
RETRIEVAL-AUGMENTED GENERATION VS. DETERMINISTIC TAX COMPUTATION IN MULTI-AGENT FINANCIAL ADVISORY: A 2×2 FACTORIAL EXPERIMENT

Aryan Brar
Royal Bank of Canada
abrar82@uwo.ca

Justin Du
Royal Bank of Canada
justin.du@mail.utoronto.ca

Avery Lor
Royal Bank of Canada
avery.lor@mail.utoronto.ca

Kylie Seto
Royal Bank of Canada
ksa210@sfu.ca

Eric Taylor
RBC Borealis
eric.j.taylor@borealisai.com

ABSTRACT

Tax-loss harvesting demonstrates consistent benefits to long-term portfolio growth; yet implementing it efficiently often involves complex considerations that are specific to the holdings within that portfolio and the individual who owns it. We introduce a custom capital gains calculation engine and a RAG-retrieved vector store of market advisory reports to provide context for a multi-agent trade recommendation system. We investigate the effects of each context provider on the quality of recommendations, measured by relative capital gains incurred during portfolio liquidation. A 2×2 repeated-measures ANOVA revealed a significant main effect of the tax optimization engine ($F(1, 29) = 9.17, p = .005, \eta_p^2 = .240$): enabling the engine *reduced* tax savings by approximately 55 percentage points relative to the no-engine conditions. The RAG main effect was not significant ($p = .841$), nor was the interaction ($p = .553$). The RAG-only condition achieved the highest descriptive mean tax savings (47.7%), and the baseline condition performed second-best (30.6%), suggesting that the pre-trained language model’s internalized financial knowledge may be sufficient for competent tax-loss harvesting recommendations without explicit tooling. These results indicate that augmenting LLM agents with domain-specific computation engines does not guarantee improved performance and may introduce conflicting optimization signals.

1 Introduction

1.1 AI in Wealth Management

The intersection of artificial intelligence and wealth management has emerged as a transformative domain in financial services. Machine learning and deep learning techniques have progressively enhanced portfolio optimization, risk assessment, and investment decision-making at both institutional and retail scales. Recent advances have moved beyond simplistic rule-based systems toward sophisticated neural architectures capable of learning complex market dynamics from high-dimensional data. Robo-advisors and AI-driven portfolio platforms have democratized access to quantitative investment strategies previously available only to institutional investors with substantial capital and technical infrastructure. However, the deployment of these systems raises critical questions about their effectiveness in real-world conditions, particularly when operating under regulatory constraints and competing objectives such as after-tax performance, transaction cost minimization, and investor alignment.

1.2 Tax-Loss Harvesting and Tax Alpha

Berkin and Ye [2003] describe Tax-loss harvesting (TLH) as the practice of selling shares below the original cost to generate tax credits. Tax credits can be used to subtract from a persons’ capital gains in the current year or use to defer

This work was completed at Royal Bank of Canada as part of the RBC Amplify program.

gains in future years. TLH represents one of the most direct mechanisms for improving after-tax portfolio returns. Tax alpha as described by Berkin, A. L and Ye, J is the “tax consequences of active management”. It represents the money added to or subtracted from a portfolio due to how capital gains, losses, and other tax events are managed and realized.

1.3 Multi-Agent Architectures in AI Systems

Mixture-of-Experts (MoE) and multi-agent architectures have emerged as powerful paradigms for handling heterogeneous, multi-objective problems in machine learning. Rather than forcing a single monolithic model to optimize conflicting objectives, MoE systems decompose the problem space into specialized expert agents, each trained on a distinct objective or task domain, with a learned router or gating mechanism determining expert allocation at inference time. Pishehvar [2026] demonstrated this approach in portfolio management by constructing four specialized experts (momentum, growth, defensive, tax-aware), each optimizing distinct investment mandates, with an intent router blending expert outputs based on active objective and market regime. This architecture naturally accommodates the heterogeneity inherent in personalized portfolio management: different investors have different time horizons, tax brackets, liquidity needs, and behavioral preferences. A reality that monolithic systems struggle to capture. Multi-agent systems also offer interpretability advantages; routing decisions can be examined to understand how objectives are being weighted, and expert specialization can be validated independently. However, the literature on multi-agent portfolio systems remains sparse, and critical questions remain about how to effectively coordinate agents when objectives are partially aligned, conflicting, or unknown at design time.

1.4 Retrieval-Augmented Generation and Knowledge Retrieval

Retrieval-Augmented Generation (RAG) has emerged as a powerful pattern for augmenting language models and decision-making systems with grounded, external knowledge. Rather than relying solely on learned parameters, RAG systems retrieve relevant documents, data, or context from a knowledge base at inference time, using that context to condition generation or decision-making. This approach has demonstrated effectiveness in reducing hallucination in large language models, improving factuality, enabling dynamic knowledge updates without retraining, and providing verifiable justification for decisions. In the financial domain, RAG principles could enhance portfolio management by retrieving relevant market news, regulatory guidance, tax documentation, and domain-specific insights to inform allocation decisions. The potential synergy between RAG and multi-agent systems is significant but underexplored: a tax agent could retrieve relevant tax code sections and recent rulings; a momentum agent could retrieve recent earnings announcements and analyst sentiment; a risk agent could retrieve historical volatility regimes and tail-risk scenarios. However, the integration of retrieval-based knowledge with learned policies, particularly in sequential decision-making contexts like portfolio rebalancing, remains largely unexplored in the academic literature.

1.5 Research Gap

Existing literature has addressed AI in portfolio management, tax-loss harvesting optimization, and multi-agent architectures largely in isolation. Few if any studies have examined the synergistic integration of these three components within a unified system. Specifically: (1) most tax-aware portfolio systems operate deterministically or use conventional optimization without deep learning; (2) multi-agent RL systems in finance typically optimize a single objective (e.g., Sharpe ratio) and do not systematically integrate tax reasoning; (3) retrieval-augmented approaches have not been tested in the portfolio management context to determine whether grounding decisions in external knowledge meaningfully improves performance; and (4) no empirical comparison exists between pure learned policies, knowledge-augmented policies, tax-specialized architectures, and their combinations. The literature provides no principled guidance on whether, and under what conditions, combining these approaches yields additive, synergistic, or diminishing returns. This gap is particularly consequential for retail investors and fintech platforms, where after-tax returns directly translate to user value, and where the practical integration of multiple AI techniques must operate within computational and regulatory constraints.

2 Related Works

2.1 Tax-Loss Harvesting: Mechanisms and Effectiveness

Tax-loss harvesting has become an increasingly important strategy for improving after-tax returns in portfolio management. Chaudhuri et al. [2020] provided an empirical evaluation of tax-loss-harvesting alpha using historical US equity data from 1926 to 2018, demonstrating that tax-loss harvesting yields an average of 1.08% annualized alpha before transaction costs, declining to 0.82% when constrained by wash-sale rules. The study revealed that tax alpha varies

substantially across market regimes, performing best during high-volatility periods with low overall returns (such as the Great Depression era, yielding 2.13% annually) and modestly during low-volatility expansion periods (0.51% during 1949–1972).

Israelov and Lu [2022] advanced the theoretical understanding of tax-loss harvesting by characterizing the optimization problem as a classic risk-reward tradeoff. They established that harvesting efficiency depends on the security’s volatility and introduced an efficient frontier between harvesting yield and active risk. Their analysis demonstrated that a -10% harvesting threshold (monthly) or -15% threshold (daily) strikes a reasonable balance, and they proposed a novel "throttled harvesting algorithm" for large positions that restricts maximum active weight deviations while maintaining adequate harvesting opportunities. Notably, they found that daily harvesting provides minimal marginal improvement over monthly harvesting when transaction costs are considered.

2.2 Tax-Aware Factor Investing and Long-Short Strategies

Krasner and Sosner [2023] examined the mechanisms underlying tax-aware long-short factor strategies, resolving an apparent paradox: how strategies can simultaneously achieve cumulative net capital losses exceeding 100% of invested capital while maintaining significant pre-tax alpha. Their key finding was that net capital losses in tax-aware long-short strategies arise primarily from gain deferral rather than increased loss realization, a distinction they argue is critical for investor understanding. The study demonstrated that tax-aware strategies achieve this through liquidating loss positions and creating new positions while minimizing gain realization, particularly by deferring short-term gains on long positions. Importantly, they showed that most portfolio turnover remains directed toward the alpha model, explaining the sustained pre-tax performance.

2.3 AI-Based and Deep Learning Approaches to Portfolio Management

Pishehvar [2026] introduced a three-phase deep reinforcement learning system for personalized, tax-aware portfolio management that addresses key limitations of prior work. Phase 1 employs self-supervised learning with a cross-asset encoder augmented by the Chronos time series foundation model, addressing "ticker lock-in" through a 50-dimensional observable metadata vector that generalizes to any publicly traded asset without retraining. Phase 2 introduces a Mixture-of-Experts (MoE) architecture with four specialized expert heads (momentum, growth, defensive, tax-aware) and a learned intent router that simultaneously serves six distinct investment objectives. A critical contribution is the inter-ticker contrastive loss that resolved representation collapse (mean cosine similarity from 0.96 to 0.24), enabling genuinely differentiated portfolio weights. Phase 3 implements lightweight personalization via a 76-parameter LoRA module that infers investment objectives from revealed trading behavior rather than questionnaires, incorporating natural language goal specification (e.g., "buy a house in 3 years").

2.4 Portfolio Optimization with Explainability and Retail Accessibility

Bachhav et al. [2024] developed an AI-based personalized portfolio allocation engine targeting retail investors, integrating a composite risk score, market regime detection, multi-factor equity scoring, and tax-aware rebalancing with FIFO lot tracking and tax-loss harvesting. Their implementation demonstrated that tax-loss harvesting reduced simulated rebalancing tax liability by 39.8%, with annualized returns of 8.2%–16.1% across risk profiles. The platform also employs Explainable AI modules to convert algorithmic decisions into plain-language narratives, addressing a key transparency gap in robo-advisory literature.

2.5 Foundational Concepts and Modern Adaptations

The present work builds upon seminal contributions to portfolio theory and tax-aware investing. Berkin and Ye [2003] established foundational methods for simulating tax-loss harvesting strategies, while Sialm and Sosner [2018] demonstrated that tax awareness in actively managed factor strategies primarily operates through gain deferral rather than loss maximization. The literature increasingly emphasizes the importance of accounting for regulatory constraints (wash-sale rules), market microstructure (transaction costs, bid-ask spreads), and behavioral finance considerations (disposition effects, investor heterogeneity) when implementing tax-efficient strategies.

Recent advances integrate technological enablement with financial theory. The decline in trading costs and rise of fintech platforms have democratized tax-aware investing, previously accessible only to high-net-worth and institutional investors. Concurrently, deep learning and reinforcement learning methods enable dynamic, personalized optimization that simultaneously balances multiple objectives, alpha generation, tax efficiency, and behavioral alignment, at scale and low cost.

3 Hypotheses

We hypothesize that portfolio optimization performance, measured by after-tax returns, follows a clear hierarchical ordering based on architectural sophistication and information integration. We are going to conduct a 2x2 factorial experiment in which a multi-agent investment advisor system generates tax-optimized trade recommendations for simulated client portfolios under four conditions: (1) a no-augmentation baseline, (2) a custom capital-gains calculation engine only (Tax Engine), (3) a RAG-retrieved vector store of market and tax advisory reports only, and (4) both components active simultaneously. The primary outcome is the percentage of reduction in projected capital-gains tax liability relative to a no-trade baseline. Within this design, we test the following hypotheses.

1. **Baseline Underperformance:** A baseline deep reinforcement learning policy trained on pre-tax returns alone, without explicit tax reasoning or external knowledge, will deliver the poorest after-tax performance relative to all augmented conditions.
2. **Tax-Aware Improvement:** A tax-specialized agent trained with explicit loss-harvesting objectives, position-level tax-lot tracking, and wash-sale constraints will substantially outperform the baseline, capturing the documented tax alpha from loss harvesting while maintaining competitive pre-tax returns.
3. **RAG Knowledge Benefit:** A retrieval-augmented policy that grounds decisions in contextual market knowledge, earnings announcements, volatility regimes, tax documentation, regulatory changes, will outperform the baseline by improving decision quality through better information access, though likely with smaller marginal gains than the tax agent alone.
4. **Synergistic Integration:** The combination of both tax-specialized reasoning and retrieval-augmented knowledge will outperform any single augmentation, yielding a cumulative effect in which tax reasoning provides structural optimization while RAG provides adaptive context sensitivity. We further hypothesize that this synergistic combination will approach or exceed the theoretical tax alpha established in prior literature, demonstrating that learned policies can internalize the insights from decades of tax optimization research.

4 Methods

This study used a 2x2 within-subjects factorial design to evaluate the independent and interactive effects of two system components: a tax optimization engine and a vectorized knowledge retrieval system on the quality of AI-generated tax-optimized trade recommendations. The independent variables were the tax engine (present vs. absent) and knowledge retrieval system (present vs. absent), yielding four conditions: both active, tax engine only, retrieval only, and a no-component baseline. All conditions ran on the same multi-agent AI architecture with only the relevant components toggled on or off via configuration settings, ensuring observed differences are attributable solely to those components. Each portfolio was evaluated under all four conditions, making portfolio the repeated-measures unit. A two-way repeated-measures ANOVA was selected because it estimates both main effects and their interaction within a single omnibus test, accounts for individual portfolio variability in the error term (increasing statistical power), and is more efficient than one-factor-at-a-time comparisons (Montgomery [2017]).

The experimental units were 30 simulated client portfolio scenarios, sized to achieve power of $1 - \beta \geq 0.80$ for a medium effect ($f = 0.25$, $\alpha = .05$; Cohen [1988]). Each synthetic portfolio scenario represented a taxable brokerage account containing between 5 and 20 security positions, with at least one unrealized capital loss to ensure a harvestable opportunity existed. Scenarios varied in total value (ranging from \$25,000 to \$500,000) and in the ratio of unrealized gains to losses, reflecting a realistic cross-section of retail client accounts. Because this was a within-subjects design, each of the 30 portfolios was evaluated under all four conditions, yielding 120 total observations (30 portfolios $\times$ 4 conditions). Portfolio holdings, account structure, and capital gains history were held constant across conditions via deterministic seeding, ensuring that any differences in outcomes are attributable to the experimental manipulation rather than portfolio characteristics.

The main experimental loop executed each of the 30 portfolio scenarios through all four conditions. For each iteration, a LangChain-based multi-agent orchestration system was invoked. An Orchestrator Agent delegated tasks to specialized sub-agents, each with access to domain-specific tools. The agent workflow operated as follows: upon receiving a request containing account context, priority information, and user intent, the Orchestrator Agent first delegated to a Holdings Analyst to retrieve current client holdings. Next, it elected from a pool of specialized worker agents: Research, Tax, Mathematics, and Guidance agents. Each worker agent operated in parallel, with access to its respective tools: public market data APIs for pricing and fundamental data, a vector database of tax rules and regulations, a tax optimization engine for cost basis computation and tax-loss harvesting strategy generation, and a vector database of portfolio management best practices for contextual guidance.

The factorial conditions correlate with the agents’ access to various components of the orchestration architecture, where the tax engine condition toggled access to sections of agents’ system prompts and the tax guidance condition toggled access to RAG tooling.

The primary dependent variable was percentage tax savings; the proportional reduction in projected federal capital gains tax liability relative to a no-trade baseline:

$$\text{Tax Savings}(\%) = \frac{T_{\text{baseline}} - T_{\text{recommended}}}{T_{\text{baseline}}} \times 100\%$$

5 Results

5.1 Descriptive Statistics

Table 1: Descriptive Statistics for Tax Savings by Condition

Tax Tools	RAG	n	M	SD	Mdn
Off	Off (Baseline)	30	30.55%	99.16	24.01
Off	On (RAG Only)	30	47.73%	45.98	35.63
On	Off (Tax Tools Only)	30	−11.03%	114.96	0.00
On	On (Tax Tools + RAG)	30	−21.02%	161.94	0.00




Figure 1: Distribution of Tax Savings Across Experimental Conditions

When the tax optimization engine was disabled, portfolios achieved positive mean tax savings regardless of RAG status (30.55% and 47.73%). When the tax optimization engine was enabled, mean tax savings became negative (−11.03% and −21.02%), indicating that the engine’s trade recommendations on average increased tax liability. Variability was highest in the Tax Tools + RAG condition ($SD = 161.94$).

5.2 Omnibus 2 × 2 Repeated-Measures ANOVA

A two-factor repeated-measures ANOVA was conducted with tax optimization engine (on/off) and tax documents knowledge retrieval (on/off) as within-subjects factors and tax savings percentage as the dependent variable. Table 2 presents the omnibus results.

Sphericity was not violated ($\epsilon = 1.0$ for all effects in this 2-level design; Greenhouse-Geisser correction is unnecessary with $df = 1$). The main effect of the tax optimization engine was significant, $F(1, 29) = 9.17, p = .005, \eta_p^2$

Table 2: Repeated-Measures ANOVA for Tax Savings Percentage

Source	SS	df	MS	F	p	η_p^2
Tax Tools	91,289.80	1, 29	91,289.80	9.17	.005	.240
RAG	389.09	1, 29	389.09	0.04	.841	.001
Tax Tools $\times$ RAG	5,536.84	1, 29	5,536.84	0.36	.553	.012




Figure 2: Distribution of Tax Savings Across Experimental Conditions

= .240, indicating a large effect. Portfolios evaluated with the tax engine enabled ($M = -16.03\%$) achieved significantly lower tax savings than those without it ($M = 39.14\%$). The main effect of RAG was not significant, $F(1, 29) = 0.04, p = .841, \eta_p^2 = .001$. The two-way interaction was not significant, $F(1, 29) = 0.36, p = .553, \eta_p^2 = .012$, indicating that the detrimental effect of the tax engine did not depend on whether RAG was also enabled. Because the interaction was not significant, main effects are interpreted directly.

5.3 Post-Hoc Pairwise Comparisons

Bonferroni-corrected pairwise comparisons were conducted to decompose the significant main effect of the tax optimization engine. Table 3 presents the within-subjects pairwise results.

Table 3: Bonferroni-Corrected Pairwise Comparisons

Contrast	Comparison	t(29)	p corrected	Hedges' g
Tax Tools (main effect)	Off vs. On	3.03	.005	0.67
RAG (main effect)	Off vs. On	-0.20	.841	-0.04
Interaction: Tax Tools = Off	RAG Off vs. On	-1.04	.618	-0.22
Interaction: Tax Tools = On	RAG Off vs. On	0.27	1.000	0.07

The significant main effect of Tax Tools ($t(29) = 3.03$, $p_{bonf} = .005$, $g = 0.67$) reflects a medium-to-large effect: disabling the tax optimization engine yielded approximately 55 percentage points more tax savings than enabling it. No other comparisons reached significance after Bonferroni correction.

5.4 Sensitivity Analysis

To assess robustness to extreme values, outliers on tax savings percentage were identified using the IQR method (bounds: [-150.00%, 250.00%]). Because the design is within-subjects, entire portfolios were excluded if any of their four observations fell outside these bounds. Seven portfolios were excluded, leaving 23 portfolios (92 observations).

Table 4: Sensitivity Analysis RM-ANOVA After Outlier Exclusion (N = 23)

Source	F(1, 22)	p	η_p^2
Tax Tools	4.38	.048	.166
RAG	1.62	.216	.069
Tax Tools × RAG	0.23	.635	.010

The pattern of results was unchanged: the tax optimization engine main effect remained significant ($p = .048$), while RAG and the interaction remained non-significant. The effect size was attenuated ($\eta_p^2 = .166$ vs. $.240$) due to the removal of extreme cases, but the direction and conclusion are consistent with the primary analysis.

5.5 Supplementary Analysis: Tax Savings Amount

A parallel 2×2 repeated-measures ANOVA was conducted on tax savings amount (dollars) as a secondary dependent variable.

Table 5: Repeated-Measures ANOVA for Tax Savings Amount (\$)

Source	F(1, 29)	p	η_p^2
Tax Tools	5.47	.026	.159
RAG	0.02	.890	.001
Tax Tools × RAG	1.14	.294	.038

The dollar-value analysis converged with the percentage-based findings: the tax optimization engine significantly reduced tax savings in absolute terms ($p = .026$), while RAG and the interaction were not significant.

5.6 Non-Parametric Robustness Check

Given the high variability and non-normal distribution of tax savings percentage, a Friedman test was conducted as a non-parametric alternative to the omnibus ANOVA. The Friedman test was significant: $\chi^2(3) = 10.47$, $p = .015$, $W = 0.116$, confirming that the conditions differed significantly in tax savings when no distributional assumptions are imposed. This is consistent with the parametric findings.

5.7 Summary

Across all analyses the primary 2×2 repeated-measures ANOVA, the sensitivity analysis with outlier exclusion, the supplementary dollar-value analysis, and the non-parametric Friedman test, results converge on a single conclusion: enabling the tax optimization engine significantly reduced tax savings (or increased tax liability), while RAG knowledge retrieval had no measurable effect and did not interact with the tax engine. The effect of Tax Tools was large ($\eta_p^2 = .240$, $g = 0.67$), robust to outlier exclusion, and consistent across both percentage and dollar-denominated outcomes.

6 Conclusion

The results largely contradicted our initial hypotheses. The omnibus ANOVA revealed a significant main effect of the tax optimization engine, $F(1, 29) = 9.17$, $p = .005$, $\eta_p^2 = .240$, representing a large effect. Portfolios evaluated with the tax engine enabled produced mean tax savings of -16.0%, compared to +39.1% when the engine was disabled a

difference of approximately 55 percentage points. The main effect of RAG was not significant, $F(1, 29) = 0.04$, $p = .841$, and the Tax Tools $\times$ RAG interaction was not significant, $F(1, 29) = 0.36$, $p = .553$, indicating that the detrimental effect of the tax engine did not depend on whether RAG was also enabled. This pattern was confirmed by a non-parametric Friedman test, $\chi^2(3) = 10.47$, $p = .015$, and held under a sensitivity analysis excluding outlier portfolios ($p = .048$).

Rather than underperforming all augmented conditions, the Baseline agent achieved a mean tax savings of 30.5%, suggesting that the pre-trained language model’s internalized financial knowledge is sufficient to generate competent tax-loss harvesting recommendations without explicit tooling. The RAG Only condition performed best (47.7%), though this advantage over Baseline was not statistically significant after Bonferroni correction. Most critically, both conditions with the tax engine enabled produced negative mean tax savings, Tax Tools Only (-11.0%) and Tax Tools + RAG (-21.0%) meaning the engine’s recommendations on average increased tax liability. Bonferroni-corrected post-hoc comparisons confirmed the significant Tax Tools main effect ($t(29) = 3.03$, $p_{\text{bonf}} = .005$, $Hedges'g = 0.67$), a medium-to-large effect.

These findings carry several important limitations. All scenarios were synthetically generated and evaluated under a single market regime, limiting ecological validity. The study did not verify compliance with wash-sale rules or assess the qualitative soundness of individual trade recommendations beyond aggregate tax impact. Additionally, while the within-subjects design increased statistical power by controlling for portfolio-level variability, the high within-condition variance (SDs ranging from 46% to 162%) suggests substantial heterogeneity in how the agent handles different portfolio configurations.

Despite these limitations, the results offer practical insight for developers of AI-driven wealth management systems: augmenting language model agents with domain-specific tooling does not guarantee improved performance and may in fact degrade it if the integration introduces friction or conflicting optimization signals. The tax engine’s negative impact likely reflects over-constrained decision-making, where rule-based cost-basis calculations conflicted with the agent’s broader reasoning about optimal trade selection. Future work should examine larger and more diverse portfolio samples, incorporate real client data under multiple market regimes, refine the tax engine’s integration to complement rather than constrain agent reasoning, and evaluate recommendation quality along dimensions beyond tax savings alone including wash-sale compliance, portfolio risk characteristics, and alignment with individual investor objectives.

References

- Tanvi Bachhav, Tejas Parekh, Kanav Rathor, Aditya Munshi, and Moumita Roy. AI-based personalised portfolio allocation engine for Indian retail investors: A multi-algorithm optimisation approach with explainable AI. *Research Article*, 2024. doi: 10.21203/rs.3.rs-9509094/v1.
- Andrew L. Berkin and Jianming Ye. Tax management, loss harvesting, and FIFO accounting. *Financial Analysts Journal*, 59(4):91–102, 2003. doi: 10.2469/faj.v59.n4.2548.
- Sanjiv E. Chaudhuri, Todd C. Burnham, and Andrew W. Lo. An empirical evaluation of tax-loss-harvesting alpha. *Financial Analysts Journal*, 76(3):99–108, 2020. doi: 10.1080/0015198X.2020.1760064.
- Jacob Cohen. *Statistical Power Analysis for the Behavioral Sciences*. Lawrence Erlbaum, 2nd edition, 1988.
- Raul Israelov and Jing Lu. Optimized tax loss harvesting: A simple algorithm and framework. Technical report, SSRN, 2022. URL <https://ssrn.com/abstract=4152425>.
- Samantha Krasner and Noah Sosner. Loss harvesting or gain deferral? a surprising source of tax benefits of tax-aware long-short strategies. *The Journal of Wealth Management*, 27(1), 2023. URL <https://ssrn.com/abstract=4584287>. Summer 2024.
- Douglas C. Montgomery. *Design and Analysis of Experiments*. Wiley, 9th edition, 2017.
- Reza Pishehvar. A three-phase foundation model for tax-aware personalized portfolio management. *arXiv preprint arXiv:2606.30997v3 [cs.AI]*, 2026. U.S. Provisional Patent Application No. 64/101,198, filed June 29, 2026.
- Clemens Sialm and Noah Sosner. Tax-aware investing. *Journal of Financial Economics*, 127(3):475–494, 2018. doi: 10.1016/j.jfineco.2017.10.009.