

Screening and Optimal Allocation of Leveraged Buyout Targets: Point-in-Time Backtesting and Chance-Constrained Mixed-Integer Programming

Prishha Ramesh Babu

Department of Computer Science and Engineering (Big Data Analytics)
SRM Institute of Science and Technology, Ramapuram Campus
Chennai, India

Abstract - Selecting acquisition targets under a capital budget is a discrete optimization problem whose inputs, company-level attractiveness scores, are themselves uncertain. This paper builds an end-to-end decision pipeline for leveraged buyout (LBO) target selection. It combines a sector-relative composite screening score built from five standard LBO ratios; a point-in-time backtest against 22 real public-to-private transactions from 2022 to 2024, using financial statements filed with the U.S. Securities and Exchange Commission (SEC) strictly before each announcement; a budget- and diversification-constrained mixed-integer program with linear programming (LP) relaxation shadow prices; and a chance-constrained extension solved by sample average approximation. Only 4.5 percent of realized targets scored in the top quintile of the universe ahead of their deals (95 percent Wilson interval 0.8 to 21.8 percent, 22 targets). Healthcare targets scored markedly higher (median percentile 76.8, three targets) than technology targets (27.6, fourteen targets), but the technology versus non-technology difference is not statistically significant under a Mann-Whitney test. Supervised recalibration of the scoring weights raised the in-sample area under the receiver operating characteristic curve to 0.739 while lowering the top-quintile hit rate from 4.5 percent to zero, a cautionary result on metric misalignment in small-sample calibration. All results are reproducible from public data sources.

Keywords - mixed-integer programming, chance-constrained optimization, sensitivity analysis, leveraged buyouts, point-in-time backtesting, SHAP

I. INTRODUCTION

In practice, screening for leveraged buyouts is a heuristic exercise: investment teams rank candidate companies based on a few financial criteria (profitability, leverage capacity, cash conversion and earnings stability) to prioritize diligence effort across a large universe. There is little publicly available, reproducible evidence on whether such heuristics are predictive of actual transactions, and even less work connecting target screening to the ultimate decision faced by a sponsor: given a fixed capital budget and a mandate of diversification, which subset of attractive targets should be pursued, and how much value is lost to constraints versus the uncertainty of the screen itself? Prior work surveys the economics of buyouts [1], links buyout activity to free cash flow and financial distress costs [2], and models takeover targets from firm financial characteristics [3].

This paper makes three contributions. First, it builds a fully specified, sector-relative screening score from public financial data and validates it against a point-in-time sample of real take-private transactions, avoiding look-ahead bias by sourcing each target's financials from SEC EDGAR filings dated before its announcement. Second, it formulates the resulting allocation decision as a mixed-integer program with a sector diversification constraint, recovers constraint shadow prices from the LP relaxation, and extends the model to a chance-constrained formulation under company-level score uncertainty using sample average approximation (SAA) [4], [5]. Third, it evaluates a supervised alternative to the hand-set weights (logistic regression) and shows it improves a standard classification metric while underperforming on the metric that matters for screening.

Uncertainty is reported explicitly: a Wilson score interval [6], [7] around the headline hit rate and a nonparametric test on the sector split. The sector pattern is directionally consistent with industry narratives about the 2022 to 2024 take-private wave but is not statistically confirmable at this sample size, a result as informative as a positive finding would have been.

II. DATA

Live screening universe. A universe of approximately 120 U.S. mid-cap public companies spanning industrials, consumer, healthcare, technology, materials, energy, and non-bank financial services was obtained through the Yahoo Finance API (income statement, cash flow statement, and balance sheet, up to five fiscal years per company). Mega-capitalization firms, banks, and regulated utilities are excluded as atypical LBO candidates given their regulatory and capital-structure constraints.

Historical target sample. Twenty-two real public-to-private transactions announced between January 2022 and December 2024 were hand-curated (e.g., Citrix Systems, CDK Global, Anaplan, Zendesk, Univar Solutions, Syneos Health, R1 RCM, Smartsheet) and identified by SEC Central Index Key. Because these companies are delisted, their pre-deal statements were retrieved from the SEC EDGAR XBRL company-facts API [8], restricting extraction to amounts reported on annual filings (Form 10-K or 10-K/A) with a reporting period of approximately twelve months (350 to 380 days), yielding up to five fiscal years of pre-deal annual financial statements per target, each with a period end-date strictly preceding that target's deal-announcement date. This is a deliberate design choice: any evaluation that uses financials

reported after a deal, or restated with hindsight, systematically overstates predictive power.

A separate, finer-grained reconstruction is used elsewhere in the companion dashboard to render a quarterly score-replay timeline for each historical target: XBRL facts are grouped by shared period-start date, consecutive cumulative end-dates are differenced to recover discrete fiscal quarters, and trailing-twelve-month aggregates are built from the four most recent quarters at each point in time. It is not the extraction method behind the annual, point-in-time backtest reported in Section IV, and is noted here only to scope what public SEC XBRL data supports beyond the results in this paper.

III. METHODOLOGY

A. Screening ratios and composite score

Five ratios standard to LBO underwriting are computed per company: earnings before interest, taxes, depreciation, and amortization (EBITDA) margin; EBITDA growth stability (the inverse of the standard deviation of year-over-year EBITDA growth); leverage headroom against a 6.0x Debt/EBITDA ceiling; free cash flow (FCF) to EBITDA conversion; and revenue volatility. Each ratio is oriented so higher is more favorable, missing values are imputed with the sector median, and each company receives a sector-relative z-score and percentile rank. Sector-relative normalization is necessary because raw ratio levels are not comparable across industries (for example, software versus consumer staples margins).

The composite LBO Attractiveness Score is a weighted sum of the five sector-percentile ranks:

$$\text{Score} = \sum_i w_i \cdot \text{percentile}_i \quad (1)$$

with hand-set default weights reflecting conventional screening priorities (FCF conversion weighted 0.28 and leverage headroom 0.26, the two highest of the five). Because the model is exactly linear in its percentile inputs, its Shapley attributions [9], [10] have the closed form $\text{SHAP}_i = w_i \cdot (\text{percentile}_i - \text{mean percentile}_i)$, which was verified to reproduce each company's score to floating-point precision, giving an exact decomposition of any score into its contributing ratios.

B. Point-in-time backtest design

Each historical target's pre-deal ratios are pooled into the live universe's ratio table and the identical sector-relative normalization and scoring procedure is re-run, so every target is evaluated against real sector peer groups under consistent methodology. A hit is a target scoring in the top quintile (80th percentile or above) of the pooled universe. Uncertainty in the hit rate is reported via a Wilson score interval [6], [7], which is more reliable than a normal approximation at small n , and the technology versus non-technology split is tested with a two-sided Mann-Whitney U test [11].

Limitation. The live comparison universe reflects the current set of public companies, not the investable universe on each deal's announcement date. Reconstructing a fully point-in-time universe (re-pulling pre-deal financials for all ~120 companies at 22 historical dates) is left to future work.

C. Capital allocation as a mixed-integer program

Acquisition cost for company i is modeled as market capitalization plus a flat 30% takeout premium. The sponsor's problem is a binary knapsack-style program solved with the CBC solver via PuLP:

$$\text{maximize } \sum_i \text{score}_i \cdot x_i \quad (2)$$

$$\text{subject to } \sum_i \text{cost}_i \cdot x_i \leq \text{Budget} \quad (3)$$

$$\sum (i \in \text{sector } s) x_i \leq \text{MaxPerSector}, \forall s; x_i \in \{0,1\} \quad (4)$$

Integer programs do not admit exact dual values, so approximate constraint shadow prices are recovered from the LP relaxation, the standard approach for sensitivity analysis in discrete optimization, and are reported as marginal values rather than exact guarantees.

D. Chance-constrained allocation

Each company's realized value is modeled as $\text{score}_i + \varepsilon_i$ with $\varepsilon_i \sim N(0, \sigma_i)$, where σ_i is calibrated to that company's historical revenue volatility with a floor for irreducible uncertainty. A constraint of the form $P(\text{portfolio value} \geq V_{\min}) \geq \beta$ is not linear under a continuous distribution, so Sample Average Approximation [4], [5] is used: $S = 300$ scenarios of noisy company values are drawn and each scenario receives a binary shortfall indicator that must activate if its portfolio value falls below $V_{\min}$. Capping the fraction of activated indicators at $(1 - \beta)$ yields a linear MILP that converges to the true chance constraint as S grows.

E. Supervised weight calibration

As an alternative to hand-set weights, a logistic regression is fit with the 22 realized targets as positives and the remainder of the universe as negatives, yielding data-driven weights over the same five ratios. Given the small positive class, this is treated as an illustration of the calibration approach, evaluated on in-sample area under the receiver operating characteristic curve (AUC) and on the same top-quintile hit-rate metric used for the hand-set weights.

IV. RESULTS

Backtest hit rate. 4.5% of the 22 realized targets scored in the top quintile ahead of their announcements (median percentile 30.8; 95% Wilson interval 0.8% to 21.8%). The wide interval reflects the small sample and tempers any precise reading of the point estimate.

Sector pattern. Healthcare targets had a median percentile of 76.8 ($n = 3$), Basic Materials 58.7 ($n = 1$), and Technology 27.6 ($n = 14$). A Mann-Whitney U test of Technology ($n = 14$) against all other sectors pooled ($n = 8$) gives $U = 42.0$, $p = 0.365$, not significant at conventional levels. The pattern is consistent with commentary that 2022 to 2024 software take-privates were driven by growth deceleration and multiple compression rather than the "stable, cash-generative, underlevered" template this screen targets, but it is reported as a suggestive observation, not a confirmed effect.

Optimal portfolio. At a \$50 billion budget with at most three targets per sector, the MILP selects 17 companies, deploying \$49.95 billion (99.9% utilization) for a total portfolio value of 991.8 score-points.

Sensitivity. The LP-relaxation shadow price indicates an additional \$1 million of budget is worth approximately 0.009 score-points at the margin. The sector constraint costs 5.4 points (0.5%) relative to the sector-unconstrained optimum, with the Consumer Cyclical cap binding most tightly (22 eligible candidates against a cap of 3).

Robust allocation. At a \$12 billion budget requiring a 90% probability of clearing 95% of the deterministic-optimal value, the deterministic-optimal portfolio already satisfies the requirement, clearing the threshold in 94.7% of the 300 simulated scenarios with no reallocation. Sweeping budgets and thresholds shows this is not isolated: in a discrete allocation setting the "safe" and unconstrained-optimal portfolios can coincide, or the constraint becomes infeasible outright, with no smooth trade-off in between.

Weight calibration. The calibrated weights reach a 0.739 in-sample AUC but a 0% top-quintile hit rate versus 4.5% for the hand-set weights. A classifier optimized for separability does not automatically optimize a specific downstream ranking objective, a cautionary note for calibrating screening models on small positive-class samples.

V. LIMITATIONS

(i) The backtest compares historical targets against today's live universe rather than the investable universe at each deal date; (ii) the calibration exercise uses a small ($n = 22$), technology-concentrated positive sample and demonstrates a method rather than a validated model; (iii) reported shadow prices are LP-relaxation approximations, an inexact proxy for integer-program duality; (iv) acquisition cost is a flat 30% premium over market capitalization, whereas realized premiums vary with competition and structure; and (v) SHAP attributions, exact for the linear model, use today's universe as the background population.

VI. CONCLUSION

A transparent, sector-relative LBO screening score built from standard underwriting ratios shows limited, sector-concentrated ability to anticipate realized 2022 to 2024 take-privates ahead of time, and the technology versus non-technology gap does not reach significance at this sample size. The allocation layer built on the score, mixed-integer selection under budget and diversification constraints extended to a chance-constrained formulation, shows that diversification and robustness constraints can be inexpensive (here under 1% of expected value) and need not always trade away performance. The principal contribution is methodological: a reproducible, point-in-time-correct backtesting design for cross-sectional screening models, paired with a tractable chance-constrained formulation for discrete allocation decisions, both of which generalize to any setting where a decision-maker must select among discrete, uncertain, budget-constrained alternatives.

DECLARATION OF AI ASSISTANCE

During the preparation of this work, the author used Claude (Anthropic) to assist with software implementation of the analysis pipeline and with drafting and editing the manuscript text. The study design, data sources, and analysis were directed

by the author. The author reviewed, edited, and verified all content and takes full responsibility for it.

DECLARATIONS

The author declares no competing interests. This research received no external funding. No human or animal participants were involved; all data are publicly available (SEC EDGAR and Yahoo Finance).

REFERENCES

- [1] S. N. Kaplan and P. Strömberg, "Leveraged buyouts and private equity," *J. Econ. Perspect.*, vol. 23, no. 1, pp. 121-146, 2009.
- [2] T. C. Opler and S. Titman, "The determinants of leveraged buyout activity: free cash flow vs. financial distress costs," *J. Finance*, vol. 48, no. 5, pp. 1985-1999, 1993.
- [3] K. G. Palepu, "Predicting takeover targets: a methodological and empirical analysis," *J. Account. Econ.*, vol. 8, no. 1, pp. 3-35, 1986.
- [4] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, 2nd ed. New York: Springer, 2011.
- [5] J. Luedtke and S. Ahmed, "A sample approximation approach for optimization with probabilistic constraints," *SIAM J. Optim.*, vol. 19, no. 2, pp. 674-699, 2008.
- [6] E. B. Wilson, "Probable inference, the law of succession, and statistical inference," *J. Amer. Statist. Assoc.*, vol. 22, no. 158, pp. 209-212, 1927.
- [7] R. G. Newcombe, "Two-sided confidence intervals for the single proportion: comparison of seven methods," *Statist. Med.*, vol. 17, no. 8, pp. 857-872, 1998.
- [8] U.S. Securities and Exchange Commission, "EDGAR application programming interfaces (APIs)," <https://www.sec.gov/edgar/sec-api-documentation> (accessed Oct. 20, 2026).
- [9] L. S. Shapley, "A value for n -person games," in *Contributions to the Theory of Games*, vol. 2, H. W. Kuhn and A. W. Tucker, Eds. Princeton, NJ: Princeton Univ. Press, 1953, pp. 307-317.
- [10] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [11] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *Ann. Math. Statist.*, vol. 18, pp. 50-60, 1947.