Valuation Dynamics of Private Investments*

Arthur Korteweg[a], Morten Sorensen[b]

[a]*USC Marshall School of Business, CA*
[b]*Tuck School of Business, Dartmouth College, NH*

August 2026

ABSTRACT

The financial performance of private investments has several features that are not present for publicly traded investments: Valuations are unobserved (latent) until exit. Exit time and failures are endogenous. Value can compound gradually over the holding period, like a traditional return, but unlike traditional returns it can also have a one-time jump component. Estimating our model for investments by private buyouts, real estate, and distressed funds provides new economic insights into the different dynamics across these different assets. Our estimation uses a Bayesian numerical Markov Chain Monte Carlo (MCMC) Gibbs Sampling procedure, with latent valuations sampled using Forward Filtering-Backward Sampling (FFBS).

Keywords: Valuation Process, One-Time Jump, Risk and Return of Private Investments, Dynamic Sample Selection, Competing Risks, Buyouts, Real-Estate, Distressed Debt, Bayesian Estimation, Gibbs Sampling, Markov Chain Monte Carlo (MCMC), Kalman Filtering, Forward Filtering-Backward Sampling (FFBS)

*We are grateful to MSCI, Wendy Hu, and IPC/PERC for generous access to the MSCI/Burgiss holdings data and for their support of this project. We thank seminar participants at INSEAD and HEC for helpful comments and insight.

Email addresses: korteweg@marshall.usc.edu (Arthur Korteweg), morten.sorensen@tuck.dartmouth.edu (Morten Sorensen)

Private investments are fundamentally different from traded securities in how value is created, observed, and realized. Consider a typical buyout transaction. A fund acquires a portfolio company, holds it for several years while trying to improve its operations and governance, and the fund eventually exits its investments through a sale or IPO of the portfolio company. Unlike for a publicly-traded stock, the investment's value is unobserved during the holding period, and the timing of the exit, the time when the valuation is finally observed, is endogenous to the valuation itself. Moreover, value can accrue both gradually over the holding period, like a traditional compounding return, but also as a one-time jump if, for example, the fund acquires the company at a discount or sells it at a premium. Such one-time gains or losses do not compound over time and are conceptually distinct from traditional accruing returns. Finally, datasets with private investments often have missing data. Most existing models of risk and return are designed for traded securities and assume that valuations are regularly observed and returns are compounding. For private investments, this is not true.

We develop and estimate an empirical valuation model for private investments. The model decomposes each investment's performance into a traditional compounding return and a one-time jump. One-time jumps have not been explicitly modeled before, and modeling and estimating these jumps is a main contribution of this study. In these investments, valuations are only observed at the time of initial investment and at exit, and the valuations are unobserved (latent) over the holding period. The exit probability increases when investments have higher (albeit latent) valuations, resulting in a selected sample of observed exits. To capture these features, our model combines three joint latent processes: a valuation process, a selection process for positive exits (e.g., IPOs or acquisitions), and a failure process for negative exits (e.g., going out of business). These features are common for private investments, such as private equity, distressed deals, and real-estate investments, and we estimate the model with data from these asset classes.

Estimation The estimation of our empirical model is numerically difficult because the valuations are only observed at entry and exit, and exits are endogenous. Evaluating the likelihood requires integrating over the resulting latent valuations and selection variables, rendering maximum likelihood (ML) estimation computationally infeasible. Instead, we adopt an iterative Bayesian Markov Chain Monte Carlo (MCMC) procedure using Gibbs sampling. The key ingredient is that the unobserved (latent) valuations are sampled using the Forward Filtering-Backward Sampling (FFBS)

algorithm (Carter and Kohn, 1994; Frühwirth-Schnatter, 1994). Other blocks of the procedure sample the latent selection variables, the failure variables, and the parameters in the valuation and selection processes from their respective conditional distributions. After convergence, iterating on the procedure generates the full posterior distribution of the model's parameters and variables.

Summary of empirical findings We find economically large one-time jumps for buyout and real estate investments, but not for distressed deals. For buyouts, the estimated market beta ranges from 0.75 to 1.01 across specifications, with an annualized volatility from 18.5% to 25.2% and an alpha between 15.68% and 18.76%. The arithmetic mean of the one-time jumps ranges from −14.2% to −12.5%, and this large negative jump is consistent with buyout funds paying a premium when acquiring portfolio companies. For private real-estate deals, the estimated market beta is higher, at 1.02 and 1.30, but the return has lower volatility. The alpha ranges from −1.04% to 5.24%, and the sign of the alpha flips when one-time jumps are included. The arithmetic mean of real estate's one-time jump is even more negative than for buyouts, at about −24.0% and −19.2%, consistent with the recent downturn in commercial real estate markets. Return dynamics of distressed deals are different. The estimated beta is from 0.49 to 0.72, lower than the betas for buyout and real estate deals, and the annual volatility is 4.1% to 7.0%. The alpha is negative, between −1.00% and −5.72%, and in contrast to the other two asset classes, the one-time jump is small and not statistically significant from zero, with an arithmetic mean from 3.12% to 3.28%.

Turning to the selection process, the latent-valuation coefficient is positive and significant for all three asset classes. However, the magnitude is smaller for distressed deals (0.26 to 0.33) than for buyouts (0.55 to 0.57) and real estate deals (0.74 to 0.89). This is consistent with the exit timing of distressed deals depending more on legal and restructuring milestones than on the valuation itself. The coefficient on the deal age is also consistent with these different dynamics. Buyout and real-estate exits have concave hazard rates, while distressed deals have an increasing exit hazard over time.

Literature Two existing studies model and estimate the cash flow dynamics of private investments, accounting for the endogenous selection of exits. Cochrane (2005) introduces a structural selection model for venture capital but finds that the model is numerically challenging to estimate using maximum-likelihood. Korteweg and Sorensen (2010) extend this model with a structural selection

process for the exits, which incorporates the information from periods *without* exits, and develop a Bayesian Markov Chain Monte Carlo (MCMC) estimation approach combining Gibbs sampling and Forward Filtering-Backward Sampling (FFBS). This paper extends Korteweg and Sorensen (2010) by including the one-time jump and the failure process (see also Korteweg and Sorensen, 2016).

A few previous studies have also found that private investments have valuation dynamics that are inconsistent with purely compounding returns. In the real-estate literature, Case and Shiller (1987, 1989) estimate repeat-sales house price indices using Generalized Least Squares (GLS). GLS estimation requires estimating the variance for each observation, and their predicted variance has a positive intercept in their holding-period regressions, implying that the return variance does not vanish as the investment horizon approaches zero. Recently, Sagi (2021) also finds that the mean and variance components do not scale linearly with the holding period for commercial real estate transactions and calibrates a search-and-matching model with transaction-specific shocks. Giacoletti (2021) provides similar evidence for the variance of residential transactions in California, and estimates a term-structure of idiosyncratic risk. Lovo and Spaenjers (2018) test an equilibrium model of art auctions with transaction-specific risk where idiosyncratic risk also does not converge to zero as holding periods shrink. Holding periods and realized returns are jointly determined, introducing an endogenous-timing problem similar to the endogenous exits captured by our selection process.

Outline Section 1 presents the empirical model, including the valuation, selection, and failure processes. Section 2 provides a general description of the Bayesian estimation procedure with additional details in the Appendix. Section 3 describes the MSCI/Burgiss holdings-level data. Section 4 reports empirical estimates and compares results across buyout, real-estate, and distressed deals. Section 5 concludes.

1. Econometric Model

The model describes the financial outcome of a private deal with an observed initial investment and possibly an observed exit value when it experiences a positive outcome. The deal is illiquid, and interim valuations are unobserved and modeled as a latent valuation process. Moreover, the

exit time, which determines when a final valuation is eventually observed, is endogenous, since an exit is more likely for a better-performing deal. Consequently, the absence of an exit is also informative, and this absence shifts the posterior distribution of the valuation downwards. Finally, the deal may be discontinued or written off in a negative exit when its valuation falls below a threshold, similar to structural models of the firm (e.g., Merton, 1974; Leland, 1994).

We model these features as follows. Initially, an observed investment is made. The initial (unobserved) valuation may be different from the invested amount, with the difference being the one-time jump (described below). From there, the unobserved (log) valuation evolves stochastically over time, driven by factor exposures and idiosyncratic shocks (for example, according to a log-CAPM process). Eventually, the investment ends in one of several ways: a positive exit with an observed exit value, a failure with a near-zero value, an exit date without a reported amount due to missing data, or no recorded exit within the sample period. The exit type and timing are endogenous to the valuations, with higher valuations being more likely to lead to a positive exit, while low valuations increase the risk of a failure. The model captures these dynamics with three latent processes: (i) a latent log-valuation process (diffusion process), (ii) a selection process for the arrival of positive exits with observed valuations, and (iii) a failure arrival process for negative outcomes, with low but unobserved valuations.

Each investment is classified according to whether it has a positive or a negative exit. A deal with a *positive exit* is an investment that is fully realized with both a recorded exit date and a reported outcome (TVPI) that exceeds 0.15. These outcomes are captured by the selection process. We denote them "positive outcomes," even though the realized TVPI may be below 1, so investors do not necessarily earn a profit on these investments. Instead the "positive" refers to the presence of a counter-party actively acquiring the investment in a transaction. Formally, the identifying assumption is that positive exits are exits where the probability increases with the latent valuation.

In contrast, a *negative exit* is an investment with a recorded exit date but a TVPI below 0.15. The timing of these exits is governed by the failure process. Intuitively, these are failed deals that end up, for example, in bankruptcy or liquidation. Formally, the identifying assumption is that the probability of a negative exit increases for deals with a *lower* valuation, and this assumption is imposed directly in the equation for the failure process.

Deals with a reported exit date but where the outcome is missing are excluded from the analysis,

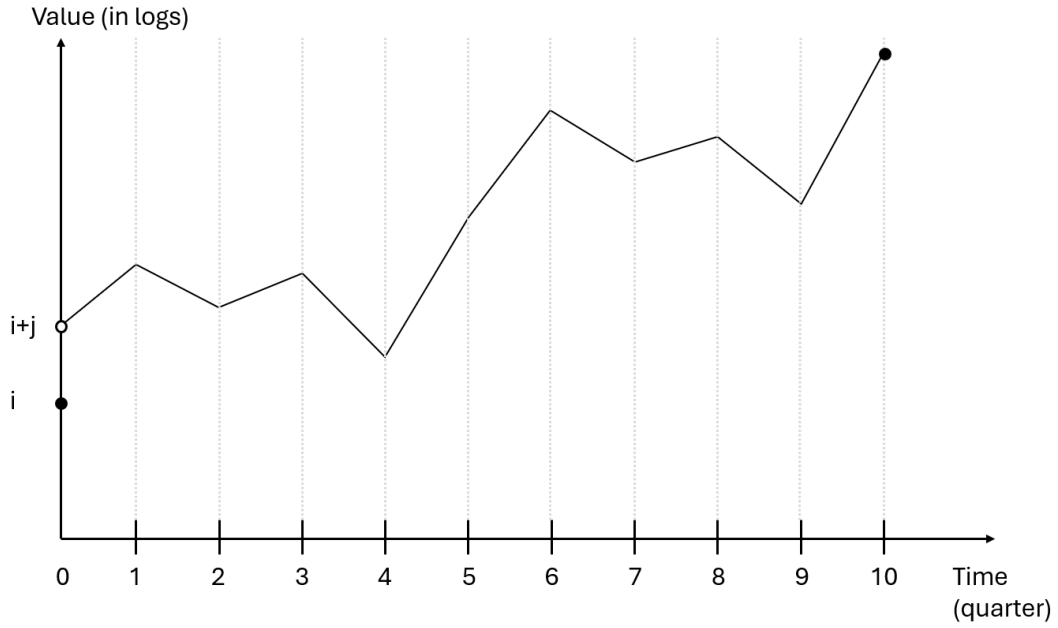
since it is not possible to determine whether they should be considered a positive or negative exit. Finally, deals without a reported exit date are also excluded since they are largely uninformative about the valuation process.

1.1 Valuation process

At time $t = 0$ an investment of the observed (log) amount i acquires an asset with latent log-value $v(0)$. Note that log values are denoted by lowercase letters, and to simplify notation we sometimes suppress index and time subscripts. The latent value, $v(0)$, may differ from the invested amount due to unobserved, deal-specific shocks:

$$v(0) \;=\; i + j, \qquad j \sim \mathcal{N}(\mu_j, \sigma_j^2). \tag{1}$$

For example, the investor may be able to acquire the asset at a discount relative to its true economic value. Imagine an investor that acquires an asset with a value of $V = \$200$ at a 10% discount, so $I = \$180$. In log terms, $v = i + j$ implies $\ln(\$200) = \ln(\$180) + \ln(\$200/\$180)$, so the one-time jump is $j = \ln(\$200/\$180) = -\ln(90\%) = 0.105$. We discuss the interpretation of the one-time jump in more detail below. To foreshadow, the estimate of the one-time jump estimated for buyout transactions is $\mu_j = 0.1215$, consistent with an average discount for a buyout transaction of $10.8\% = 1 - \exp(-\mu_j + \sigma_j^2/2) = 1 - \exp(-0.1215 + 0.0147/2)$.



Example of a value process. This graph shows an example of the time series of value for a single hypothetical asset. The investor invests an amount I at time zero. In logs this is denoted by $i \equiv \log(I)$. The investment's latent log-valuation is $v(0)$, which differs from i by an amount j. After $T = 10$ quarters, the investment is exited at a log-value $v(T)$. Filled-in circles are valuations that are observed to the econometrician.

Continuous-time diffusion After the initial investment, the (latent) valuation evolves according to a continuous-time diffusion process with systematic and idiosyncratic components:

$$dv(t) - r\,dt \;=\; \delta\,dt \;+\; \beta' df(t) \;+\; \sigma_v\,dW(t), \tag{2}$$

where r is the risk-free rate, $df(t)$ is a vector of factor returns (e.g., excess market log returns), β is the vector of factor loadings, and $W(t)$ is a standard Brownian motion. The drift δ is the expected excess log growth rate after adjusting for r and factor exposures. In our specification, there is one factor, the log return on the public stock market in excess of the log risk-free rate, i.e., $df(t) = d\log[M(t)] - r dt$, with market index $M(t)$ and risk-free rate r.

One-period discretization The one-period discretization of the continuous-time process is (in the estimation, a period is three months, i.e., one quarter):

$$v(t) \;=\; v(t-1) \;+\; r \;+\; \delta \;+\; \beta' f(t) \;+\; \varepsilon_v(t), \tag{3}$$

where $f(t) = \int_{s=t-1}^{t} df(s)$ is the realized vector of factor returns, and $\varepsilon_v(t) \sim \mathcal{N}(0, \sigma_v^2)$ is an idiosyncratic innovation that is independent of the factor returns.

Realized Return (TVPI) Distribution A deal with a positive exit is sold at time T for the observed amount, $V(T) = \exp(v(T))$. The data contain the investment's multiple, called TVPI (Total Value to Paid-In capital), given as the paid-out amount divided by the invested amount, TVPI $= V(T)/I$.

Combining with the discretized process, the log multiple ($\ln(\text{TVPI}) = v(T) - i$) is then:

$$\ln(\text{TVPI}) \;=\; v(T) - i \;=\; j \;+\; T\,(r+\delta) \;+\; \beta' \sum_{t=1}^{T} f(t) \;+\; \sum_{t=1}^{T} \varepsilon_v(t), \tag{4}$$

with variance

$$\text{Var}[\ln(\text{TVPI})] \;=\; \sigma_j^2 \;+\; T\left(\beta' \Sigma_f \beta + \sigma_v^2\right), \tag{5}$$

where Σ_f is the covariance matrix of the factor vector $f(t)$. The variance expression shows that the dispersion of outcomes scales with the duration of the investment, through accumulated systematic and idiosyncratic shocks. However, due to the one-time jump variance, the dispersion does not vanish as the investment's duration approaches zero, because even very short investments still experience one jump.

Timing of the one-time jump In the model, j is the one-time jump, and the identifying assumption is that every investment experiences exactly one such jump, regardless of its duration. This implies that the mean and variance of a deal's performance do not vanish as the deal's duration shrinks to zero, unlike for a traditional compounding return process.

In the model, we specify the jump as occurring at the time of the initial investment, with a typical example being an investor who acquires an asset at a discount relative to its economic value. More generally, however, a deal may experience several distinct one-time effects. A buyout fund may replace the portfolio company's management after acquiring it, and management is replaced only once, regardless of how long the investment is held. At exit, the fund may sell the asset at a premium or a discount relative to its economic value. These are economic effects that arise once, and only once, for each investment, irrespective of the holding period. While such effects may occur at the beginning, the middle, or the end of the investment, we model them as a single initial jump that captures their combined effect.

The individual jumps are not separately identified in the valuation process. To see this, Suppose there are two jumps, with an initial jump of j_0, as before, and a second jump j_T at exit, so that the exit log-value is $v(T) + j_T$. Iterating the transition equation from $t = 1$ to T gives

$$\ln(\text{TVPI}) \;=\; \left(j_0 + j_T\right) \;+\; T\,(r+\delta) \;+\; \beta' \sum_{t=1}^{T} f(t) \;+\; \sum_{t=1}^{T} \varepsilon_v(t). \tag{6}$$

Comparing this with the single-jump expression above, the two jumps enter only through their sum. With j_0 and j_T independent and normally distributed, the realized multiple identifies $\mu_{j_0} + \mu_{j_T}$ and $\sigma^2_{j_0} + \sigma^2_{j_T}$, but neither component separately in the valuation process. Accordingly, μ_j and σ^2_j should be read as the mean and variance of the *total* one-time component of a deal's performance, rather than of an entry effect alone. The selection process is the only feature of the model that could distinguish them, since j_0 shifts the latent valuation, and hence the exit hazard, throughout

the deal's life, whereas j_T is realized only at exit, but we do not pursue that separation here.

Outcomes of negative exits Negative exits have near zero outcomes, which are problematic for the valuation process specified in logs, because small errors in near-zero outcomes, while economically irrelevant, are important for the log return. Geometric returns are sensitive to differences in recovery amounts that are economically negligible. To illustrate, consider an investment of \$100 that returns \$400 with probability 80% and otherwise fails, returning only a nominal amount x. The arithmetic return is nearly unaffected by x, while the log return is not. The below table shows the expected multiple (TVPI) and the expected log multiple depending on the amount, x, recovered from a failed deal:

Recovery x	E[TVPI]	E[ln TVPI]
\$10.00	3.22	0.65
\$1.00	3.20	0.19
\$0.10	3.20	−0.27
\$0	3.20	$-\infty$

A reduction in the recovery from \$10 to \$0.10 reduces the expected multiple from 3.22 to 3.20, which is economically irrelevant, and a further reduction to \$0 leaves the expected multiple effectively unchanged at 3.20. However, in log returns, the reductions in the ln multiple (i.e., log return) from 0.65 to −0.27 and then to negative infinity, are significant and turn an apparently profitable investment into a loss-making one.

This insight motivates two features of our specification. First, we do not treat the reported recovery of a failed deal as an accurate observation of its liquidation value due to the large effect of small observation errors, as observed above. Instead, we treat a negative exit as a censored observation, recording that the valuation fell below the failure boundary without using the recovery amount itself. Second, we report and focus on arithmetic performance, which is the economically relevant concept, and where the precise recovery in failures is immaterial, as observed above.

Contrast with traditional jump literature Our one-time jump differs from the jumps typically studied in the asset pricing literature. Merton (1976), among others, has modeled discontinuities in asset prices as random jumps arriving according to a Poisson arrival process with intensity λ. In this setting, the expected number of jumps over a holding period of length T is λT, so

the jump contributes $\lambda T(\mu_J^2 + \sigma_J^2)$ to the variance of the log return over that period, which scales with duration exactly as the diffusion does (see also, e.g., Eraker et al., 2003). While the arrival process is not always Poisson (e.g., Duffie et al., 2000), the common property in this literature is that the jump's contribution to the variance of a holding-period return grows with the length of the period and vanishes as the period shrinks to zero. This differs from our one-time jump in two respects. Economically, we model a single repricing that every investment experiences exactly once, regardless of the length of the holding period. Econometrically, we observe only the total change in value over $[0, T]$. A jump component whose variance is proportional to T enters the variance of ln(TVPI) additively with σ_v^2 and cannot be separated from it by the duration of the deal. In contrast, the one-time jump is identified because it does *not* scale with duration, and it enters as the constant σ_j^2. Put differently, we identify our jump from the fact that the dispersion of deal outcomes does not vanish for short deals.

1.2 Selection Process for Positive Exits

Exit decisions typically involve selling a portfolio company or a property to a willing buyer, and such sales are endogenous to the investment's valuation, resulting in a selected sample of observed valuation outcomes (Korteweg and Sorensen, 2010). To address this selection issue, we follow Korteweg and Sorensen (2010) and model an exit process for the exit of an investment and the resulting observation of its valuation depending on the valuation itself. An investment has a positive exit when:

$$w(t) > 0, \tag{7}$$

where $w(t)$ is the latent selection variable,

$$w(t) = \gamma_v \, v(t) + x'(t)\gamma + \eta(t). \tag{8}$$

Here $x'(t)$ is a vector of observed covariates for the asset at time t, such as a constant term, time fixed effects, and the investment's age, where *age* refers to the time since the initial investment.

The coefficients to be estimated are (γ_v, γ), and a positive γ_v implies that investments with higher valuations are more likely to experience a positive exit. Formally, we define positive exits as exits that are more likely for higher valuations, i.e., where $\gamma_v > 0$. Conversely, negative exits are

increasingly likely for lower valuations. More generally, the selection issue arises whenever $\gamma_v \neq 0$, and the estimated coefficients in the valuation equation will generally be biased, as discussed in Cochrane (2005) and Korteweg and Sorensen (2010). The error term $\eta(t) \sim \mathcal{N}(0,1)$ is assumed i.i.d. and with the variance normalized to one, without loss of generality following standard arguments for the probit model.

1.3 Failure Process for Negative Exits

Failures (negative exits) are modeled analogously but for the lower tail of the valuation process. An investment fails when its latent value falls below a boundary that triggers a write-off or liquidation of the investment. This failure process is governed by the latent variable $z(t)$, and a failure occurs when:

$$z(t) < 0, \tag{9}$$

where $z(t)$ follows the latent failure equation

$$z(t) = v(t) + \phi + \xi(t). \tag{10}$$

Here ϕ is a fixed intercept that sets the failure threshold, and $\xi(t)$ is an idiosyncratic error term capturing unobserved factors influencing the failure decision, assumed i.i.d. $\xi(t) \sim \mathcal{N}(0, \sigma_\xi^2)$. In general, this specification is only identified up to scale (like a probit model), and we normalize the coefficient on $v(t)$ to one without loss of generality. Moreover, due to the unobserved outcomes of failures, the remaining parameters in the failure process are only weakly identified, and we fix, rather than estimate, these parameters. Specifically, the main specification sets $\phi = -\ln(0.15)$ and $\sigma_\xi = 0.2$.

The specification of the failure process is motivated by structural models of firm value and credit risk. For example, in the Merton (1974) and Leland (1994) models a firm defaults when its value falls below the face value of its debt. Subsequent extensions of these models have derived similar failure thresholds, and our specification captures this economic logic. Specifically, the failure boundary is represented by the constant threshold $-\phi$, and failure occurs with increasing probability as the latent log valuation $v(t)$ declines relative to this boundary, that is, when $v(t) < -\phi - \xi(t)$. Thus, investments with higher valuations are less likely to fail, and a lower value of ϕ corresponds to a

higher threshold and more defaults.

In our empirical model, the failure process functions primarily as a lower guard rail on the valuation process. It prevents the latent valuation paths with near-zero outcomes from descending to arbitrarily low values, and it imposes the economic condition that failures are approached from above, as the result of deteriorating performance.

2. Estimation

Evaluating the likelihood is infeasible because it requires high-dimensional integration over the latent valuation paths and latent selection/failure variables across all observations in the dataset. Instead, we adopt a Bayesian estimation procedure using Gibbs sampling, which is a data augmentation method for sampling from the augmented joint posterior distribution (see Geman and Geman, 1984; Tanner and Wong, 1987; Gelfand and Smith, 1990; Johannes and Polson, 2010; Korteweg and Sorensen, 2010, 2016). This iterative conditional sampling method turns the intractable integration problem into a sequence of tractable draws from the conditional distributions of stochastic variables, and the model is designed to exploit the Gaussian-linear structure of the valuation process and the probit structure of the selection and failure processes to produce tractable conditional distributions for these draws. The econometric model can be divided into blocks, with different variables and parameters, that are sampled conditional on the previous draws of the variables and parameters in the other blocks.

This block-wise structure is a key advantage of Gibbs sampling. Each conditional draw is analytically tractable (Gaussian or truncated Gaussian for latent variables; conjugate or conditionally Gaussian-Inverse Gamma for parameters), yet the overall algorithm produces the full joint posterior. More details, including the FFBS extension for the initial jump, are in Appendix A. We describe the intuition behind each block below.

1. *Latent valuation path* $\{v(t)\}$ *(state block).* Except for investments and positive exits, all interim valuations are unobserved and treated as latent states. Conditional on the current draws of the parameters and selection/failure variables, the valuation process reduces to a linear Gaussian state-space model. We sample the entire valuation trajectory $\{v(t)\}$ from the implied conditional posterior distributions using Forward Filtering-Backward Sampling (FFBS)

(Carter and Kohn, 1994; Frühwirth-Schnatter, 1994). Our FFBS implementation accommodates the one-time jump at inception and incorporates the observed valuations of positive exits.

2. *Jump parameters* (μ_j, σ_j^2). Given the sampled initial latent valuations across deals, the jump mean and variance are drawn using standard Gaussian-Inverse Gamma conjugacy, and our implementation follows Gelman et al. (2013, Sec. 3.3). Note that this block is skipped in the "no jump" specifications where the jump is turned off.

3. *Valuation equation parameters* $(\delta, \beta, \sigma_v^2)$. Conditional on the latent valuations, the valuation parameters are sampled using Bayesian linear regression with conjugate priors.

4. *Latent selection/failure variables* $\{w(t), z(t)\}$. Conditional on $\{v(t)\}$ and parameters, the latent selection and failure variables are sampled from truncated Gaussian distributions. In periods with observed exits (or observed failures), the latent variables are drawn from the appropriate truncated distribution. In periods with non-exits (and non-failures), they are drawn with the opposite truncation. This is the standard data augmentation model for Bayesian estimation of the probit model (Albert and Chib, 1993).

5. *Selection coefficients* (γ_v, γ). Conditional on the drawn selection variables, the selection parameters (γ_v, γ) are then sampled using a Bayesian linear regression with conjugate priors. The gamma parameters include the loading on the valuation, γ_v, as well as time fixed effects to capture time variation in exit markets.

2.1 Implementation

The priors are diffuse. In the valuation equation, (δ, β) has a normal prior with mean $(0, 1)$, corresponding to zero drift and a market-neutral beta of one, and a prior variance of 25. The variance of the error term in the valuation process has the standard flat, non-informative reference prior. Coefficients in the selection process, including the time fixed effects, have independent normal priors with mean zero and variance 25. The mean and variance of the one-time jump have a diffuse conjugate Gaussian-Inverse Gamma prior, centered at $\mu_0 = 0$ and $s_0^2 = 0.03$, and is specified with two pseudo-observations for the jump mean and variance.

The main specifications are run for 350,000 Gibbs iterations, of which the first third is discarded as burn-in, leaving roughly 233,000 draws for inference. The high number of iterations is necessary because the sampler mixes slowly, especially for the one-time jump. A single run takes approximately 8.5 hours for the buyout sample (N = 7,243), 4.5 hours for real estate (N = 4,447), and 1.3 hours for distressed debt (N = 859).

Convergence is assessed by inspecting the sample path and calculating the implied effective sample size. The effective sample size uses the initial positive sequence estimator of Geyer (1992). Figure 5 plots the draws and the corresponding posterior histograms for the buyout specification. The three main specifications behave similarly. Median effective sample sizes are roughly 2,200 for buyouts, 1,900 for real estate, and 860 for distressed deals. The selection coefficients are estimated most precisely, and the one-time jump parameters are estimated less precisely, with one-lag autocorrelations between 0.99 and 1.00 and effective sample sizes of 115, 114, and 127. This is due to the fundamental identification problem, discussed in Section 1, where the data contain the total entry-to-exit change in the value (TVPI), but not how it is divided between the one-time jump and the accumulated drift, and the sampler traverses that ridge slowly. However, despite the identification problems for the jump parameters, the total arithmetic return is better identified, because it is relatively constant across the ridge. This is seen in the simulation evidence below, where the annualized arithmetic return α is recovered essentially exactly, and the arithmetic entry jump is recovered within half a standard error, even though the log parameters from which they are constructed are individually less well estimated.

2.2 Simulation Evidence

We assess the properties of the estimation procedure by estimating it using simulated data with known parameters. Specifically, we generate twenty-four independent datasets, each with N = 7,200 observations, corresponding to the size of the buyout sample's 7,243 observations. The parameter values used to generate the data are the estimates from the buyout sample, as reported in column (1) of Table 5. For each deal in the simulated data, we estimate the one-time log jump $j \sim \mathcal{N}(\mu_j, \sigma_j^2)$, and its valuation then evolves as specified by the valuation process. The deal fails when the latent index $z(t) = v(t) + \phi + \xi(t)$ crosses zero, and it has a positive exit governed by the selection process, so holding periods are endogenous. As in the empirical specifications, the failure

parameters are fixed rather than estimated, and they are set to $\phi = -\ln(0.15)$ and $\sigma_\xi = 0.2$, leading to a failure rate of 11.4% across the simulated datasets. Moreover, to simplify the construction of the simulated data, they are simulated without time fixed effects.

For each dataset, we estimate the model in the same way as for the actual data, with 350,000 iterations and discarding the first third for burn-in (note that the estimates with simulated data are run locally, on substantially faster computers, allowing us to run twenty-four parallel estimations). Table 4 reports the true value of each parameter, the average of the estimates across the twenty-four datasets, the standard deviation of the estimates across the twenty-four datasets in parentheses, and the average within-dataset posterior standard deviation. Figure 1 plots the draws for a representative dataset. Table 4 also reports the ratio of the standard deviation of the estimates across datasets to the average posterior standard deviation within each dataset. According to the Bernstein–von Mises theorem, the posterior distribution should have the same variance as the sampling distribution of the estimator itself, asymptotically (Gelman et al., 2013, Ch. 4), so their ratio should converge to one. The between-dataset standard deviation is itself estimated from twenty-four datasets, with a relative precision of about 15%, so ratios between roughly 0.7 and 1.3 are within two standard errors of one.

The estimated parameters are close to the true ones across both the valuation and selection processes. The annualized arithmetic abnormal return α matches its true value to four decimals (0.1290 against 0.1290), β and δ are within 0.7 standard errors, σ_v^2 is half a percent low, and all four selection coefficients are recovered to within 1.5 standard errors and 0.8% in level terms. The one exception is the jump variance σ_j^2, which comes out above its true value, while the jump mean μ_j and the arithmetic jump $\exp(\mu_j + \sigma_j^2/2) - 1$ are both recovered well. The likely mechanism is that the jump variance is estimated from jumps that are themselves inferred rather than observed, and therefore absorbs their posterior uncertainty. An implication is that while the estimates recover the typical jump in arithmetic terms, they overstate the variation of the log jump across deals. This is the same arithmetic-versus-log asymmetry discussed in Section 1. The data identify the combination $\mu_j + \sigma_j^2/2$ rather than the log location and scale separately, which is why the arithmetic quantities are the best-recovered estimates in the table. The final row makes this point in economic terms. Simulating the estimated model forward reproduces the average TVPI observed in the simulated samples to within a tenth of a percent (3.451 against 3.455).

3. MSCI/Burgiss Holdings Data

We use data from MSCI (formerly Burgiss), a leading source of cash-flow and valuation data for private markets. Most of the data are collected from reports from general partners (GPs) to their limited partners (LPs), who are often MSCI/Burgiss clients, and the data have been extensively validated and cross-checked to ensure consistency and accuracy (Brown et al., 2020).

The data contain an extensive sample of (anonymized) individual private investments (holdings in the MSCI/Burgiss terminology) by thousands of funds across multiple asset classes. We use three samples in our analysis, containing the holdings of funds that MSCI/Burgiss classifies as buyout, real estate, or distressed debt funds. We restrict the sample to investments in North America, with holding status *exited*, non-missing investment dates, a reported TVPI, and a holding period between five and sixty quarters. We exclude investments made after the end of 2020 to allow sufficient time for exits (exit data are current through the second quarter of 2025). In terms of sample size, the entire samples of North American holdings of buyout, real estate, and distressed debt funds first invested through the end of 2020 number 16,004, 19,824, and 7,395, respectively. Of these, 4,174, 9,648, and 1,469 cannot be classified as either a positive or a negative exit, and a further 4,297, 5,544, and 4,801 have no recorded exit date at all, leaving 7,533, 4,632, and 1,125 holdings with an observed outcome. The requirement of a five to sixty quarter holding period leaves final samples of 7,243 buyout, 4,447 real estate, and 859 distressed debt holdings. Negative exits account for 15.6%, 6.7%, and 3.6% of these deals.

At the fund level, *Vintage* is the year of the fund's inception, and *Size (USDmm)* is the fund's committed capital in millions of US dollars. At the deal level, *Total Inv (USDmm)* is the total amount invested into the deal by the fund. Note that this variable has many missing values, and it is included in the descriptive statistics but not used in the actual estimation, where we instead use the TVPI as our performance measure. The variable *Start (qtr)* is the quarter of the investment.[1] For investments with a recorded exit, either positive or negative, *End (qtr)* is the exit quarter, and *Holding Period (qtr)* is the holding period in quarters, calculated as the difference between the starting and ending quarters. Investment performance is given as the reported gross value-to-paid-in multiple (*Gross TVPI*), defined as the total amount generated by the investment, before subtracting fees such

[1] Quarters are coded using Stata's standard for quarter numbering, where quarter zero is Q1 of 1960.

as management fees and carried interest, divided by the total amount invested. Additionally, the data report gross internal rate of return (*Gross IRR*), and we report descriptive statistics for this variable, but it is also not used in the actual estimation. Note that we do not observe individual cash flows, and the TVPI and IRR measures are calculated by MSCI/Burgiss. Brown et al. (2020) discuss the data collection and coverage of the MSCI/Burgiss data in detail.

To classify exits, we treat an investment with a TVPI over 15% as a *positive exit*, and we use ln(TVPI) as its outcome for the estimation. Note that a positive exit does not imply that the investment is profitable for the investor (that requires a TVPI above one). Instead, a positive exit means that the investment is sold, typically in an acquisition or initial public offering (IPO), for a substantial amount in a transaction with an active counter-party. The defining property is that the probability of a positive exit is increasing in the valuation. In contrast, an investment with a TVPI below 15% is coded as a *negative exit*, or failure. Such an investment is typically liquidated for a nominal amount, and the defining property of a failure is that the probability is increasing for *lower* valuations. Due to uncertainty about the timing and recovery for negative exits, we treat these outcomes as censored at the barrier of 15%. This also resolves the zero-TVPI problem inherent in the log specification, where the log of zero is negative infinity (see, e.g., Cochrane, 2005).

The MSCI/Burgiss holdings data cover a broad range of private investments. In the MSCI taxonomy, the Tier-1 categories are Private Equity, Private Credit, and Private Real Assets. To get a broad representation of private investments, we focus on the Tier-2 subcategories with more observations within each of the Tier-1 categories. For Private Real Assets and Private Credit, these categories are Real Estate and Distressed Debt. For Private Equity, the most common Tier-2 category is Venture Capital, followed by Buyout. However, in this Tier-1 category, we use buyout deals because our econometric model is not suited to venture capital investments, due to the frequent syndication among investors and staged investments across multiple rounds (see Korteweg and Sorensen, 2010; Cochrane, 2005). Our dataset therefore contains holdings-level observations for buyout, real estate, and distressed debt funds. Tables 1 through 3 report summary statistics, and Figures 2 through 4 show the corresponding distributions by asset class.

4. Results

4.1 Estimates for Buyout Funds

Buyout Investments As defined by MSCI/Burgiss, *Buyout* investments are a subset of the broader Private Equity category, which includes company-focused "[e]quity investments in privately held companies or companies that are expected to become private as a result of the transaction." Specifically, buyout investments are "[e]quity investments in mature companies that result in a change of control. These are typically larger transactions that use leverage."

Buyout funds constitute the largest and most mature segment of our dataset, with an average vintage year of 2008.0 and a mean fund size of $2.25 billion (Table 1). Individual buyout deals are also substantially larger than those in the other two categories, with an average total invested capital of $113 million against $31 million for real estate and $48 million for distressed debt. Panel A of Table 2 shows that buyout deals span a broad cross-section of sectors, with Industrials, Consumer Discretionary, and Information Technology together accounting for just over half of all transactions.

Our buyout sample starts in 1984, although coverage is sparse until the late 1990s (Panel B of Table 2). We restrict the sample to investments made through the end of 2020, to allow time for later deals to realize an outcome. The distribution of holding periods is shown in Panel B of Figure 2. A buyout deal is typically held for around four years, but the distribution is right-skewed, consistent with average and median periods of 5.9 and 5.3 years, or 23.6 and 21 quarters (Tables 1 and 3).

Buyout investments perform well, with a mean gross TVPI of 2.93, a median of 2.33, and a standard deviation of 4.06 (Panel A of Table 3). These are gross TVPIs, as reported by MSCI/Burgiss, and they do not subtract the funds' fees, such as management fees and carried interest. The histogram in Panel A of Figure 2 shows the dispersion in buyout outcomes, with 15.6% of buyout deals having negative exits with a gross TVPI at or below 15%, and a substantial share of these are complete write-offs with a reported TVPI of zero.

Panels C and D of Figure 2 show how these outcomes vary with the holding period. Mean gross TVPI declines steadily as the holding period increases. However, the variance of the gross TVPI is close to flat across holding periods, which is inconsistent with a pure diffusion that would imply that dispersion widens linearly as the holding period grows.

Buyout Estimates Table 5 reports parameter estimates for the 7,243 buyout deals under three nested specifications. Column (1) is the full model, containing both the one-time jump at entry and the failure process. Column (2) removes the one-time jump, and column (3) removes the failure process as well, leaving a specification with neither. Comparing the columns therefore isolates the contribution of each component in turn. All parameters are well estimated, and the finite-sample behavior of the estimator at this sample size is documented in Table 4 and Figure 1.

In the valuation process, the estimated market beta is 0.6974 in the full specification and essentially unchanged at 0.6744 once the one-time jump is removed, but it rises to 1.4235 when the failure process is removed as well. The first two figures sit towards the lower end of the range reported in the literature. Korteweg (2023) surveys existing estimates, ranging from below one to substantially above one.[2] The estimated quarterly variance σ^2 ranges from 0.0733 to 0.0934, implying an annualized idiosyncratic volatility between 54% and 61%.[3] This dispersion matches the cross-sectional spread of TVPI outcomes seen in Panel A of Figure 2 and the relatively flat variance-by-duration profile in Panel D, which is difficult to reconcile with a geometric Brownian motion but arises naturally in a jump–diffusion setting.

The Bayesian estimation approach, with simulated posterior distributions, makes it trivial to estimate non-linear transformations of the parameters, such as the annual alpha and arithmetic jump mean. For each iteration, we compute these two statistics from the parameter draws, and Table 5 reports the mean and standard deviation of their posteriors, and the full posterior distributions are shown in Figure 5.[4] The annual alpha is economically large, at 12.92% in the main specification, although this figure is not directly comparable to the net-of-fee, fund-level alphas

[2] See Korteweg (2023), Figure 1, for a compilation of CAPM beta estimates across studies.

[3] The estimated σ^2 is the variance of the quarterly innovation ε in the valuation equation. Innovations are independent by assumption, so variance scales linearly in time. Annualized variance is $4\sigma^2$ and the corresponding volatility is $\sqrt{4\sigma^2}$.

[4] At each retained draw we compute the quarterly alpha $\alpha = \delta + \frac{1}{2}\sigma_v^2 - \frac{1}{2}\beta(1-\beta)\sigma_M^2$ and the arithmetic jump $\exp(\mu_j + \frac{1}{2}\sigma_j^2) - 1$, and summarize the resulting posterior distributions. Because α is a per-period quantity, the table reports four times its posterior mean and four times its posterior standard deviation, whereas the jump is a one-time shift in the level of the valuation at entry and is reported as drawn. In the expression for α, the term $\frac{1}{2}\sigma_v^2$ converts the log drift δ into an arithmetic expected return, and the final term applies the corresponding convexity adjustment to the CAPM benchmark, since the deal's market exposure contributes $\frac{1}{2}\beta^2\sigma_M^2$ while the benchmark charges β times the market's own $\frac{1}{2}\sigma_M^2$. We set $\sigma_M^2 = (0.15)^2/4$, the quarterly variance implied by an annual market volatility of 15%. The reported alpha is insensitive to this choice: the market variance enters only through the term $\frac{1}{2}\beta(1-\beta)\sigma_M^2$, and in our main specification in column (1) varying the annual market volatility between 10% and 20% changes the annualized alpha by about 0.3 percentage points.

surveyed in Korteweg (2023), because our alpha is gross of fees and does not account for the one-time jump, which we treat separately.

Table 5 also reports estimates of the one-time jump, in column (1). The jump mean is large and significantly positive, $\mu_j = 0.1215$ with a posterior standard deviation of 0.0242, and an estimated jump variance of $\sigma_j^2 = 0.0147$. Applying the Jensen correction (see footnote 4), the implied arithmetic mean jump is 13.79%. Economically, a positive jump could mean, for example, that buyout funds tend to acquire companies at a discount relative to their true economic value, and the estimate would then correspond to an average acquisition discount of 10.8%, as discussed in Section 1.

In column (1)'s selection equation, the coefficient on the valuation is positive, and equals 0.5585 with a posterior standard deviation of 0.0121, confirming that better-performing deals are more likely to have a positive exit. This is unsurprising, but the endogenous timing creates the dynamic sample-selection problem that can potentially bias risk and return estimates, as discussed by Korteweg and Sorensen (2010).

The linear age term is positive and significant, and the indicator for deals aged eight quarters or less is negative and significant, implying that buyout deals have a low probability of exit during the first two years, and that this probability rises as the deal matures. The pattern is consistent with Panel B of Figure 2 and our median holding period of 5.3 years. This holding period is close to the median of five years reported by Brown et al. (2020), who also find that both multiples and public-market equivalents decline for holding periods beyond roughly three years, a pattern they interpret as buyout funds targeting a holding period of five years or less. Ivashina (2023) likewise reports average holding periods around five years, but somewhat shorter when valuations are high. The declining performance at longer horizons is also visible in our data (Panel C of Figure 2), although we attribute this pattern to the faster exits of better-performing deals and the resulting attrition in the pool of longer-term deals.

4.2 Estimates for Real Estate Funds

Real Estate Investments In the MSCI/Burgiss taxonomy, *Private real estate* investments are a subset of the broader Private Real Asset category, which includes asset-focused "[i]nvestments that provide ownership interests in underlying tangible assets, including land, properties, and

commodities, as well as long-life assets/structures that provide essential products or services." Specifically, Private Real Estate is "[i]nvestments in properties and/or land that will be zoned for properties."

The real-estate funds in our data are younger and smaller than buyout funds, with an average vintage year of 2010.6 and a mean fund size of $1.41 billion (Table 1). Individual real-estate deals are also smaller, averaging $31.4 million of invested capital against $113 million for buyout deals (Panel A of Table 2).

Real-estate deals have shorter holding periods (mean of 19.9 and median of 17 quarters) than buyout deals (mean of 23.6 and median of 21 quarters), as seen in Tables 1 and 3. Outcomes are also less dispersed, with a mean gross TVPI of 1.76 and a standard deviation of 1.19 (Panel B of Table 3), against 2.93 and 4.06 for buyouts. Negative exits are correspondingly rarer, at 6.7% of real-estate deals against 15.6% for buyout. The narrower distribution is visible in Panel A of Figure 3, and reflects the more stable cash flows from real-estate assets and the lower operational risk compared to buyout transactions.

Real-Estate Estimates Table 6 reports parameter estimates for the 4,447 real-estate deals. Column (1) is the full model, with both the one-time jump at entry and the failure process, column (2) removes the jump, and column (3) removes the failure process as well. The estimates are stable and reasonably precise across the first two columns, while column (3) departs from both, as it does for buyout deals.

In the valuation process, the estimated market beta is 0.9308 in the full specification and essentially unchanged at 0.9369 once the one-time jump is removed, rising to 1.4181 when the failure process is removed as well. Real-estate betas are close to one and higher than our buyout estimates of around 0.70. As noted by Korteweg (2023), betas for real-estate funds depend on whether they are estimated from NAV-based returns, cash flows, or secondary prices, and values close to one are plausible after addressing smoothing and stale pricing. The estimated quarterly variance σ^2 ranges from 0.0331 to 0.0356, implying an annualized idiosyncratic volatility between 36% and 38%, well below the 54% to 61% estimated for buyout deals. The lower volatility is consistent with the narrower dispersion of TVPI outcomes in Panel A of Figure 3.

The annual alpha, computed from the posterior draws as described in footnote 4, is 4.24% in the main specification in column (1) and 3.72% once the one-time jump is removed. In specification (3),

it falls to 0.44%, and is no longer distinguishable from zero, when the failure process is removed as well. Real-estate alphas are therefore positive and economically meaningful, but substantially smaller than the 12.92% we estimate for buyout deals. Table 6 also reports the one-time jump, in column (1). The jump is close to zero, with posterior mean of $\mu_j = -0.0242$ and standard deviation of 0.0267, implying an arithmetic mean jump of −2.07%.

Turning to the selection equation, the coefficient on the valuation is positive and significant, at 0.7517 in the full specification and 0.7399 without the jump. A positive coefficient implies that higher-valued properties are more likely to have a positive exit, and this selection is stronger than for buyout deals, where the corresponding coefficient is 0.5585. The linear age term is positive and significant at 0.0214 in column (1), and the indicator for deals aged eight quarters or less is negative and significant (−0.4136), so real-estate deals follow the same exit pattern as buyout deals, with a low probability of exit during the first two years that rises as the deal matures.

4.3 Estimates for Distressed Deals

Distressed Investments In the MSCI/Burgiss taxonomy, *Distressed debt* investments are a subset of the broader Private Credit category, which includes "[d]ebt investments comprised of loans, bonds, credit derivatives, and other related securities of companies, government entities, or tangible assets." Specifically, these funds make "[d]ebt investments into companies/tangible assets that are under stress or distress, resulting in a substantial discount to the securities' par value."

Distressed funds are younger and have the smallest number of observations in our data, with only 859 holdings, average vintage year of 2012.5, and a mean fund size of $1.64 billion (Table 1). Deals average $48.2 million of invested capital, which is larger than the $31.4 million for real-estate deals but below the $113 million for buyout deals. Distressed funds have a broader sector allocation. Panel A of Table 2 shows their holdings concentrated in Financials (194 deals), Real Estate (131), and Industrials (128), with a further 99 in Energy. Distressed holdings also have shorter holding periods, with a mean of 18.8 and median of 16 quarters (Tables 1 and 3). Performance is also more modest, with a mean gross TVPI of 1.53 and a median of 1.30 (Panel C of Table 3), against 1.76 and 1.71 for real estate and 2.93 and 2.33 for buyout deals. However, negative exits with an almost total loss of the invested capital are rare with only 3.6% of deals, compared to 6.7% for real estate and 15.6% for buyout deals.

Distressed Estimates Table 7 reports parameter estimates for the 859 distressed deals. In the valuation process, the estimated market beta is 0.4112 in the full specification and essentially unchanged at 0.4105 without the one-time jump, which are the lowest betas of the three asset classes. This is consistent with distressed funds' returns being driven more by the resolution of distress than by movements in the market. The estimated quarterly variance σ^2 ranges from 0.0267 to 0.0304, implying an annualized idiosyncratic volatility between 33% and 35%, again the lowest of the three classes. The estimates are, however, less precise, reflecting the smaller sample size. The annual alpha is positive in all three specifications, with 3.56% in the full model, which is similar in magnitude to real estate but well below the 12.92% alpha for buyout deals.

Table 7 also reports the one-time jump. The main estimate in column (1) is not distinguishable from zero. The jump mean is $\mu_j = -0.0012$, with standard deviation 0.0509, implying an arithmetic jump mean of 1.11% with a standard deviation of 0.0520. In the selection equation the coefficient on the latent valuation is positive and significant, at 0.3191 in the main specification in column (1). Better-performing deals are again more likely to exit, but the effect is weaker, consistent with distressed outcomes depending more on the legal and restructuring processes than on the underlying valuation itself. The age coefficients are broadly similar to those for the other two classes. Distressed deals also have a low probability of exit during the first two years that rises as the deal matures, but the coefficients are smaller than for the other two asset classes, leading to a less pronounced pattern.

5. Conclusion

This paper develops and estimates a valuation model for private investments in illiquid deals. Traditional asset-pricing models for traded securities assume that valuations evolve continuously and compounding, typically as a geometric Brownian motion. In contrast, we model the valuation of private investments as a combination of a compounding return and a one-time jump component, with valuations observed only at endogenous exit dates. Our model combines a latent log-valuation process, a selection process to capture the endogenous exit and selection of observed outcomes, and a failure process that imposes that failures arise from poor performance.

To estimate the model, we use a Bayesian Markov Chain Monte Carlo (MCMC) approach using Gibbs sampling and Forward Filtering-Backward Sampling (FFBS) (Carter and Kohn, 1994;

Frühwirth-Schnatter, 1994; Korteweg and Sorensen, 2010). Conditional on the parameters and latent exit indicators, the valuation process forms a linear Gaussian state-space system, allowing us to efficiently sample the entire valuation trajectory. This way, the algorithm delivers the joint posterior distribution of both the latent valuations, latent selection, and failure processes, the parameters governing these processes, as well as non-linear transformations of these parameters, such as the implied alpha.

We estimate the model using holdings-level data from MSCI/Burgiss for three asset classes: buyout, private real estate, and distressed debt. An observation is an individual investment by a fund, which allows us to characterize the different value creation dynamics across these investment types. Outcomes are measured using gross deal TVPI gross of management fees and carried interest.

Across the three samples, real-estate deals carry the highest market beta, at 0.9308 in the main specification, followed by buyout deals at 0.6974 and distressed deals at 0.4112. Idiosyncratic risk runs the other way. Buyout deals have by far the largest quarterly variance, with σ^2 between 0.0733 and 0.0934 and an annualized idiosyncratic volatility of 54% to 61%, matching the wide dispersion of gross TVPIs in Figure 2. Real estate follows at 36% to 38%, consistent with its narrower distribution of outcomes (Figure 3), and distressed deals are lowest at 33% to 35%, reflecting a credit-like payoff whose returns depend more on the resolution of distress than on aggregate equity-market fluctuations. In short, buyout deals carry the most idiosyncratic risk, real-estate deals the most systematic risk, and distressed deals the least of both.

The jump component differs across asset classes. For buyout deals the one-time entry jump is large and positive, with a posterior mean of $\mu_j = 0.1215$ and an implied arithmetic mean of 13.79%, more than five posterior standard deviations from zero. A positive jump is consistent with buyout funds acquiring portfolio companies below their economic value, as discussed in Section 4. For real-estate and distressed deals the estimated jumps are small and not statistically different from zero, with arithmetic means of -2.07% and 1.11% respectively.

The Bayesian framework also delivers posterior distributions for non-linear functions of the primitive parameters, such as the annual alpha and arithmetic mean of the one-time jump. Buyout deals have large positive alphas, 12.92% in the main specification and between 12.92% and 15.48% across the three specifications. Real-estate and distressed alphas are also positive, but much

smaller and less precisely estimated, at 4.24% and 3.56% in their main specifications. These alphas are gross of fees and are computed with the entry jump held separate, so they are not directly comparable with the net-of-fee, fund-level alphas surveyed in Korteweg (2023).

The selection process confirms that exit timing is endogenous to the latent valuation in all three asset classes. The coefficient on the valuation is positive and significant everywhere, so higher-valued deals are more likely to exit. The effect is strongest for real estate (0.7517), then buyout (0.5585), and weakest for distressed deals (0.3191), where legal and restructuring dynamics play a larger role. The age profile has the same shape in all three classes, with a low probability of exit during the first two years followed by a rising hazard as the deal matures. What varies is the strength of that pattern, which is most pronounced for real estate and weakest for distressed deals.

Finally, the model underscores the central role of illiquidity, exit timing, and execution constraints in portfolio choice. In public markets, investors can rebalance frequently and scale exposures quickly. In private markets, however, execution speed is limited by deal flow and exit-market conditions. Our selection process provides a reduced-form representation of these frictions. Incorporating these features into dynamic allocation models is a promising avenue for future research, and we hope our estimates provide empirical inputs for such work.

References

Albert, J. H., Chib, S., 1993. Bayesian analysis of binary and polychotomous response data. Journal of the American Statistical Association 88, 669–679.

Anderson, B. D., Moore, J. B., 1979. Optimal Filtering. Prentice Hall.

Brown, G., Harris, R. S., Hu, W., Jenkinson, T., Kaplan, S. N., Robinson, D. T., 2020. Private equity portfolio companies: A first look at Burgiss holdings data, working paper, University of North Carolina, University of Virginia, University of Oxford, University of Chicago, and Duke University.

Carter, C. K., Kohn, R., 1994. On Gibbs sampling for state space models. Biometrika 81, 541–553.

Case, K. E., Shiller, R. J., 1987. Prices of single-family homes since 1970: New indexes for four cities. New England Economic Review pp. 45–56.

Case, K. E., Shiller, R. J., 1989. The efficiency of the market for single-family homes. American Economic Review 79, 125–137.

Cochrane, J. H., 2005. The risk and return of venture capital. Journal of Financial Economics 75, 3–52.

Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. Econometrica 68, 1343–1376.

Eraker, B., Johannes, M., Polson, N., 2003. The impact of jumps in volatility and returns. Journal of Finance 58, 1269–1300.

Frühwirth-Schnatter, S., 1994. Data augmentation and dynamic linear models. Journal of Time Series Analysis 15, 183–202.

Gelfand, A. E., Smith, A. F. M., 1990. Sampling-based approaches to calculating marginal densities. Journal of the American Statistical Association 85, 398–409.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., Rubin, D. B., 2013. Bayesian Data Analysis. Chapman and Hall/CRC, Boca Raton, FL, third ed.

Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-6, 721–741.

Geyer, C. J., 1992. Practical Markov chain Monte Carlo. Statistical Science 7, 473–483.

Giacoletti, M., 2021. Idiosyncratic risk in housing markets. The Review of Financial Studies 34, 3695–3741.

Ivashina, V., 2023. Private equity financing. In: Eckbo, B. E., Phillips, G. M., Sorensen, M. (eds.), *Handbook of the Economics of Corporate Finance*, Elsevier, vol. 1 of *Handbook of the Economics of Corporate Finance*, pp. 139–160.

Johannes, M., Polson, N. G., 2010. MCMC methods for financial econometrics. In: Ait-Sahalia, Y., Hansen, L. P. (eds.), *Handbook of Financial Econometrics*, Elsevier, vol. 1, pp. 1–72.

Kalman, R. E., 1960. A new approach to linear filtering and prediction problems. Journal of Basic Engineering 82, 35–45.

Korteweg, A., 2023. Risk and return in private equity. In: Eckbo, B. E., Phillips, G. M., Sorensen, M. (eds.), *Handbook of the Economics of Corporate Finance*, Elsevier, vol. 1 of *Handbook of the Economics of Corporate Finance*, pp. 343–410.

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

Korteweg, A., Sorensen, M., 2016. Estimating loan-to-value distributions. Real Estate Economics 44, 41–86.

Leland, H. E., 1994. Corporate debt value, bond covenants, and optimal capital structure. The Journal of Finance 49, 1213–1252.

Lovo, S., Spaenjers, C., 2018. A model of trading in the art market. American Economic Review 108, 744–774.

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

Merton, R. C., 1976. Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics 3, 125–144.

Sagi, J. S., 2021. Asset-level risk and return in real estate investments. The Review of Financial Studies 34, 3647–3694.

Tanner, M. A., Wong, W. H., 1987. The calculation of posterior distributions by data augmentation. Journal of the American Statistical Association 82, 528–540.

A. Estimation procedure

The model for each of the N assets is as follows:

$$v(t) = v(t-1) + r + \delta + \beta' f(t) + \varepsilon_v(t), \tag{A.1}$$

$$w(t) = v(t)\gamma_v + x'(t)\gamma + \eta(t), \tag{A.2}$$

$$z(t) = v(t) + \phi + \xi(t), \tag{A.3}$$

for time $t = 1 \ldots T$. Residuals are i.i.d. with the following distributions

$$\begin{bmatrix} \varepsilon_v(t) \\ \eta(t) \\ \xi(t) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_v^2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \sigma_\xi^2 \end{bmatrix} \right) \tag{A.4}$$

For notational simplicity we suppress the asset index subscript, as well as the time-dependence of the log risk-free rate, r.

The initial latent log-value $v(0) = i+j$, where i is the observed logarithm of the initial investment. The unobserved log jump $j \sim \mathcal{N}\left(\mu_j, \sigma_j^2\right)$ is independent from the residuals of the dynamic process. Together, (A.2) and (A.3) determine the outcome observed in each period. An investment has a positive exit if $w(t) > 0$, and it fails if $z(t) < 0$. When neither condition holds, the investment remains active and its valuation is unobserved.

We are interested in the joint posterior distribution $F(\theta|data)$ of the parameter vector $\theta = \left(\mu_j, \sigma_j^2, \delta, \beta, \sigma_v^2, \gamma_v, \gamma\right)$. For estimation it is convenient to expand this distribution with the latent valuations, selection and failure variables, and jumps, resulting in the augmented posterior $F\left(v(0), \ldots, v(T), w(0), \ldots, w(T), z(0), \ldots, z(T), j, \theta|data\right)$.

We draw from the augmented posterior using a Gibbs sampler (Gelfand and Smith, 1990) with five blocks, containing the valuations, the jumps, the selection variables, the failure variables, and the parameters of the valuation, selection, and jump processes. The Hammersley-Clifford theorem implies that repeated sampling from these blocks, conditional on the most recent draws from the other blocks, produces a Markov chain whose stationary distribution is the augmented posterior. The remainder of this appendix provides details for simulating draws from each block.

Note that the parameters of the failure process, ϕ and σ_ξ^2, are not sampled and therefore have no block of their own. As discussed in Section 1, the outcomes of failed investments are not observed,

so these parameters are only weakly identified, and we fix them at $\phi = -\ln(0.15)$ and $\sigma_\xi = 0.2$ rather than estimating them. They enter the sampler only as known constants, through the failure equation (A.3) and the draws of the latent failure variables.

A.1 Draw latent valuations

The first block draws the time series of latent log-valuations $v(0) \dots v(T)$ conditional on the observed data and the current draws of the selection variables $w(0) \dots w(T)$, jumps, and parameters θ. Based on the assumption that the error terms are independent across assets, we can sample the time series of $v(t)$ separately for each asset.

The idea is to interpret the model as a state space. The state variable is $v(t)$, (A.1) is the transition rule, and (A.2) and (A.3) are (noisy) observation equations. We use Forward Filtering-Backward Sampling (FFBS), an efficient algorithm to sample state variables in a linear state space model (Carter and Kohn, 1994; Frühwirth-Schnatter, 1994).

Intuitively, FFBS first propagates forward the information contained in the observed data and in the latent selection and failure variables, forming a sequence of conditional distributions for $v(t)$, and then samples the entire valuation path in a single backward sweep. Drawing the path jointly in this way avoids the strong serial dependence that would arise if the valuations were sampled one period at a time, and this is what makes the algorithm efficient.

Conditional on the parameters and on the latent selection and failure variables, the system is linear and Gaussian, so the valuation path has the joint conditional density

$$p\left(v(0) \dots v(T) \mid \theta, data\right) \;\propto\; p\left(v(0)\right) \prod_{t=1}^{T} p\left(v(t) \mid v(t-1)\right) \prod_{t=0}^{T} p\left(w(t), z(t) \mid v(t)\right), \tag{A.5}$$

and FFBS samples directly from this $(T+1)$-dimensional Normal distribution. Note that $w(t)$ and $z(t)$ enter as observation densities in *every* period, so periods in which neither an exit nor a failure occurs are informative about $v(t)$ as well.

The following definitions will prove helpful in describing how the algorithm works. Let $m(t|\tau) \equiv E\left(v(t)|w^\tau, \theta\right)$ and $s(t|\tau) \equiv \mathrm{Var}\left(v(t)|w^\tau, \theta\right)$. These are the mean and variance of $v(t)$, both conditional on $w^\tau \equiv \{w(0) \dots w(\tau)\}$ and the current draw of the parameters.

The first stage in FFBS is to compute the distribution of $v(t)$ conditional on available information up to time t, from $t = 0 \dots T$. Since all conditional distributions are Normal, they are fully characterized by $m(t|t)$ and $s(t|t)$, see, e.g., Kalman (1960); Anderson and Moore (1979).

For the first period, we observe the initial investment. The jump is unobserved but Normal with mean μ_j and variance σ_j^2, so $m(0|0) = i + \mu_j$ and variance $s(0|0) = \sigma_j^2$. We filter forward one period at a time using the Kalman filter (as the model is linear). Starting with $t = 1$, this iterates over two steps.

Step 1: Forecast

$$m(t|t-1) = m(t-1|t-1) + r + \delta + \beta' f(t), \tag{A.6}$$

$$s(t|t-1) = s(t-1|t-1) + \sigma_v^2. \tag{A.7}$$

Step 2: Update

$$m(t|t) = m(t|t-1) \cdot (1 - K \cdot H) + K \cdot \begin{bmatrix} w(t) - x'(t)\gamma \\ z(t) - \phi \end{bmatrix}, \tag{A.8}$$

$$s(t|t) = s(t|t-1) \cdot (1 - K \cdot H), \tag{A.9}$$

where $H \equiv \begin{bmatrix} \gamma_v \\ 1 \end{bmatrix}$, and

$$K \equiv s(t|t-1) \cdot H' \cdot \left(HH' \cdot s(t|t-1) + \begin{bmatrix} 1 & 0 \\ 0 & \sigma_\xi^2 \end{bmatrix} \right)^{-1} \tag{A.10}$$

is the 1×2 Kalman gain vector, which determines how much information from each of the two observation equations is incorporated into $v(t)$. Notably, the first element of K is larger (in absolute terms) when γ_v is large (also in absolute terms), because a large value of γ_v means that the selection equation contains more information about v.

If the time-t valuation is observed, then instead of (A.8) and (A.9), m(t | t) is equal to the observed log value and s(t | t) = 0.

Next, iterate over these two steps, incrementing t one period at a time until T. There are three possible scenarios for the final period of a given asset:

- observed exit time at observed value: $m(T|T)$ equals the observed exit value and $s(T|T) = 0$. No further sampling beyond T.
- observed exit time at unobserved value: T is the known exit time, and $m(T|T)$ and $s(T|T)$ are as computed from (A.8) and (A.9). No further sampling beyond T.

To summarize, the forward filter provides the distribution $p(v(t)|w^t)$ for all $t = 0 \dots T$, suppressing the conditioning on θ from here on out for notational simplicity. However, what we need is the distribution of $v(t)$ conditional on the full time series w^T (also for all t). Backward sampling, the second stage of the FFBS algorithm, solves this problem. Per lemma 2.1 in Carter and Kohn (1994),

$$p\left(v(0) \dots v(T)|w^T\right) = p\left(v(T)|w^T\right) \prod_{t=0}^{T-1} p\left(v(t)|w^t, v(t+1)\right). \tag{A.11}$$

This equation implies that we can start with sampling $v(T)|w^T$, for which we know the distribution from the final step of the forward filter. Then we sample $v(T-1)$ conditional on w^{T-1} and the draw of $v(T)$. Next, we sample $T-2$, and continue sampling backwards until we reach time 0.

The distribution $p\left(v(t)|w^t, v(t+1)\right)$ can be found by starting from $p\left(v(t)|w^t\right)$ and interpreting $v(t+1)$ as another noisy observation of $v(t)$. This is essentially another filter update step, but this time using equation (A.1) as the observation equation. It follows that

$$p\left(v(t)|w^t, v(t+1)\right) \sim \mathcal{N}\left(m(t|t) + G \cdot [v(t+1) - m(t+1|t)], s(t|t) \cdot [1-G]\right), \tag{A.12}$$

with the gain

$$G \equiv \frac{s(t|t)}{s(t|t) + \sigma_v^2}. \tag{A.13}$$

After repeating the above algorithm for all assets, move on to the next block.

A.2 Draw jump parameters

Conditional on the valuation draws from the previous step, the jump for each asset can be computed from the difference between $v(0)$ and i. Let J be the $N \times 1$ vector that stacks $v(0) - i$ across all assets (note that $v(0)$ is the log-valuation at the time of the initial investment, which may be at different calendar dates for different assets). The jump parameters μ_j and σ_j^2 are the mean and variance of J. Their posterior distributions can be derived from a Bayesian regression of J on a constant.

Prior: We use a conjugate prior, the Normal-scaled inverse chi-squared:

$$\sigma_j^2 \sim \frac{\nu_0 s_0^2}{\chi^2(\nu_0)}, \tag{A.14}$$

$$\mu_j|\sigma_j^2 \sim \mathcal{N}\left(\mu_0, \sigma_j^2/\kappa_0\right), \tag{A.15}$$

where $\chi^2(\nu)$ is the chi-squared distribution with ν degrees of freedom.

Hyperparameters: We set $\mu_0 = 0$, $\kappa_0 = \nu_0 = 2$, and $s_0^2 = 0.03$.

Posterior: Since the only explanatory variable is a constant, the posterior is

$$\sigma_j^2|J \sim \frac{\nu s^2}{\chi^2(\nu)}, \tag{A.16}$$

$$\mu_j|\sigma_j^2, J \sim \mathcal{N}\left(\mu, \sigma_j^2/(\kappa_0 + N)\right), \tag{A.17}$$

with

$$\nu = \nu_0 + N, \tag{A.18}$$

$$s^2 = \frac{\nu_0 s_0^2 + (J - \bar{J})'(J - \bar{J}) + \frac{\kappa_0 N}{\kappa_0 + N}\left(\bar{J} - \mu_0\right)^2}{v_0 + N}, \tag{A.19}$$

$$\mu = \frac{\kappa_0}{\kappa_0 + N}\mu_0 + \frac{N}{\kappa_0 + N}\bar{J}, \tag{A.20}$$

where $\bar{J} = J'1/N$ is the sample average.

A.3 Draw latent selection variables

This step follows Korteweg and Sorensen (2010), and is similar to sampling the augmented posterior distribution of a probit model (Albert and Chib, 1993). For each period t, starting with the time of the first investment and repeating until exit (or the final sampling period), the selection variable $w(t)$ is sampled conditional on the current realizations of the valuation, the model parameters, and the information whether or not a valuation for the asset is observed at time t. Since $\eta(t)$ is independent over time, each period can be sampled separately.

Let $\mathcal{LTN}(\mu, \sigma^2)$ and $\mathcal{UTN}(\mu, \sigma^2)$ be the lower and upper-truncated normal distributions with mean μ and variance σ^2, truncated at zero. Depending on whether $v(t)$ is observed, $w(t)$ is sampled from

$$w(t)|x(t), v(t), \gamma \sim \begin{cases} \mathcal{LTN}\left(v(t)\gamma_v + x'(t)\gamma, 1\right) & \text{if } v(t) \text{ is observed,} \\ \mathcal{UTN}\left(v(t)\gamma_v + x'(t)\gamma, 1\right) & \text{if } v(t) \text{ is not observed.} \end{cases} \tag{A.21}$$

This step is repeated for all assets before moving on to the next block.

A.4 Draw latent failure variables

Sampling the latent failure variable $z(t)$ works similarly to the selection variable $w(t)$. For each period t from initial investment until exit (or the end of the sample period), we draw $z(t)$

conditional on the current valuation draws, model parameters, and whether a failure is observed in period t. Sampling is done period by period and asset by asset, since $\xi(t)$ is independent across both time and assets.

The sampling distribution is

$$z(t)|v(t),\phi,\sigma_\xi^2 \sim \begin{cases} \mathcal{UTN}\left(v(t)+\phi,\sigma_\xi^2\right) & \text{if failure is observed,} \\ \mathcal{LTN}\left(v(t)+\phi,\sigma_\xi^2\right) & \text{if failure is not observed.} \end{cases} \tag{A.22}$$

Note that failure occurs when $z(t) < 0$, whereas a positive exit is observed when $w(t) > 0$. For this reason, the truncation (upper versus lower) appears reversed here, compared to the previous block.

A.5 Draw valuation and selection equation parameters

This step follows Korteweg and Sorensen (2010). The parameters of the valuation evolution (A.1) and the selection equation (A.2) can be sampled separately from each other, since their residuals are independent and the two equations do not have any parameters in common. Conditional on the current draws of v(t) and w(t), the posterior distributions of the parameters for both equations are found from Bayesian linear regressions.

For the valuation equation, let Y be the stacked vector of excess returns $v(t)-v(t-1)-r$ across all periods for which there are draws of $v(t)$ and its lagged value, and across all assets. Similarly, let X be the matrix with a constant term in the first column and the stacked values of $f'(t)$ in the other columns. Thus, if we observe excess returns for $N(t)$ assets at time t then Y is of dimension $\sum_t N(t) \times 1$ and X has $\sum_t N(t)$ rows and $K+1$ columns, with K the number of factors in $f(t)$.

Prior: The prior distribution is the Normal-scaled inverse chi-squared:[5]

$$\sigma_v^2 \sim \frac{\nu_0 s_0^2}{\chi^2(\nu_0)}, \tag{A.23}$$

$$\delta,\beta|\sigma_v^2 \sim \mathcal{N}\left(\mu_0,\sigma_v^2 A^{-1}\right). \tag{A.24}$$

Hyperparameters: We set $\mu_0 = \left[\delta_0,\beta_0'\right] = [0,1_K']$, $A = 0.04 \cdot I_{K+1}$, $\nu_0 = 0$.

[5] We specify a scaled inverse chi-squared distribution as the prior distribution for variance, whereas Korteweg and Sorensen (2010) use an inverse gamma distribution. They are equivalent: The scaled inverse chi-squared with ν degrees of freedom and scale s^2, i.e., $\frac{\nu s^2}{\chi^2(\nu)}$, corresponds to an inverse gamma distribution with shape parameter $\alpha = \frac{\nu}{2}$ and scale parameter $\beta = \frac{\nu s^2}{2}$.

Posterior:

$$\sigma_v^2|Y,X \sim \frac{\nu s^2}{\chi^2(\nu)}, \tag{A.25}$$

$$\delta,\beta|Y,X \sim \mathcal{N}\left(\mu, \sigma_v^2 (A + X'X)^{-1}\right), \tag{A.26}$$

with

$$\nu = \nu_0 + \sum_t N(t), \tag{A.27}$$

$$s^2 = \frac{v_0 s_0^2 + (Y - X\mu)' (Y - X\mu) + (\mu - \mu_0)' A (\mu - \mu_0)}{v_0 + \sum_t N(t)}, \tag{A.28}$$

$$\mu = (A + X'X)^{-1} (A\mu_0 + X'Y). \tag{A.29}$$

The selection equation follows the same regression setup but is substantially simpler due to the variance of $\eta(t)$ being unity. Let W be the stacked vector of $w(t)$ across all relevant periods and across all assets, and let X be the matrix of corresponding stacked $[v_t, x'(t)]$.
Prior: The prior distribution for the regression coefficients is

$$\gamma_v,\gamma \sim \mathcal{N}\left(c_0, \Omega^{-1}\right). \tag{A.30}$$

Hyperparameters: We set $c_0 = 0'_{M+1}$ and $\Omega = 0.04 \cdot I_{M+1}$, where M is the number of variables in $x(t)$.
Posterior:

$$\gamma_v,\gamma|W,X \sim \mathcal{N}\left(c, (\Omega + X'X)^{-1}\right), \tag{A.31}$$

with

$$c = (\Omega + X'X)^{-1} (\Omega c_0 + X'W). \tag{A.32}$$

B. Derivation of Updating Step

Below is an alternative, detailed derivation of the updating step in the valuation block of the FFBS procedure. The updating conditions on the latent selection variable, w, and there are separate derivations with and without additionally conditioning on the failure process, z. For this derivation, the notation and setup is:

Valuation:

$$v \sim \mathcal{N}(\mu_v, \sigma_v^2).$$

Selection Process:

$$w \;=\; \gamma x \;+\; \gamma_v\, v \;+\; \eta, \qquad \eta \sim \mathcal{N}(0,1).$$

Failure Process:

$$z \;=\; \phi \;+\; v \;+\; \xi, \qquad \xi \sim \mathcal{N}(0, \sigma_\xi^2).$$

The vector (v, w, z) is jointly Normal distributed with

$$\mu_{vwz} \;=\; \begin{pmatrix} \mu_v \\ \gamma x + \gamma_v\, \mu_v \\ \phi + \mu_v \end{pmatrix}, \qquad \Sigma_{vwz} \;=\; \begin{pmatrix} \sigma_v^2 & \gamma_v\, \sigma_v^2 & \sigma_v^2 \\ \gamma_v\, \sigma_v^2 & 1 + \gamma_v^2 \sigma_v^2 & \gamma_v\, \sigma_v^2 \\ \sigma_v^2 & \gamma_v\, \sigma_v^2 & \sigma_v^2 + \sigma_\xi^2 \end{pmatrix}.$$

The conditional distribution of $v \mid w$, i.e., without conditioning on the failure Process is:

$$v \mid w \sim \mathcal{N}(\mu, \sigma^2),$$

with

$$\mu = \mu_v + \frac{\gamma_v\, \sigma_v^2}{1 + \gamma_v^2 \sigma_v^2} \Big(\gamma_v (v - \mu_v) + \eta \Big),$$

and

$$\sigma^2 = \sigma_v^2 - \frac{\gamma_v^2 \sigma_v^4}{1 + \gamma_v^2 \sigma_v^2} = \frac{\sigma_v^2}{1 + \gamma_v^2 \sigma_v^2}.$$

The conditional distribution of $v \mid w, z$, i.e., with conditioning on the failure process is:

$$v \mid w, z \sim \mathcal{N}(\mu, \sigma^2),$$

where

$$\mu = \mu_v + \begin{pmatrix} \gamma_v \sigma_v^2 & \sigma_v^2 \end{pmatrix} \begin{bmatrix} 1 + \gamma_v^2 \sigma_v^2 & \gamma_v \sigma_v^2 \\ \gamma_v \sigma_v^2 & \sigma_v^2 + \sigma_\xi^2 \end{bmatrix}^{-1} \begin{pmatrix} \gamma_v (v - \mu_v) + \eta \\ v - \mu_v + \xi \end{pmatrix},$$

and the conditional variance is

$$\sigma^2 = \sigma_v^2 - \begin{pmatrix} \gamma_v \sigma_v^2 & \sigma_v^2 \end{pmatrix} \begin{bmatrix} 1 + \gamma_v^2 \sigma_v^2 & \gamma_v \sigma_v^2 \\ \gamma_v \sigma_v^2 & \sigma_v^2 + \sigma_\xi^2 \end{bmatrix}^{-1} \begin{pmatrix} \gamma_v \sigma_v^2 \\ \sigma_v^2 \end{pmatrix}.$$

Expanding the inverse yields the more compact expression

$$\sigma^2 = \frac{\sigma_v^2 \sigma_\xi^2}{\sigma_v^2 + \sigma_\xi^2 + \gamma_v^2 \sigma_v^2 \sigma_\xi^2},$$

Table 1: Summary Descriptive Statistics for Funds and Holdings Fund statistics weighted by deal. The sample includes North American private-market holdings (deals) from the MSCI (Burgiss) database, split into three asset classes: Buyout (N = 7,243), Real Estate (N = 4,447), and Distressed (N = 859). A deal is included if it (i) was first observed (entered) on or before 2020Q4, (ii) has a holding period of 5–60 quarters, and (iii) reached a realized outcome, *i.e.*, an observed exit date together with an observed exit valuation. Deals that are still active/unrealized, and exited deals without a reported valuation, are excluded. Failure shares are 15.6% / 6.7% / 3.6% for Buyout / Real Estate / Distressed.

	Buyout		Real Estate		Distressed	
	Mean	Std.dev.	Mean	Std.dev.	Mean	Std.dev.
Fund						
Vintage	2008.0	6.1	2010.6	5.5	2012.5	4.2
Size (USDbn)	$2.25	$3.50	$1.41	$2.03	$1.64	$1.12
Holding						
Total Inv (USDmm)	$113.0	$166.2	$31.4	$86.5	$48.2	$76.6
Start (qtr)	201.7	24.2	208.9	22.0	218.0	15.4
End (qtr)	225.4	23.6	228.9	20.6	236.8	14.6
Holding Per. (qtr)	23.6	11.6	19.9	10.6	18.8	11.1
Gross TVPI	2.9	4.1	1.8	1.2	1.5	1.5
Gross IRR	24.8%	138.5%	18.4%	37.2%	15.0%	36.8%
Observations	7,243		4,447		859	

Table 2: Holdings: Sectors and Investment Year Panel A shows the number of holdings (deals) in the estimation sample by sector, as classified by MSCI/Burgiss and asset class. Panel B shows the number of of holdings (deals) in the estimation sample by investment (entry) year and asset class

Panel A: Sectors of Holdings (Deals)

Holding Sector	Buyout	Real Estate	Distressed	Total
Communication Servi..	502	0	47	549
Consumer Discretion..	1,195	13	76	1,284
Consumer Staples	421	0	25	446
Energy	255	0	99	354
Financials	504	80	194	778
Health Care	1,039	5	29	1,073
Industrials	1,682	4	128	1,814
Information Technol..	1,093	1	33	1,127
Materials	387	0	40	427
Real Estate	57	4,337	131	4,525
Unknown	67	7	23	97
Utilities	41	0	34	75
Total	7,243	4,447	859	12,549

Panel B: Holdings (Deals) by Investment Year

Year	Buyout	Real Estate	Distressed	Total
1984	9	0	0	9
1985	0	0	0	0
1986	0	0	0	0
1987	5	0	0	5
1988	1	0	0	1
1989	7	0	0	7
1990	3	0	0	3
1991	4	0	0	4
1992	9	0	0	9
1993	8	1	0	9
1994	14	1	0	15
1995	24	7	0	31
1996	40	14	0	54
1997	67	20	0	87
1998	79	39	0	118
1999	122	48	0	170
2000	205	45	0	250
2001	91	21	0	112
2002	125	32	0	157
2003	193	79	0	272
2004	254	132	3	389
2005	315	200	5	520
2006	407	257	17	681
2007	552	282	26	860
2008	403	161	19	583
2009	261	43	24	328
2010	411	158	67	636
2011	382	268	66	716
2012	504	364	81	949
2013	324	286	89	699
2014	455	299	62	816
2015	414	363	83	860
2016	397	311	63	771
2017	401	276	58	735
2018	344	301	42	687
2019	242	269	54	565
2020	171	170	100	441
Total	7,243	4,447	859	12,549

Table 3: Holdings Detailed Statistics

Panel A: Holdings by Buyout Funds

	Obs.	Mean	Std.dev.	min	p25	p50	p75	max
Total Inv (USDmm)	4,638	$112.97	$166.19	$0.00	$23.06	$53.62	$134.57	$2,391.10
Start (qtr)	7,243	201.75	24.24	98	186	204	221	243
End (qtr)	7,243	225.35	23.60	110	211	229	244	261
Holding Period (qtr)	7,243	23.61	11.65	5	15	21	31	60
Gross TVPI	7,243	2.93	4.06	0.00	0.92	2.33	3.77	90.00
Gross IRR	5,831	0.25	1.39	-1.00	0.02	0.22	0.44	50.55

Panel B: Holdings by Real Estate Funds

	Obs.	Mean	Std.dev.	min	p25	p50	p75	max
Total Inv (USDmm)	2,094	$31.44	$86.46	$0.00	$8.91	$16.30	$30.49	$2,355.59
Start (qtr)	4,447	208.94	22.03	135	190	212	226	243
End (qtr)	4,447	228.87	20.64	157	218	231	246	261
Holding Period (qtr)	4,447	19.93	10.64	5	12	17	26	60
Gross TVPI	4,447	1.76	1.19	0.00	1.21	1.71	2.17	21.17
Gross IRR	3,082	0.18	0.37	-1.00	0.08	0.19	0.31	4.88

Panel C: Holdings by Distressed Funds

	Obs.	Mean	Std.dev.	min	p25	p50	p75	max
Total Inv (USDmm)	335	$48.20	$76.58	$0.01	$5.00	$18.22	$55.04	$528.49
Start (qtr)	859	218.02	15.45	179	207	217	231	243
End (qtr)	859	236.84	14.64	189	227	239	248	261
Holding Period (qtr)	859	18.82	11.14	5	10	16	26	60
Gross TVPI	859	1.53	1.53	0.00	1.04	1.30	1.62	18.00
Gross IRR	550	0.15	0.37	-1.00	0.02	0.15	0.28	2.82

Table 4: Estimates with Simulated Data The table reports estimates from twenty-four independently simulated datasets. Each dataset has N = 7,200 observations and is generated from parameters calibrated to those estimated using the Buyout sample (Specification 1 in Table 5). Entry values are normalized and entry dates are confined so that every holding resolves inside the sample. Realized failure rates are in the range 10.8–11.9% across datasets. The estimation runs 350,000 iterations per dataset, of which the first third is discarded as burn-in. Calendar fixed effects are not included in this design. The True column shows the true parameters in the data-generating process. The Estimates column shows average posterior means across the twenty-four datasets, with the between-dataset standard deviation in parentheses. Within SD is the average within-dataset posterior standard deviation. Alpha and its Within SD are annualized. Ratio is the between-dataset standard deviation of the point estimates (net of MCMC noise) divided by the average within-dataset posterior standard deviation. A ratio of 1.0 means that the posterior standard deviation for each dataset equals the standard deviation of the point estimates across datasets. A ratio below 1.0 means the estimates vary less across datasets than the posterior standard deviation implies, so the posterior understates the precision of the estimates.

	True	Estimates	Within SD	Ratio
Valuation Process				
Delta	-0.0038	-0.0036	0.0014	0.88
		(0.0012)		
Beta	0.6970	0.7014	0.0395	0.94
		(0.0373)		
Sigma^2	0.0733	0.0729	0.0012	0.87
		(0.0011)		
Alpha	0.1290	0.1290	0.0062	0.89
		(0.0056)		
Jump Mean	0.1220	0.1159	0.0207	0.95
		(0.0198)		
Jump Sigma^2	0.0147	0.0219	0.0087	0.65
		(0.0057)		
Jump	0.1381	0.1357	0.0241	0.93
		(0.0226)		
Selection Process				
Constant	-2.4700	-2.4715	0.0199	0.92
		(0.0183)		
Valuation	0.5590	0.5616	0.0111	1.13
		(0.0126)		
Age	0.0266	0.0268	0.0007	0.81
		(0.0006)		
Age ≤ 8	-0.2770	-0.2783	0.0227	0.90
		(0.0204)		

Table 5: Estimates for Buyout Deals The table shows estimates for buyout holdings (deals) in the estimation sample (see Table 1) with time fixed effects in the selection/exit process. Specifications in the three columns are (1) the full model, with both the one-time entry jump and the probabilistic failure process (run R50); (2) the same model with the entry jump removed, retaining the failure process (run R53); (3) both the entry jump and the failure process removed (run R46). Each reports the posterior mean with the posterior standard deviation in parentheses. ***, **, and * indicate that the posterior mean lies more than 2.576, 1.96, and 1.645 posterior standard deviations from zero (equivalently, the 99%, 95%, and 90% credible interval excludes zero). Significance markers are omitted for the variance parameters (Sigma^2, Jump Sigma^2), which are positive by construction.

	(1) Jump + Failure	(2) Failure, No Jump	(3) No Failure, No Jump
Valuation Process			
Delta	-0.0038** (0.0015)	0.0019* (0.0010)	-0.0143*** (0.0014)
Beta	0.6974*** (0.0337)	0.6744*** (0.0325)	1.4235*** (0.0520)
Sigma²	0.0733 (0.0011)	0.0750 (0.0010)	0.0934 (0.0016)
Alpha (annual)	0.1292*** (0.0068)	0.1548*** (0.0044)	0.1364*** (0.0052)
One-time Jump			
Jump Mean	0.1215*** (0.0242)	—	—
Jump Sigma²	0.0147 (0.0058)	—	—
Jump (Arithmetic)	0.1379*** (0.0284)	—	—
Selection Process			
Constant	-2.4714*** (0.0224)	-2.4740*** (0.0227)	-2.1306*** (0.0184)
Valuation	0.5585*** (0.0121)	0.5911*** (0.0106)	0.2039*** (0.0066)
Age	0.0266*** (0.0007)	0.0266*** (0.0007)	0.0233*** (0.0006)
1(Age ≤ 8 qtr)	-0.2769*** (0.0241)	-0.2495*** (0.0240)	-0.3774*** (0.0211)
Time Fixed Effects (reference: 2015–2019)			
1980-1994	-0.7745*** (0.1941)	-0.7721*** (0.1997)	-0.7788*** (0.1763)
1995-1999	-0.2914*** (0.0703)	-0.2959*** (0.0715)	-0.2095*** (0.0601)
2000-2004	-0.1726*** (0.0372)	-0.1675*** (0.0378)	-0.0202 (0.0285)
2005-2009	-0.1322*** (0.0237)	-0.1291*** (0.0241)	-0.0579*** (0.0201)
2010-2014	-0.1354*** (0.0182)	-0.1346*** (0.0185)	-0.1169*** (0.0158)
2020-2024	0.2314*** (0.0178)	0.2233*** (0.0178)	0.2886*** (0.0156)
Memo			
Holdings (N)	7,243	7,243	7,243
Failures in sample	15.6%	15.6%	15.6%
Iterations	350,000	350,000	200,000
Minimum n_eff	115	1,117	1,662

Table 6: Estimates for Real Estate Holdings

	(1) Jump + Failure	(2) Failure, No Jump	(3) No Failure, No Jump
Valuation Process			
Delta	-0.0057*** (0.0018)	-0.0071*** (0.0010)	-0.0184*** (0.0012)
Beta	0.9308*** (0.0357)	0.9369*** (0.0355)	1.4181*** (0.0444)
Sigma²	0.0331 (0.0009)	0.0330 (0.0007)	0.0356 (0.0008)
Alpha (annual)	0.0424*** (0.0080)	0.0372*** (0.0040)	0.0044 (0.0044)
One-time Jump			
Jump Mean	-0.0242 (0.0267)	—	—
Jump Sigma²	0.0057 (0.0018)	—	—
Jump (Arithmetic)	-0.0207 (0.0261)	—	—
Selection Process			
Constant	-2.0327*** (0.0260)	-2.0338*** (0.0260)	-1.8620*** (0.0232)
Valuation	0.7517*** (0.0264)	0.7399*** (0.0194)	0.3387*** (0.0131)
Age	0.0214*** (0.0010)	0.0215*** (0.0010)	0.0192*** (0.0009)
1(Age ≤ 8 qtr)	-0.4136*** (0.0267)	-0.4170*** (0.0258)	-0.4772*** (0.0246)
Time Fixed Effects (reference: 2015–2019)			
1995-1999	-0.6361*** (0.1633)	-0.6370*** (0.1620)	-0.5870*** (0.1544)
2000-2004	-0.1045** (0.0491)	-0.1043** (0.0486)	-0.1315*** (0.0450)
2005-2009	-0.1815*** (0.0331)	-0.1833*** (0.0332)	-0.1544*** (0.0294)
2010-2014	-0.1815*** (0.0231)	-0.1832*** (0.0230)	-0.1484*** (0.0208)
2020-2024	0.1943*** (0.0209)	0.1959*** (0.0207)	0.2304*** (0.0191)
Memo			
Holdings (N)	4,447	4,447	4,447
Failures in sample	6.7%	6.7%	6.7%
Iterations	350,000	350,000	200,000
Minimum n_eff	114	1,243	1,054

Table 7: Estimates for Distressed Holdings

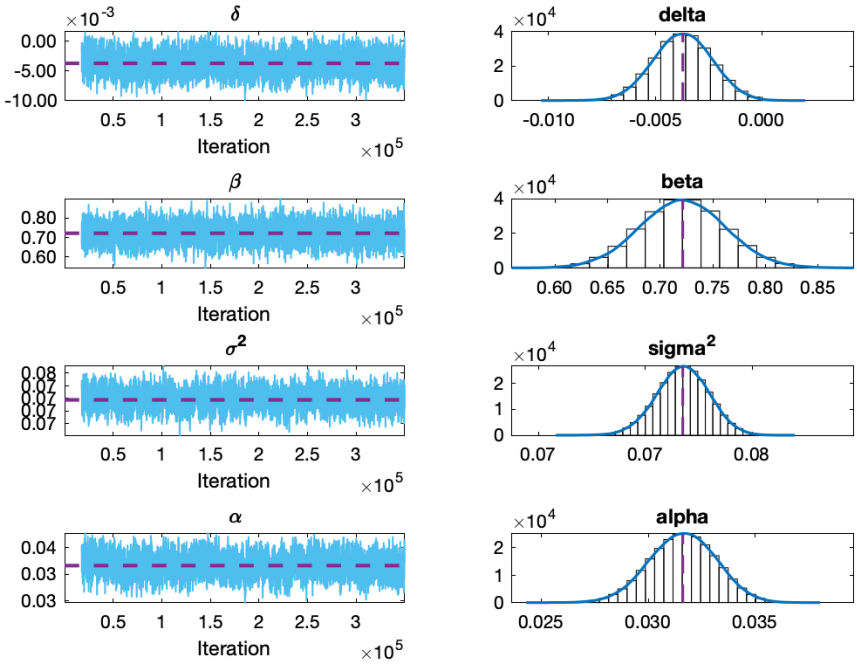
	(1) Jump + Failure	(2) Failure, No Jump	(3) No Failure, No Jump
Valuation Process			
Delta	-0.0038 (0.0041)	-0.0038 (0.0032)	-0.0082** (0.0037)
Beta	0.4112*** (0.1068)	0.4105*** (0.1040)	0.5569*** (0.1199)
Sigma²	0.0267 (0.0017)	0.0283 (0.0013)	0.0304 (0.0015)
Alpha (annual)	0.0356** (0.0176)	0.0388*** (0.0136)	0.0252* (0.0148)
One-time Jump			
Jump Mean	-0.0012 (0.0509)	—	—
Jump Sigma²	0.0217 (0.0124)	—	—
Jump (Arithmetic)	0.0111 (0.0520)	—	—
Selection Process			
Constant	-1.8783*** (0.0526)	-1.8810*** (0.0525)	-1.8757*** (0.0508)
Valuation	0.3191*** (0.0564)	0.3254*** (0.0387)	0.1610*** (0.0322)
Age	0.0112*** (0.0022)	0.0113*** (0.0022)	0.0131*** (0.0021)
1(Age ≤ 8 qtr)	-0.2021*** (0.0521)	-0.2013*** (0.0523)	-0.1928*** (0.0510)
Time Fixed Effects (reference: 2015–2019)			
2005-2009	-0.5529*** (0.1746)	-0.5509*** (0.1741)	-0.5631*** (0.1702)
2010-2014	-0.0959* (0.0494)	-0.0957* (0.0495)	-0.0793* (0.0480)
2020-2024	0.3970*** (0.0396)	0.3974*** (0.0396)	0.3967*** (0.0388)
Memo			
Holdings (N)	859	859	859
Failures in sample	3.6%	3.6%	3.6%
Iterations	350,000	350,000	200,000
Minimum n_eff	127	1,162	1,434

Table 8: Robustness of Buyout Estimates The table show estimates for the Buyout estimation sample. Each cell reports the posterior mean with the posterior standard deviation in parentheses. ***, **, and * indicate that the posterior mean lies more than 2.576, 1.96, and 1.645 posterior standard deviations from zero. Markers are omitted for the variance parameters (Sigma², Jump Sigma²), which are positive by construction. Time fixed effects are estimated in every column but are not reported. Columns (1)–(3) vary the failure barrier, from TVPI 0.10 (run R56), 0.15 (run R50) and to 0.20 (run R57), where φ = −log(TVPI). Columns (4)–(6) vary the failure variance, with σ²_ξ = 0.01 (run R58), 0.04 (run R50), and 0.09 (run R59). Columns (2) and (5) are therefore the same baseline run so each sweep reads left to right.

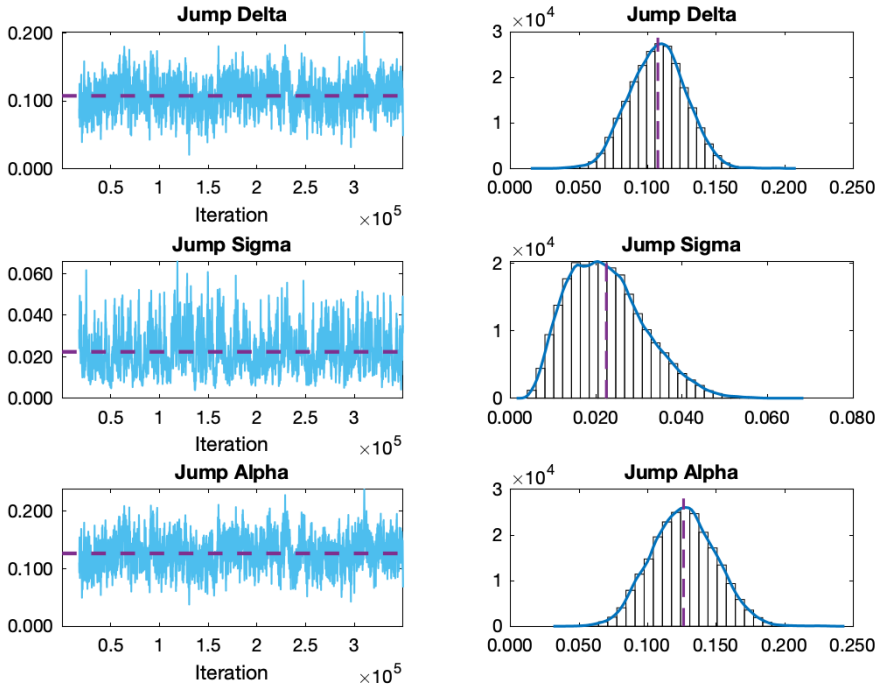
	(1) TVPI 0.10	(2) TVPI 0.15	(3) TVPI 0.20	(4) σ²_ξ 0.01	(5) σ²_ξ 0.04	(6) σ²_ξ 0.09
Valuation Process						
Delta	0.0044***	-0.0038**	-0.0110***	-0.0013	-0.0038**	-0.0073***
	(0.0016)	(0.0015)	(0.0014)	(0.0015)	(0.0015)	(0.0015)
Beta	0.7800***	0.6974***	0.6282***	0.6735***	0.6974***	0.7152***
	(0.0367)	(0.0337)	(0.0330)	(0.0360)	(0.0337)	(0.0332)
Sigma²	0.0926	0.0733	0.0607	0.0776	0.0733	0.0674
	(0.0013)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)
Alpha (annual)	0.2008***	0.1292***	0.0748***	0.1476***	0.1292***	0.1032***
	(0.0072)	(0.0068)	(0.0064)	(0.0064)	(0.0068)	(0.0068)
One-time Jump						
Jump Mean	-0.1805***	0.1215***	0.3742***	0.0577**	0.1215***	0.2192***
	(0.0246)	(0.0242)	(0.0236)	(0.0226)	(0.0242)	(0.0244)
Jump Sigma²	0.0132	0.0147	0.0361	0.0115	0.0147	0.0221
	(0.0044)	(0.0058)	(0.0096)	(0.0055)	(0.0058)	(0.0078)
Jump (Arithmetic)	-0.1593***	0.1379***	0.4808***	0.0657***	0.1379***	0.2593***
	(0.0203)	(0.0284)	(0.0371)	(0.0249)	(0.0284)	(0.0326)
Selection Process						
Constant	-2.4732***	-2.4714***	-2.4471***	-2.4753***	-2.4714***	-2.4650***
	(0.0231)	(0.0224)	(0.0224)	(0.0229)	(0.0224)	(0.0224)
Valuation	0.6390***	0.5585***	0.4763***	0.5788***	0.5585***	0.5282***
	(0.0119)	(0.0121)	(0.0126)	(0.0120)	(0.0121)	(0.0126)
Age	0.0280***	0.0266***	0.0252***	0.0269***	0.0266***	0.0261***
	(0.0008)	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0007)
1(Age ≤ 8 qtr)	-0.2235***	-0.2769***	-0.3233***	-0.2647***	-0.2769***	-0.2954***
	(0.0253)	(0.0241)	(0.0234)	(0.0246)	(0.0241)	(0.0238)
Memo						
Failures in sample	15.6%	15.6%	15.6%	15.6%	15.6%	15.6%
Minimum n_eff	104	115	132	40	115	199

Figure 1: Estimated Parameters using Simulated Data (R131) Posterior draws for run R131, one of the twenty-four simulated datasets pooled in Table 4 (N = 7,200; 350,000 Gibbs iterations, of which the first third is discarded as burn-in). Left panels plot draws against iteration number, and right panels show corresponding histograms with an overlaid kernel density. Dashed lines show posterior means. Panel A shows parameters for the valuation process, Panel B for the one-time jump specification, and Panel C for the selection process.

Panel A: Parameters for Valuation Process



Panel B: Parameters for One-time Jump



Panel C: Parameters for Selection Process

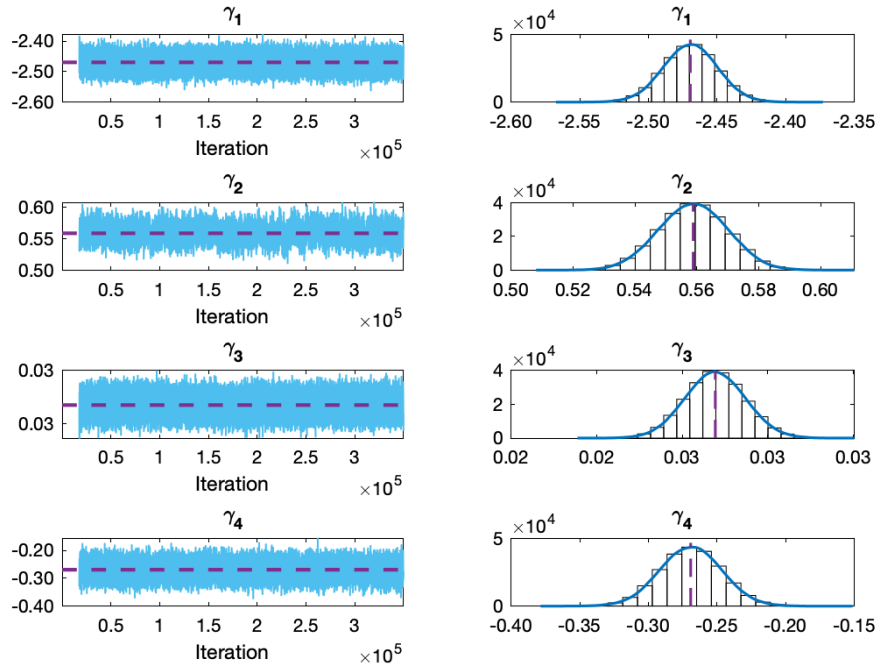
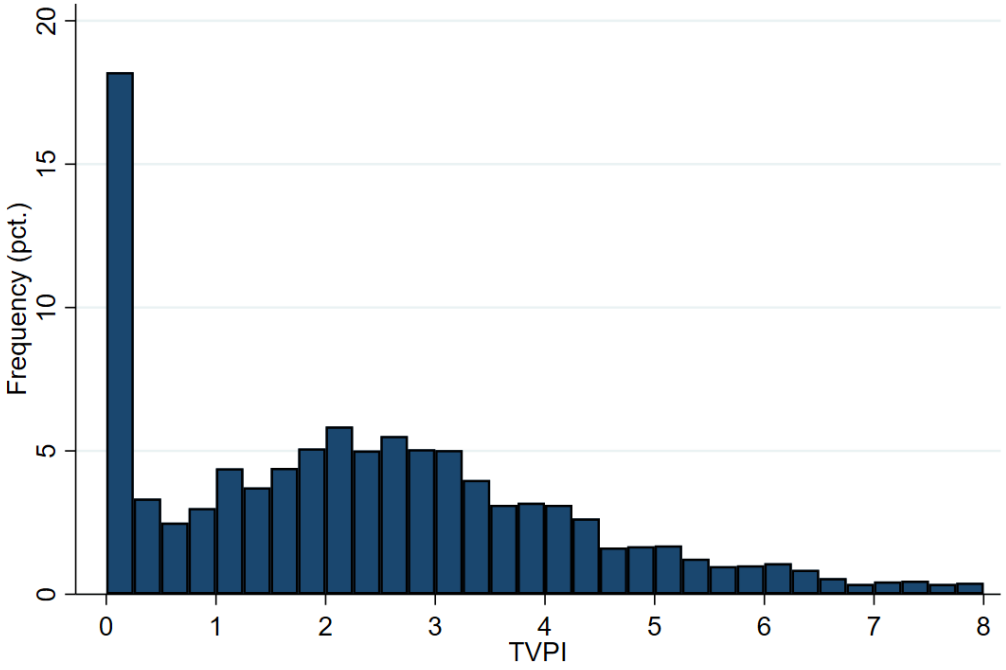
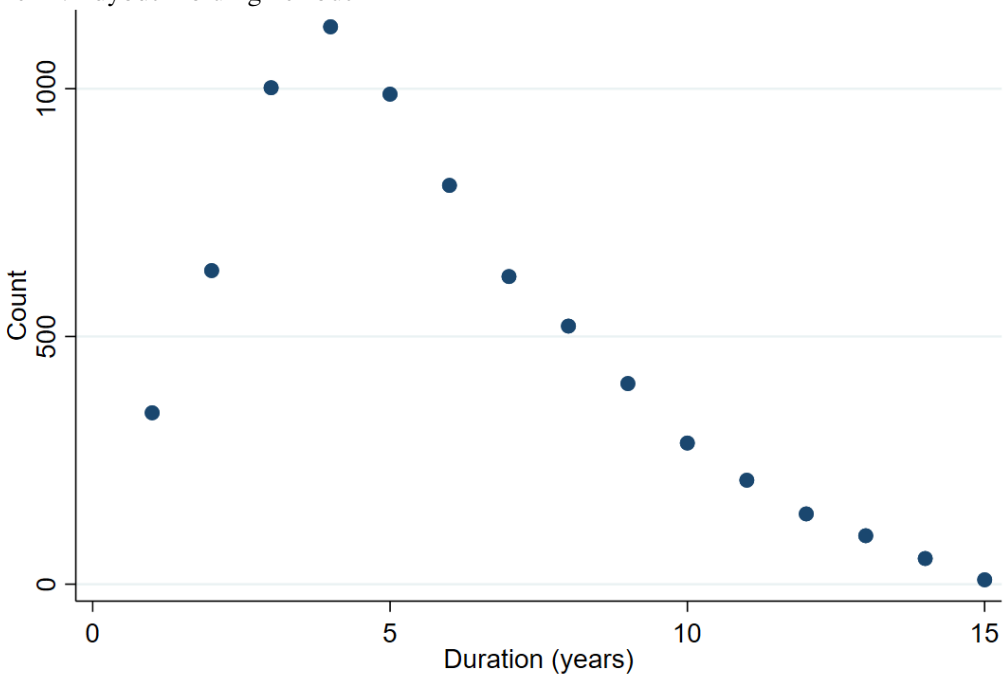


Figure 2: Holdings by Buyout Funds

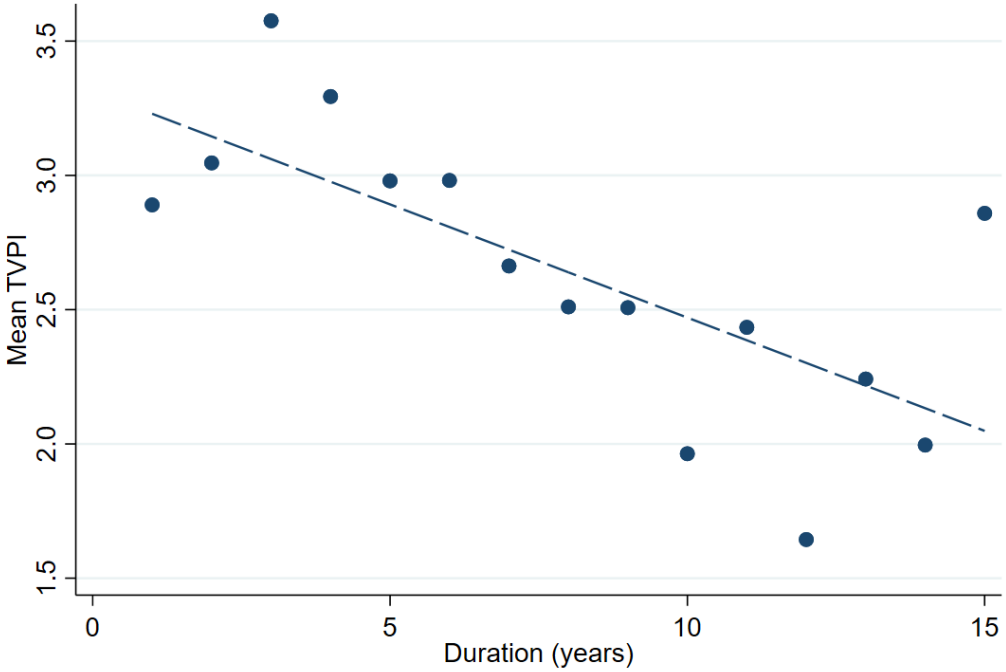
Panel A: Buyout Gross TVPI



Panel B: Buyout Holding Periods



Panel C: Buyout Mean Gross TVPI by Holding Period



Panel D: Buyout Gross TVPI Variance by Holding Period

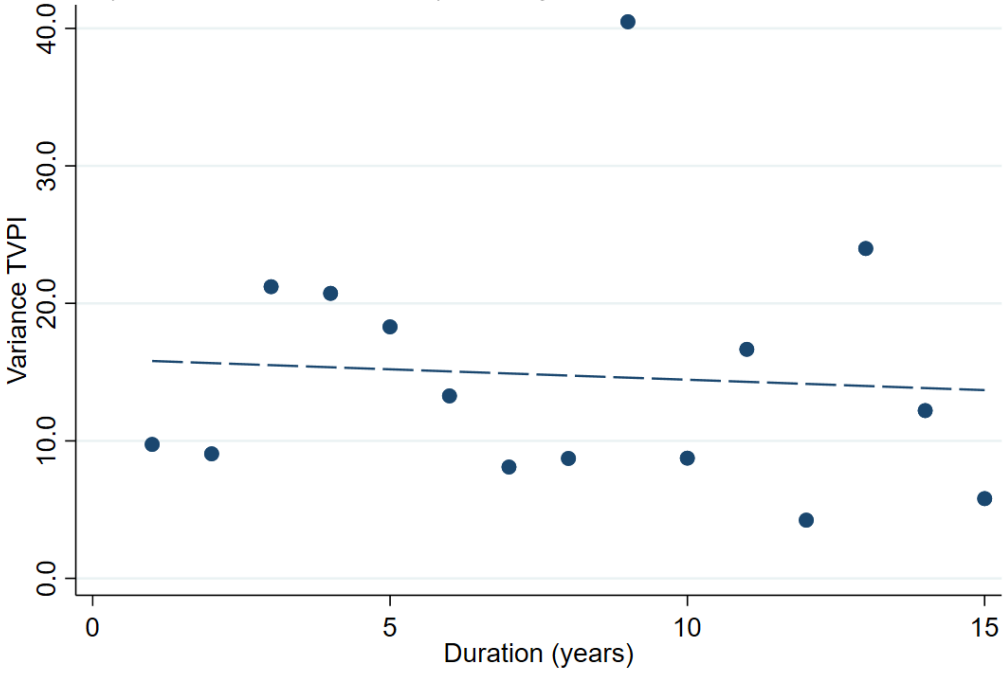
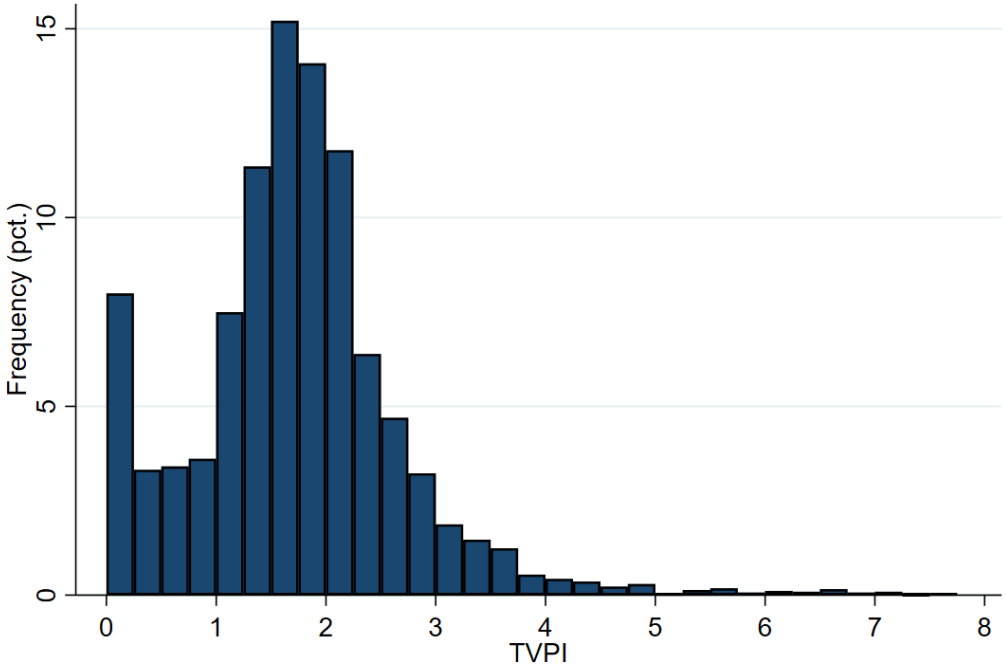
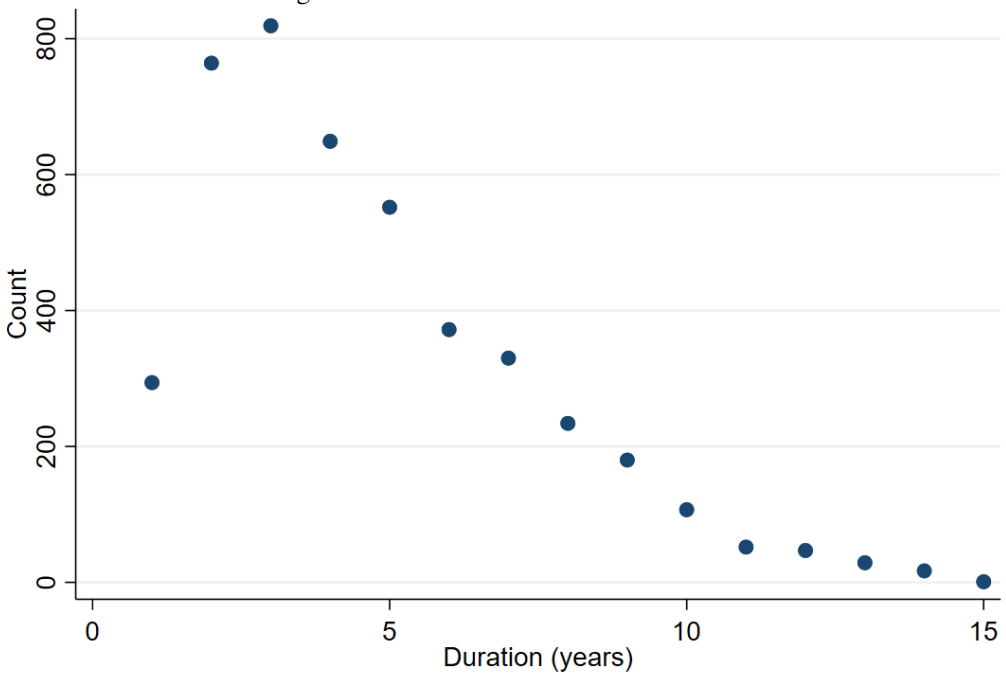


Figure 3: Holdings by Real Estate Funds

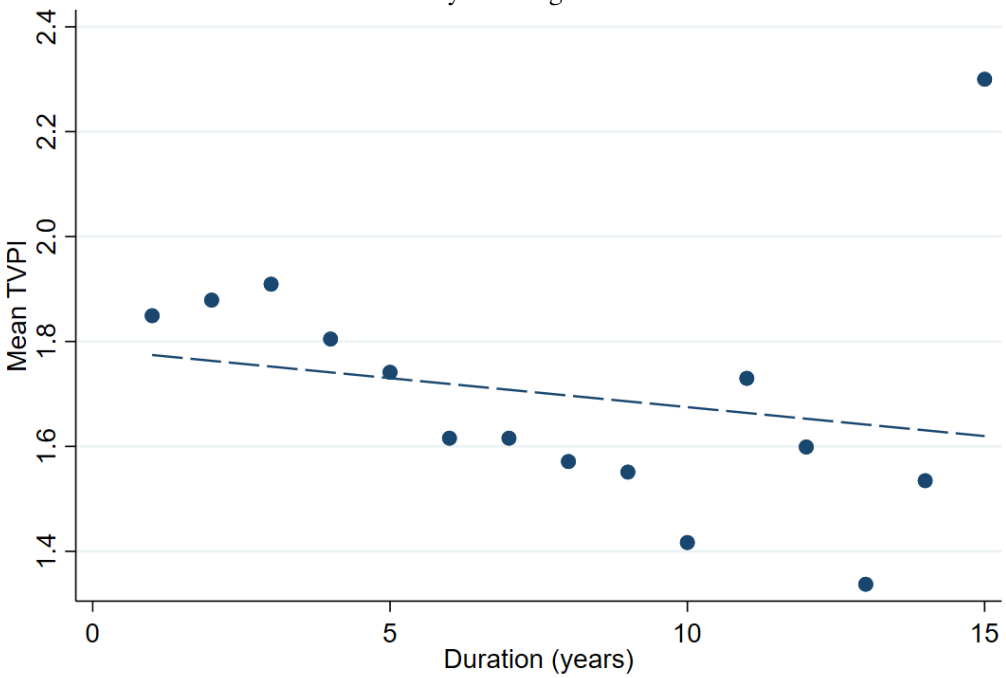
Panel A: Real Estate Gross TVPI



Panel B: Real Estate Holding Period



Panel C: Real Estate Mean Gross TVPI by Holding Period



Panel D: Real Estate Gross TVPI Variance by Holding Period

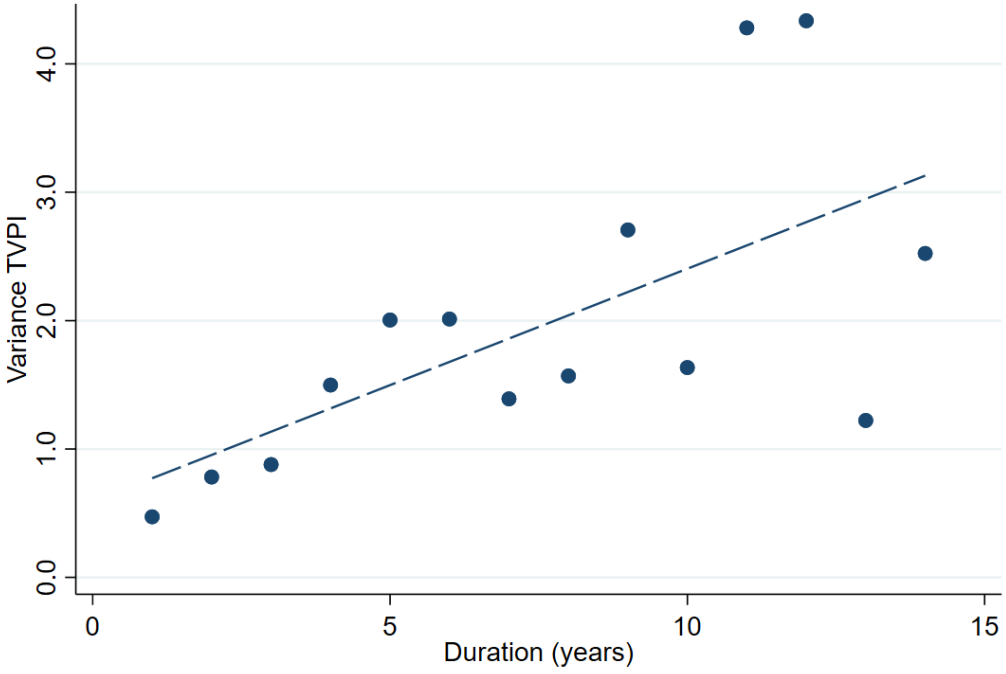
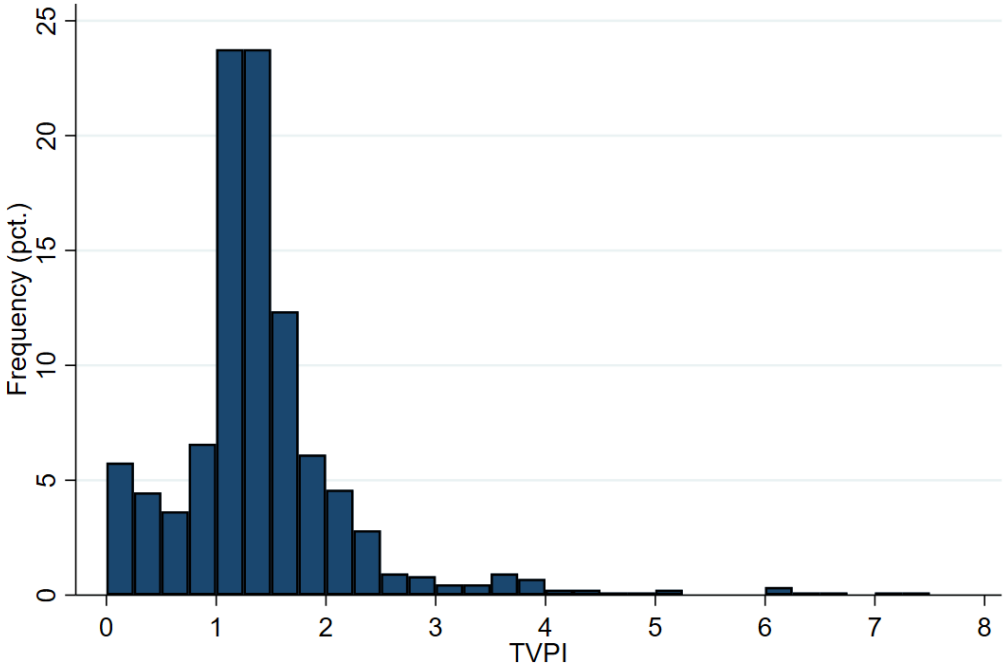
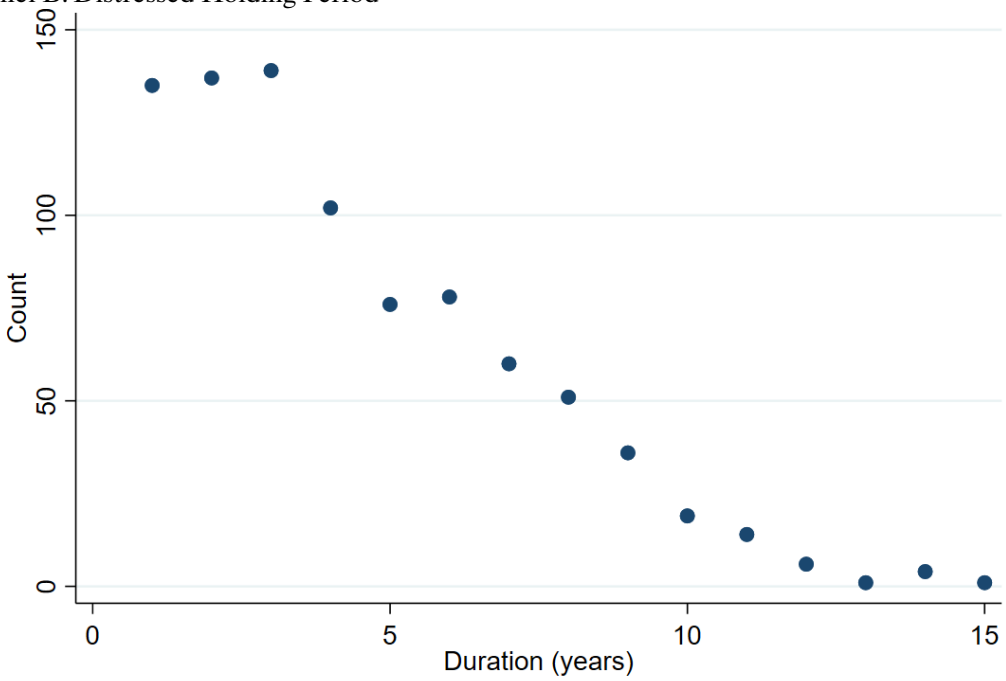


Figure 4: Holdings by Distressed Funds

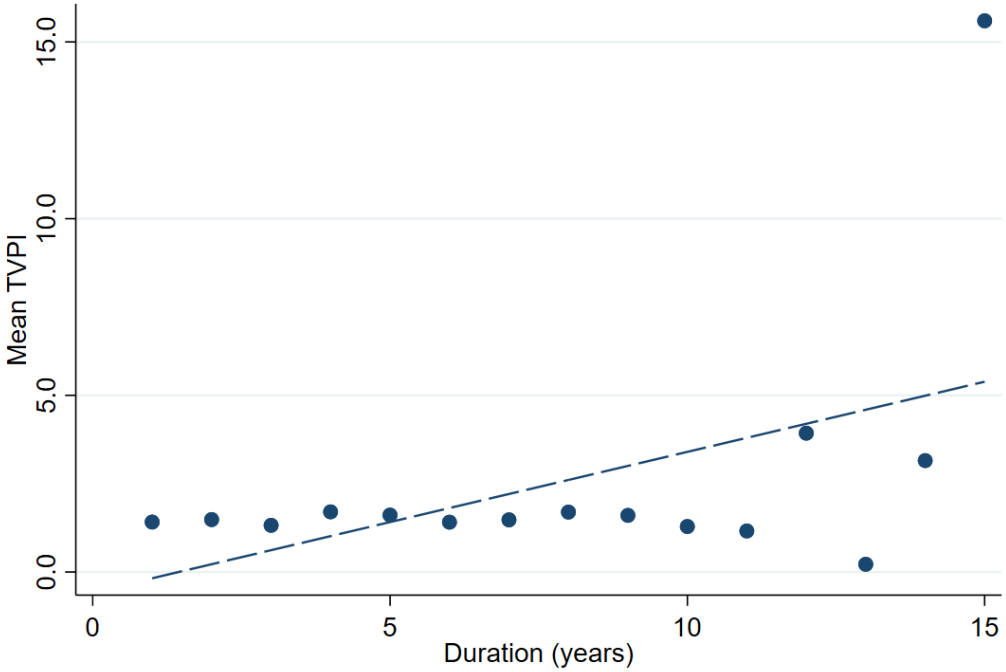
Panel A: Distressed Gross TVPI



Panel B: Distressed Holding Period



Panel C: Distressed Mean Gross TVPI by Holding Period



Panel D: Distressed Gross TVPI Variance by Holding Period

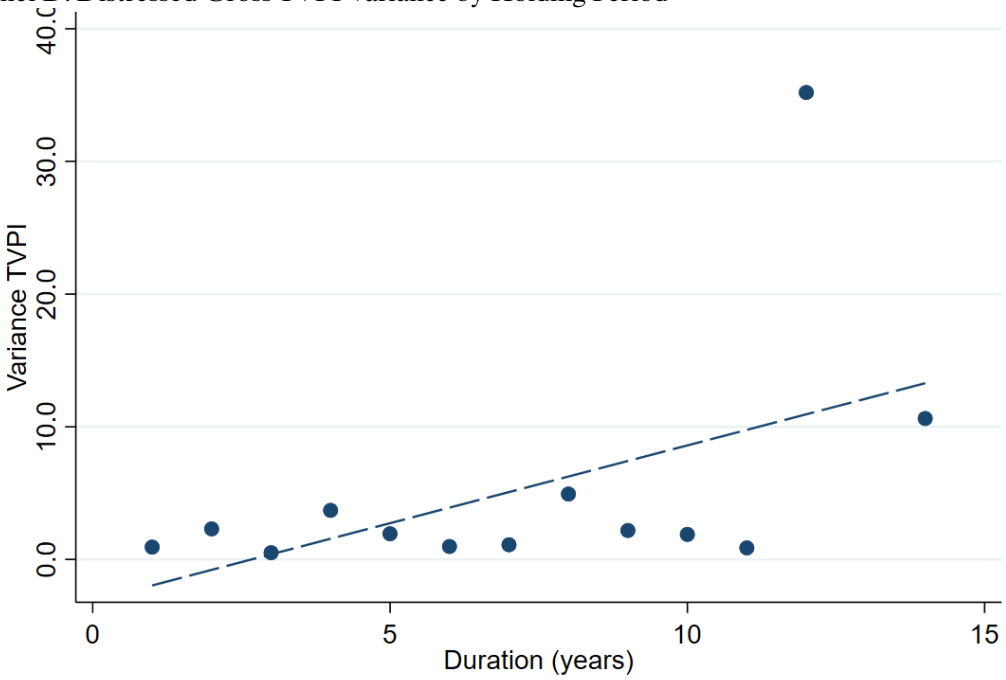
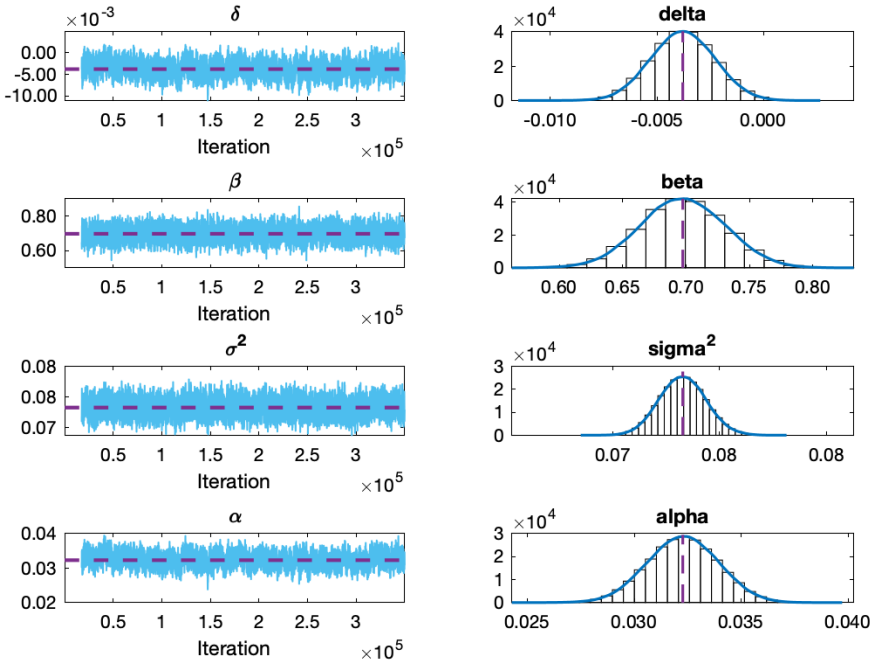
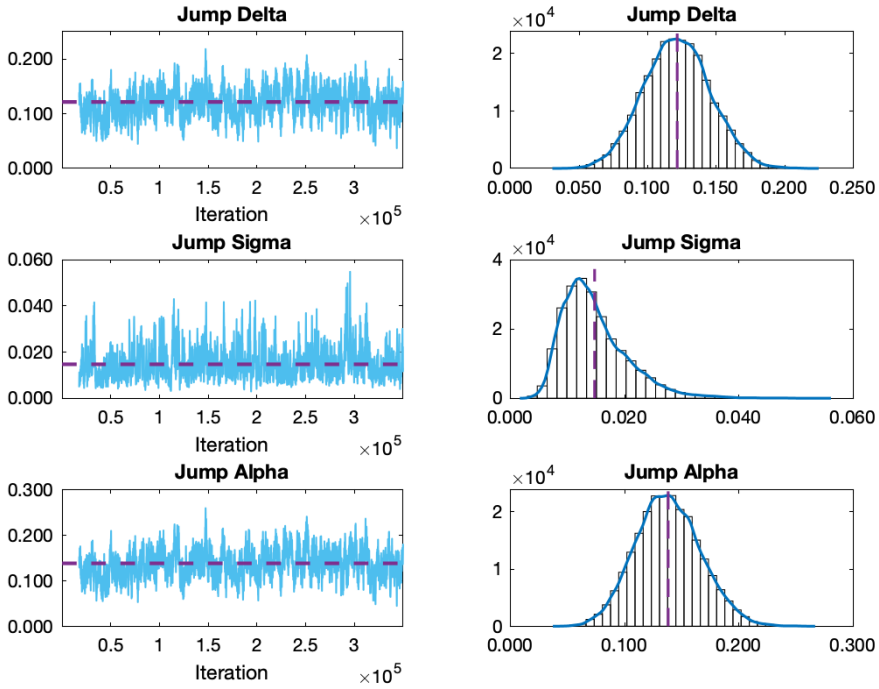


Figure 5: Estimated Parameters for Buyout Sample Posterior draws from specification (1) for buyout holdings (Table 5) for the 350,000 Gibbs iterations, of which the first third is discarded as burn-in. Left panels plot draws by iteration, and right panels show histograms with an overlaid kernel density. Dashed lines show posterior means. Panel A shows parameters for the valuation process, Panel B for the one-time jump specification, and Panel C for the selection process (time fixed effects are omitted).

Panel A: Estimated Parameters for Valuation Process



Panel B: Estimated Parameters for One-time Jump



Panel C: Estimated Parameters for Selection Process

