SERIES:
Research Insights for Investment Professionals

ISSUE: NO. 1 | DATE: JUNE, 2026

TRANSFER LEARNING FOR EXPECTED RETURNS

A practitioner-focused synthesis of original research

EXECUTIVE SUMMARY

Estimating expected stock returns is central to portfolio construction, manager evaluation, cost-of-capital analysis, and asset pricing. The practical difficulty is that the best expected-return proxies are only reliable for a limited set of firms. Option-implied measures work mainly for large stocks with liquid options markets. Analyst-based implied cost-of-capital measures work mainly where analyst coverage is deep. Outside those data-rich segments, the proxies either disappear or lose much of their link to future returns.

This study develops a transfer-learning framework that extends high-quality expected-return information from firms where proxies work well to the broader U.S. equity universe. The method first learns how firm characteristics map into high-signal expected-return proxies in a reliable "teacher" domain. It then applies a disciplined correction, using realized returns only as a calibration device, so that the predictions remain useful outside that original domain. The resulting measures cover U.S. equities from 1957 to 2023 and deliver near-unbiased forecasts relative to existing alternatives.

The practical value is clear. Investment professionals often need expected-return estimates precisely where traditional inputs are weakest: smaller, less liquid, and less-followed firms. This framework is designed for that setting. It expands market coverage, preserves the economic content of stronger proxy measures, and reduces the risk of drawing conclusions from a narrow and unrepresentative subset of the market.

Key Takeaways for Investment Professionals

- Strong expected-return proxies are most useful in large, liquid, well-covered firms.
- Extrapolating those proxies mechanically to the rest of the market can be misleading.
- Transfer learning can extend expected-return estimates across the full cross-section while keeping the original proxy signal intact.
- The framework delivers economically meaningful signals, including a 0.93% monthly FF6 alpha in 2012-2023 for the option-implied learned expected-return strategy.

CFA Society Netherlands

ABOUT THE VBA RESERVES The research in this series is funded through the VBA reserves of CFA Society Netherlands. These reserves were created following the merger between VBA Beleggingsprofessionals and CFA Society Netherlands and are dedicated to supporting initiatives that promote knowledge and professional practice related to investment analysis, portfolio management, financial issues, and related disciplines. Through this funding, the Society aims to contribute to practitioner-oriented research and knowledge sharing for the development of the investment profession in the Netherlands.

The analysis also shows that sample coverage can materially affect empirical conclusions. A prominent example is the idiosyncratic volatility premium. In analyst-covered samples, the relation between idiosyncratic volatility and expected returns appears positive. Once learned expected returns are extended to the full cross-section, that relation falls sharply and largely disappears. For practitioners, that is an important warning: some established findings may reflect data coverage as much as underlying economics.

WHY THIS RESEARCH MATTERS NOW

Expected-return estimation is becoming more important, not less. Asset allocators need better estimates to support active risk budgets, manager selection, and strategic tilts. Corporate finance users need them for discount rates and capital budgeting. At the same time, markets remain segmented between firms with rich data coverage and firms with sparse coverage. That makes it risky to treat evidence from large, liquid, heavily followed firms as if it applies uniformly to the broader market.

The study addresses a structural problem in current investment practice: measurement quality varies sharply across firms. Option-implied expected returns are strongest where options are liquid. Analyst-based measures are strong where analyst coverage is dense. But many practical investment decisions concern firms outside these segments. Small caps, less liquid names, and firms with thin analyst coverage are often precisely where valuation disagreement, financing frictions, and return opportunities are most pronounced.

The rise of machine learning makes the study especially relevant. Much recent work trains models directly on realized returns. That can help prediction, but such signals are often harder to interpret as expected returns. This framework takes a different route. It starts with economically meaningful expected-return proxies and uses machine learning to generalize them, rather than replacing them with a pure return-forecasting black box. For practitioners, that distinction matters because expected-return estimates are used not only to rank stocks, but also to justify portfolio tilts, estimate hurdle rates, and communicate investment views.

If coverage issues are ignored, investors risk two errors. First, they may underuse information from firms where direct proxy measurement is unavailable. Second, they may overgeneralize relationships estimated in selected samples. Both errors can distort portfolio construction and empirical interpretation.

Relevance in Today's Context

- Coverage gaps remain a core challenge in expected-return measurement.
- Investors increasingly need signals that work beyond large-cap, high-liquidity universes.
- Machine learning is most useful here when it extends interpretable signals rather than replacing them.
- Findings from narrow samples should be tested before they are implemented broadly.

RESEARCH QUESTION AND APPROACH

The core question is straightforward: can information from reliable expected-return proxies be extended to firms and periods where those proxies are unavailable or weak? The study answers this using a teacher-student transfer-learning design.

The teacher domain is the part of the market where a given proxy is reliable. For option-implied measures, that is mainly the S&P 500, where option markets are liquid and coverage reaches roughly 99%. In that domain, the proxy has a regression slope of 1.11 against subsequent realized returns, close to the ideal benchmark of one. Outside the S&P 500, the same slope falls to 0.12 and coverage drops below 4%. Analyst-based implied cost-of-capital proxies show a similar pattern: reasonable performance in well-covered firms, much weaker performance elsewhere.

The framework proceeds in two stages. First, a Gradient Boosted Trees model learns the mapping from 153 firm characteristics to the expected-return proxy within the teacher domain. This stage captures the proxy's economic structure from observable characteristics. Second, the model applies a linear transfer correction using realized returns to restore the link between predicted expected returns and subsequent returns outside the teacher domain. Realized returns are not used as direct training labels for the baseline model; they are used only to calibrate the correction. That design prevents the approach from collapsing into a generic return-prediction exercise.

The sample covers U.S. equities from 1957 to 2023, with strict out-of-sample testing in 2012-2023 for the modern evaluation. The method is also applied historically to periods before option-implied data were broadly available, allowing the authors to recover long-run expected-return estimates where direct proxy data do not exist.

The framework should be interpreted with appropriate caution. It does not claim to observe expected returns directly. It still evaluates performance using realized returns, which are noisy. Nor does it imply that any one proxy is universally correct. Instead, it provides a structured way to preserve useful proxy information while adjusting for the fact that proxy quality changes across domains.

What This Research Does (and Does Not) Show

- It shows that strong expected-return proxies can be extended beyond their original coverage domains.
- It shows that characteristics-only imputation is not enough; domain correction is necessary.
- It does not claim that realized returns are a perfect measure of expected returns.

APPLICATIONS FOR INVESTMENT PROFESSIONALS

The first application is portfolio construction. Sorting stocks on learned expected returns produces economically meaningful long-short returns. In the modern out-of-sample period from 2012 to 2023, the option-implied learned expected-return signal generates a 0.93% monthly alpha relative to the Fama-French six-factor model, with a t-statistic of 3.66. Even after adding a neural-network factor designed to absorb generic machine-learning return predictability, the alpha remains positive at 0.49% per month. This suggests that the transfer-learned signal contains information beyond standard characteristics-based return prediction.

Historical application is also important. In 1957-1995, when option-implied expected-return measures were not broadly available, the learned option-based strategy still delivers a 0.67% monthly FF6 alpha, with a t-statistic of 4.36. That matters for investors and researchers who want longer time series of expected-return estimates for backtesting, strategic research, or economic interpretation.

A key practical lesson is that naive extrapolation fails. When the model simply predicts the original proxy from firm characteristics without the transfer correction, the resulting portfolios can produce strongly negative alphas. In other words, broadening coverage is not enough. The method must also account for the fact that the proxy-behavior relationship changes outside the original sample.

For investment teams, the framework can be used in at least two ways. First, it can support cross-sectional stock ranking where direct expected-return proxies are missing. Second, it can serve as a robustness tool: if a trading strategy is identified in a particular sample and relies on "special data", our framework offers a method to test the relationship in a large cross-section of firms including firms without "special data".

Practical Use Cases

- Extend expected-return estimates to small-cap and less-followed stocks where direct proxy data are weak or unavailable.
- Build broader stock-ranking signals without relying solely on realized-return machine learning.
- Reassess trading signals and factor tilts for potential sample-selection bias.

CONCLUSIONS AND NEXT STEPS

The main conclusion is that expected-return measurement has a domain problem. Existing proxies work best where data are richest, but many practical decisions concern firms outside those environments. This study provides a disciplined way to transfer information from high-quality proxy domains to the broader market while preserving economic interpretability.

For practitioners, the contribution is less about adding another opaque alpha signal and more about improving measurement. Better expected-return estimates can improve portfolio optimization, and reduce the risk of naïve extrapolation from large-caps tot the full cross-section.

There are also clear next steps. A rolling implementation may improve real-time use. The same framework could be extended to international equities or to other finance settings with the same structure: a high-quality label for a selected subset and a noisy outcome for the full population. Possible examples include analyst forecasts, credit-risk measures, and ESG signals.

The appropriate professional stance is constructive but cautious. The framework improves coverage and appears economically meaningful, but it does not eliminate estimation risk. Expected-return estimates should still be treated as inputs to judgment rather than mechanical outputs. Even so, the study makes a strong case that the next generation of expected-return modeling should focus not only on prediction accuracy, but also on where the labels are reliable and how to extend them responsibly.

Points for Professional Discussion

- Which investment processes would benefit most from broader expected-return coverage?
- How should firms govern machine-learning tools when the objective is measurement rather than pure prediction?
- Which empirical signals in current practice may partly reflect data-coverage bias?
- Can this framework be extended reliably beyond U.S. equities or beyond expected-return estimation?

About the Original Research

This summary is based on the following original research paper:

- **Title:** Transfer Learning for Expected Returns
- **Authors:** Julio Crego, Jens Soerlie Kvaerner, and Marc Stam
- **Affiliations:** Nova School of Business and Economics; BI Norwegian Business School; Tilburg University
- **Source manuscript:** https://papers.ssrn.com/sol3/Delivery.cfm/4345646.pdf?abstractid=4345646&mirid=1, 2 February 2023. Last revised 13 April 2026
- **Status:** working paper currently under review
- **Note:** this practitioner summary is based on a work in progress. The final peer-reviewed version may change following reviewer feedback and the publication process.