio or pool of startups that initially received VC-funding 4 years prior. We simulate underlying asset values for 10,000 startups from the calibrated model. The startups progress through stage 1, conditional on what was the actual realized index return (using both the SPX and NDX) over stage 1 (i.e. the period from 4-years before quarter t to 2-years before quarter t). Using the optimal continuation rule at the end of stage 1, a portion of the initial startup pool continues to stage 2. At this point, there is a wide distribution of startup values due to the substantial idiosyncratic risk in individual startup returns, as shown in Figure 3. Then, we simulate terminal underlying asset values for each startup, conditional on what was the actual realized index return in the 2-year period leading up to the current quarter. Applying the optimal stage 2 decision rule, firms with terminal underlying asset values exceeding the modified strike price, K_2^*, invest K_2 to proceed to an IPO, with the remaining firms skipping the time 2 investment to be acquired.[6] We calculate the conditional IPO probability for this quarter as the fraction of the initial IPO pool that optimally choose their decisions at each stage and eventually choose to IPO.

Figure 7 plots the time series of empirical and model-implied IPO probabilities. The overall time series pattern in the empirical IPO probabilities is well-matched by the model-implied series using either the SPX or NDX. The model that conditions on the NDX index returns is better able to match the high IPO rates around the 2000 time period, but otherwise the time series patterns are similar. There are two striking features of the empirical time series. First is the high IPO rate sustained through the 1990s, where there are relatively few VC-backed startup firms. Second is the relatively low IPO rate since the 2001-2002 crash in tech stock prices. The conditional model is unable to match the high IPO rates in the very early sample period when market indices perform well, but not well enough to expect such high IPO rates. The model predicts IPOs during the tech run-up and since then quite well.

[6] The model makes additional predictions about startup returns that we have not yet tested, including predictions about return distributions conditional on both market realizations and exit types.

Figure 7: Conditional IPO probability: Model vs. Data

(A) Time-series of IPO probability

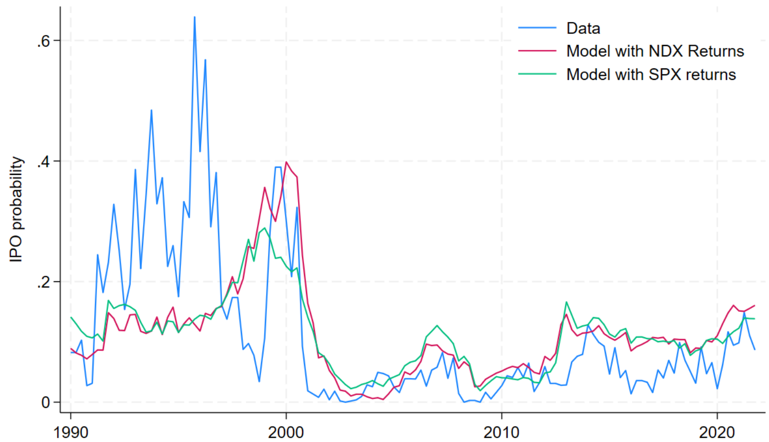

(B) IPO probability vs. Nasdaq-100

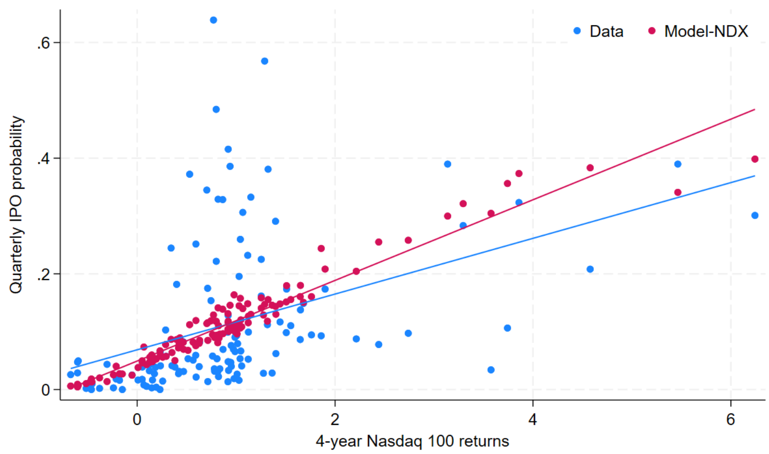

(C) IPO probability vs. S&P 500

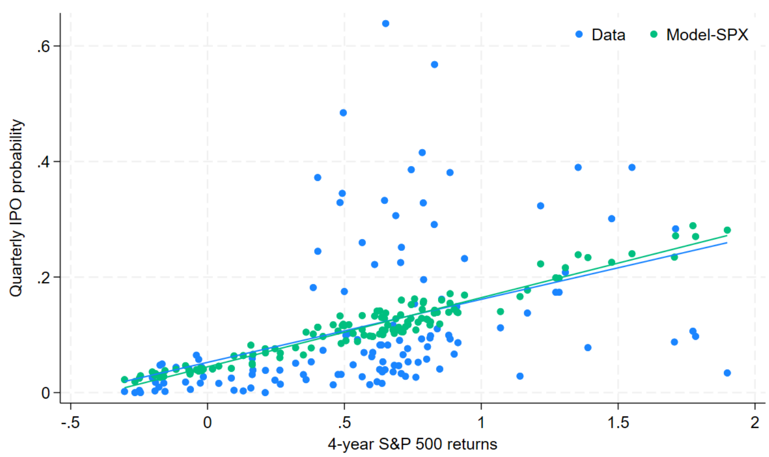

Note: This figure evaluates how well the model explains the empirical pattern of IPO probabilities. Panel A compares the time series of data and model IPO probabilities over time. The empirical IPO probability is calculated as described in Section 3. Panel B provides binscatters of IPO probabilities on the returns of the Nasdaq 100 over the past 4 years. Panel C uses returns of the S&P500.

A related analysis involves regressing quarterly empirical and model-implied IPO rates on recent market index returns to compare these sensitivities, which is a central prediction of the state-contingent model. In particular, the beta of the underlying asset links the startup firm's IPO probability to the recent performance of the market index return, as shown in Figure 3. Table 6 reports the results from these regressions with SPX realized returns shown in Panel A and NDX realized returns in Panel B. In regressions of the empirical IPO rates on recent 4-year market index returns, the slope coefficients are reliably positive and the R^2 is around 0.16 for both indices. The model relying on the market factor being calibrated to the SPX performs slightly better in terms of matching the empirical slope coefficients.

Table 6: IPO probability vs. Market Returns

	Quarterly IPO probability					
	(1)	(2)	(3)	(4)	(5)	(6)
	Data	Model-SPX	Model-NDX	Data	Model-SPX	Model-NDX
4-year Nasdaq 100 return	0.0482**	0.0465***	0.0697***			
	(2.55)	(8.29)	(26.35)			
4-year S&P 500 return				0.109**	0.120***	0.143***
				(2.22)	(16.48)	(10.92)
R^2	0.169	0.720	0.931	0.170	0.935	0.769
N	128	128	128	128	128	128

Note: Regressions of IPO probability on market returns.

Table 7: IPO probability: Model vs. Data

Sample Period	Model	Coefficient	t-Statistic	Adj R^2
1990-2021	Model with SPX Returns	1.0680	6.4780	0.2439
1990-2021	Model with NDX Returns	0.7832	6.1906	0.2271
2003-2021	Model with SPX Returns	0.5197	6.6686	0.3669
2003-2021	Model with NDX Returns	0.5232	7.4667	0.4220

Note: This table shows regression of the empirical IPO probabilities on the model-implied IPO probabilities.

We view these results to be 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.

4.3 VC Vintage Returns

A second model evaluation exercise involves constructing model-implied VC vintage returns and comparing these with pooled VC fund returns organized by vintage calculated by Cambridge Associates (CA). We consider the period from 1985 through 2016, to match the annual vintages reported by CA where there is enough time for the VC funds to have mostly liquidated their investments. The construction of the model-implied vintage returns proceeds similarly to the previous time series analysis, where we perform a simulation for a new pool of startups at each vintage start date. We calculate the returns to an early-stage investor who holds their investment through both stages 1 and 2, but does not add to their investments that continue into stage 2. This calculation follows the example described in Section 2.4.

Specifically, at the beginning of each year we simulate the underlying assets for a pool of 10,000 startups that proceed through stage 1, conditional on the index return realized over the next 2 years using both the SPX and NDX. Again, based on the optimal continuation rule at the end of stage 1, we determine which individual startups proceed to stage 2 and which are liquidated for zero. Then, we simulate values at the end of stage 2 conditional on the realized market return for years 3 and 4 following the current vintage year. At this point, we have a simulated 4-year portfolio return for an early-stage investor that vintage of startups. The average holding periods for the actual vintage returns vary substantially over time and often exceed our 4-year startup firm life-cycle. To better match the average holding period within a vintage, whenever the actual average holding period exceeds five years, we invest the proceeds at the end of stage 2 into the index for an additional two years.

Table 8 summarizes actual and model-implied VC vintage returns. The actual VC return data are based on the pool of venture capital funds active at that point in time. Our portfolio of individual startups closely resembles this object before fees are paid, but can be distorted since performance fees are determined fund-by-fund. The vintage return data include the annualized IRR net of fees, the TVPI (total value returned divided by total paid-in-capital), the holding period of the average dollar invested in the vintage. Additionally, we crudely approximate a pre-fee return by adding back a 2% fixed management fee and a 20% performance fee on returns in excess of an 8% hurdle rate. This can serve as a lower bound on the gross return as performance dispersion across funds typically results in higher actual performance fees (and, therefore, higher gross returns). To illustrate, consider a scenario with no dispersion: two funds in a vintage both achieve an 8% return. There would be no performance fee paid. Now, suppose one outperforms the hurdle rate, while the other one underperforms by the same margin. In this case, the total fee paid is increasing in performance dispersion. Importantly, dispersion matters only so far when some funds underperform while others exceed the hurdle rate. At the same, this means our approximation becomes more accurate for highly successful vintages (e.g., 1995, 1996, 1997), where the majority of funds likely achieved returns above the hurdle rate. Our imputed pre-fee IRR of 31.5% suggests an average annual fee of 6.5%.

Table 8: VC Vintages: Model vs. Data

	VC Data					Model-NDX		Model-SPX	
Vintage	Net IRR	Net TVPI	Holding period	Gross IRR	Gross TVPI	Pooled IRR	Pooled TVPI	Pooled IRR	Pooled TVPI
1985	12.90	2.68	8.12	16.13	3.37	29.59	4.74	52.66	12.66
1986	14.57	2.89	7.80	18.21	3.69	32.45	5.40	51.15	11.92
1987	18.31	2.72	5.95	22.89	3.41	30.66	4.98	41.42	8.00
1988	18.88	2.46	5.20	23.60	3.01	41.18	7.92	51.09	11.90
1989	19.16	2.59	5.43	23.95	3.21	36.87	6.57	43.95	8.90
1990	33.11	3.15	4.01	41.39	4.01	36.17	6.38	38.72	7.12
1991	27.89	3.17	4.69	34.86	4.07	38.25	3.65	34.25	3.25
1992	32.53	3.07	3.98	40.66	3.89	28.92	2.76	38.86	3.72
1993	46.71	4.13	3.70	58.39	5.48	47.31	4.71	47.50	4.73
1994	59.31	5.41	3.63	74.14	7.47	49.87	5.05	61.12	6.74
1995	88.48	6.07	2.85	110.60	8.32	89.49	12.89	87.25	12.29
1996	101.41	4.97	2.29	126.76	6.52	96.92	15.04	67.36	7.84
1997	92.59	3.12	1.74	115.74	3.80	59.51	6.47	49.16	4.95
1998	11.88	1.50	3.61	14.85	1.65	13.51	1.66	23.25	2.31
1999	-0.85	0.95	6.01	1.15	1.07	-18.07	0.45	-8.17	0.71
2000	0.03	1.00		2.03	1.13	-16.65	0.34	9.69	1.74
2001	3.61	1.25	6.29	5.61	1.41	-4.67	0.75	14.66	2.27
2002	1.07	1.07	6.36	3.07	1.21	12.46	2.02	24.34	3.70
2003	8.97	1.73	6.38	11.21	1.97	17.04	2.57	26.67	4.13
2004	6.92	1.63	7.30	8.92	1.87	15.85	2.42	22.48	3.38
2005	9.75	1.93	7.07	12.19	2.25	11.01	1.87	13.22	2.11
2006	9.06	1.95	7.70	11.33	2.28	16.51	2.50	17.17	2.59
2007	14.32	2.60	7.14	17.90	3.24	15.62	2.39	13.90	2.18
2008	13.32	2.17	6.20	16.65	2.60	18.10	2.71	15.33	2.35
2009	14.06	2.30	6.33	17.58	2.79	41.41	8.00	44.35	9.05
2010	24.13	3.99	6.40	30.16	5.41	33.15	5.57	39.70	7.43
2011	22.94	3.77	6.43	28.68	5.05	29.72	4.77	38.61	7.09
2012	20.91	3.60	6.75	26.14	4.79	36.72	6.53	44.21	8.99
2013	19.25	2.84	5.93	24.06	3.59	31.75	5.23	36.18	6.38
2014	20.57	2.97	5.82	25.71	3.79	34.03	5.80	35.21	6.11
2015	16.12	2.30	5.57	20.15	2.78	36.11	6.36	29.29	4.67
2016	17.92	2.29	5.03	22.40	2.76	40.52	7.70	46.51	9.89
Mean	24.99	2.76	5.54	31.47	3.50	30.67	4.88	35.97	5.97
Stdev	25.45	1.21	1.61	31.60	1.73	23.54	3.20	18.68	3.40

Note: This table compares the actual VC vintage returns with the model-implied VC returns. “Model-NDX” uses the Nasdaq 100 as the market factor, and “Model-SPX” uses the S&P 500 as the market factor.

There are several striking patterns in the time series of VC vintage returns. First, the average net of fee annualized IRR for VC funds is a substantial 25% with only 4 vintages reporting IRRs below

5%. The model-implied vintage IRRs average 30.77% when calibrated to the NDX and average 36% when calibrated to the SPX. The NDX-calibrated model has three vintages with negative returns, while the SPX-calibrated model has only one vintage realizing negative returns. One of the most striking features of the empirical VC returns is the remarkably high returns realized by the 1995, 1996, and 1997 vintages, with IRRs averaging 94%. For these same vintages, the NDX-calibrated model has vintage IRRs averaging 82%, while the SPX-calibrated model earns 68%.

Regressions of empirical vintage returns on model-implied vintage returns, shown in Table 9, yield slope coefficients near 1 and R^2 values of 0.55 for the SPX-calibrated model an 0.74 for the NDX-calibrated model. The high R^2 values suggest that the model captures a substantial portion of the variation in VC vintage returns despite being calibrated only to unconditional exit probabilities. The slope coefficients close to 1 indicate that the model not only captures the patterns, but also the magnitude of return variation across vintages. The NDX-calibrated model performs slightly better in matching the extreme returns of the late 1990s vintages, consistent with the findings from our IPO probability analysis.

Table 9: VC Vintages – Excess Returns

	VC Index	VC Index	VC Index	VC Index	Model (SPX)	Model (NDX)
Intercept	0.049	0.044	-0.129	-0.037	0.176	0.083
	(1.20)	(1.35)	(-2.08)	(-1.06)	(11.14)	(7.35)
SPX	2.305				2.096	
	(5.82)				(13.69)	
NDX		1.469				1.624
		(8.06)				(25.46)
Model (SPX)			1.049			
			(6.17)			
Model (NDX)				0.923		
				(9.38)		
Adj R^2	0.514	0.674	0.545	0.737	0.857	0.954

Note: This table shows regressions of excess returns, defined as the IRR minus the risk-free rate, of VC vintages on horizon-matched excess returns of equity indices and model-implied excess returns of VC vintages.

The general patterns in actual VC vintage returns are well-matched by the model, which is highly encouraging for the basic design of the structural model given it was calibrated solely to unconditional startup exit rates with no reference to startup or VC fund returns.

5 Discussion

The model's success in matching both IPO dynamics and vintage returns suggests that the compound option structure combined with market factor exposure captures key economic forces driving startup outcomes. In particular, the strong performance during the late 1990s tech boom in both validation exercises highlights the model's ability to capture how market conditions affect both exit decisions and returns. The results also suggest that a substantial portion of VC fund returns can be explained by the fundamental economics of startup investing.

A striking feature of our results is the model's ability to match time series patterns with constant parameters, calibrated to unconditional moments. Despite being calibrated only to unconditional exit probabilities, the model explains approximately 23% to 42% of the variation in quarterly IPO rates and 55% to 74% of the variation in vintage returns. This predictive success stems from the model's core economic mechanism: the compound option structure translates market return variation into state-contingent exit decisions, which in turn drive investment returns.

Our stage-specific beta estimates provide novel insights into the evolution of market risk exposure through a startup's lifecycle. As Table 10 shows, we find that market risk exposure generally declines as firms mature, with average betas of 2.4 in stage 1, 2.0 in stage 2, and 1.8 at IPO. This reflects the fact that the leverage on the underlying 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 estimates seem quantitatively more realistic. Our estimates 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 as summarized by 10. A key point to emphasize is that the systematic risk of VC investments is derived from the systematic risk of newly-listed firms, which we estimate to be 1.8. If one accepts the option-like characteristic of venture capital (independent of whether one agrees with our parameter calibration), then 1.8 serves as a lower bound for the venture capital beta, as the option-like structure inherently amplifies leverage and, consequently, beta in earlier stages. Under our calibration, the all-stage beta is 2.1 since there is considerable leverage in the earlier stages. Remarkably, our beta lines up closely when regression VC vintage returns on the horizon-matched S&p 500 returns.

Table 10: Market Betas of VC Investments

	Stage 1 (Early)	Stage 2 (Late)	Mezzanine/IPO	All-stage Beta
SPX Beta (This Paper)	2.4	2.0	1.8	2.1
NDX Beta (This Paper)	1.9	1.5	1.4	1.6
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
Cochrane (2005)				1.9
Woodward (2009)				2.2
Driessen, Lin, and Phalippou (2012)				2.7
Ewens, Jones, and Rhodes-Kropf (2013)				1.2
Brown, Ghysels, and Gredil (2023)				1.4–1.6
Ang, Chen, Goetzmann, and Phalippou (2018)				1.8
Korteweg and Nagel (2024)				2.4

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

The model does face some limitations that suggest directions for future research. First, the high IPO rates in the early sample period (pre-1997) are not well explained by the model using market returns alone. Second, the model generates somewhat less return convexity than observed during the tech boom and crash – actual VC funds earned higher returns during the boom and experienced smaller losses afterward. These limitations appear connected and suggest a potential structural change in exit behavior around the 2000 period, when the nature of successful exits may have begun shifting from IPOs toward acquisitions.

There are several directions for future research that emerge from these findings. First, the role that model properties like the exogenous timing of investment decisions and the 2-stage development structure can be explored more rigorously. Second, our treatment of acquisitions can be refined to distinguish between high-value exits and more modest outcomes. Additional deal-level datasets may allow us to test the model's predictions about liquidation rates and different types of acquisitions.

Finally, there are several additional model predictions regarding returns that we have not yet considered. Among these are the properties of startup firm conditional return distributions and VC fund return distributions beyond vintage-level means. The structural model also describes the conditional portfolio return as a function of realized market returns, which implies an index option portfolio replicating strategy that can provide a feasible VC investment benchmark that can be evaluated with historical option price data.

6 Conclusion

This paper develops and tests a structural model of startup firm valuation that integrates standard option pricing theory with systematic risk exposure through a market factor. The model's success in matching time series patterns in both IPO rates and in venture capital returns suggests that this relatively simple framework captures key economic forces driving startup outcomes. While the model faces some limitations in matching early-sample IPO rates and extreme market events, its overall predictive success suggests the fundamental approach is sound.

The results yield several insights for our understanding of entrepreneurial finance. First, the strong link between market conditions and exit decisions provides a structural link for the widely observed pattern of "hot" and "cold" IPO markets and recent multi-year market returns. Second, the model's ability to match VC fund returns suggests that a substantial portion of these returns reflects systematic risk exposure. Third, our stage-specific beta estimates provide new evidence on how systematic risk evolves through a startup's lifecycle.

Several promising directions for future research emerge from this work. The framework could be extended to incorporate additional stages, more nuanced exit decisions, or endogenous investment timing. The model's predictions about state-contingent returns could be tested against index option-based replication strategies. More broadly, the success of this relatively simple framework suggests that the structural option-based model of startup firm valuation may help explain other patterns in entrepreneurial finance and private markets like the valuation of employee stock options and the variety of investor claims that differ in their liquidation preferences.

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.

Berk, Jonathan B, Richard C Green, and Vasant Naik, 1999. Optimal investment, growth options, and security returns. The Journal of finance 54, 1553–1607.

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.

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.

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.

Childs, Paul D., and Alexander J. Triantis, 1999. Dynamic r&d investment policies. Management Science 45, 1359–1377.

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.

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, 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.

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

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

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

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.

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.

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.

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.

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

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.

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.

Woodward, Susan E, and Robert E Hall, 2004. Benchmarking the returns to venture. .