

Extended Abstract

Motivation The prediction of future stock performance is a notoriously difficult problem due to the complex interplay of quantitative, behavioral, and stochastic factors. While many quantitative methods exist, there is growing interest in leveraging the reasoning capabilities of Large Language Models (LLMs) on unstructured data like corporate earnings call transcripts. This project aims to develop a model capable of predicting future stock performance by applying state-of-the-art reinforcement learning techniques to an LLM. A core objective is not only to investigate the effectiveness of these techniques to improve the baseline performance of a reasoning LLM, but also to compare their predictive capabilities with those of human analysts.

Method Our methodology is a multi-stage pipeline designed to incrementally build financial reasoning capabilities into a base LLM. We begin by collecting and processing over 120,000 earnings transcripts. We establish performance baselines using a linear TF-IDF model and the pre-trained Qwen 3 1.7B LLM. In our successful pipeline, the next training step involves Supervised Fine-Tuning (SFT) on the Qwen model using synthetic reasoning traces. We used Gemini Pro 2.5, a "frontier" LLM, to generate reasoning traces via two techniques: (a) rejection sampling and (b) reasoning generation in hindsight knowing the ground truth label. The second method we call "STaR" for Self Taught Reasoner as it was inspired by Zelikman et al. (2022). The final stage employs Reinforcement Learning with Verifiable Rewards (RLVR), using Group Relative Policy Optimization (GRPO), to optimize the fine-tuned model for generalization and improved performance.

Implementation The project utilizes a dataset of earnings call transcripts from Financial Modeling Prep, covering 3,000 U.S. stocks over 20 years. Target labels ("STRONG BUY" to "STRONG SELL") are derived from 1-month forward returns, which are percentile-ranked within their respective sector to isolate company-specific performance. For computational efficiency, all SFT and RLVR training stages leverage Low-Rank Adaptation (LoRA), and model inference is accelerated using the vLLM framework. The base model for fine-tuning is Qwen3 1.7B (Yang et al. (2025)), while reasoning traces are generated with Gemini Pro 2.5 (DeepMind (2023)).

Results Our initial experiments revealed that a simple linear TF-IDF baseline (F1 score: 0.2587) significantly outperformed both the base Qwen 1.7B LLM (F1: 0.0847) and a human analyst consensus benchmark (F1: 0.1741). Given limited compute, we were unable to outperform the linear baseline with a reasoning LLM pipeline. However, we are able to demonstrate a working pipeline where each step, (a) baseline Qwen3, (b) supervised fine-tuning on synthetic data, and (c) RLVR warm-started with the prior step, is a successive improvement. Given our compute budget (one Nvidia H100), we were only able to fine-tune on 414 synthetic examples, and train with RLVR on 1,000 samples. Given that the entire prepared training set was 90,000 samples, there is substantial opportunity to increase the size of the training sample. Our pipeline (including inference on the full test set) takes approximately 20 hours to run.

Discussion The performance of the TF-IDF model confirmed valuable predictive signals in the transcript text. Our results demonstrate that a multi-stage training process—first using SFT on reasoning traces, then optimizing with RLVR—can improve a reasoning model's predictive capability. Despite this success, the study has limitations, including potential data leakage and the need for more rigorous ablation studies. Future work should also explore integrating external tools and data sources and applying this pipeline to larger models and a larger sample of the training set.

Conclusion This project presents a novel pipeline for enhancing the financial reasoning capabilities of Large Language Models by combining supervised fine-tuning with reinforcement learning on synthetic reasoning traces. While the pipeline did not surpass a simple TF-IDF baseline within current compute constraints, it consistently improved model performance at each stage, establishing a promising foundation for future development. Our findings suggest that LLMs, when equipped with structured reasoning supervision and carefully tuned through RL, hold significant potential in financial forecasting tasks. With expanded training data, compute resources, and additional ablation studies, this approach could evolve into a viable alternative or complement to traditional and human analyst-driven methods.

Reinforcement Learning for Predicting Future Stock Performance

Jonathan Larkin

Stanford Center for Professional Development
jonathan.r.larkin@gmail.com

Ram Komarraju

Stanford Center for Professional Development
ram.komarraju@gmail.com

Tamika Bassman

Stanford Center for Professional Development
tbassman@alumni.stanford.edu

Abstract

Predicting future stock performance remains a difficult challenge due to the interplay of financial, behavioral, and stochastic factors. While traditional quantitative models have shown success, the potential of large language models (LLMs) to reason over unstructured financial text, such as earnings call transcripts, is still underexplored. In this work, we propose a novel pipeline that enhances a base LLM’s forecasting ability through supervised fine-tuning on synthetic reasoning traces and reinforcement learning with verifiable rewards (RLVR). We train and evaluate our approach using a corpus of 120,000 earnings call transcripts covering 3,000 U.S. companies. To benchmark performance, we establish two baselines: (1) a simple linear TF-IDF model trained on the *entire* 90,000-example training set, and (2) a zero-shot Qwen3 1.7B LLM. Notably, the TF-IDF model, despite its simplicity, achieves the highest macro F1 score (0.2597), indicating strong signal in the raw transcript text when trained on the full dataset. In contrast, our LLM pipeline, limited by compute to less than 2% of the training set, including synthetic data, was unable to surpass the TF-IDF baseline but consistently improved with each stage of fine-tuning and RL, ultimately outperforming a human analyst consensus benchmark. These findings underscore both the promise and the limitations of LLMs for financial forecasting under resource constraints. Our work demonstrates that, even with limited data, reinforcement learning-enhanced LLMs can learn interpretable reasoning patterns and gradually improve in accuracy. With larger-scale training and expanded compute, this methodology holds promise as a complement to traditional models in financial analysis.

1 Introduction

There is no shortage of professional analysts today using quantitative methods to predict stock returns and thereby make financial investment decisions. However, given the complex interplay of stochastic, quantitative, and behavioral factors, the task of predicting future stock returns for the purpose of making investment decisions remains among the most challenging problems in the industry and in the area of financial machine learning. In particular, it remains an open research question whether and to what extent reasoning large language models – that is, LLMs capable of using “chain-of-thought” step-by-step thinking to approach a problem – can accurately forecasting the future performance of stocks.

In the United States, publicly listed companies report their earnings results quarterly and host live conference calls at which the CEO and CFO present and take questions from professional analysts.

As such, the text transcripts of these calls are rich with backward and forward looking information pertaining to the financial performance of the company. The extensive length of these transcripts, however, poses a challenge to the ease with which they can be utilized for informing future predictions in either an automated or manual manner.

In this work, we present a novel methodology which utilizes state-of-the-art techniques in reinforcement learning to address the task of evaluating how well reasoning LLMs can be fine-tuned to accurately predict future stock performance at the one month investment horizon. Central to our approach is the use of reinforcement learning with verifiable rewards following supervised fine-tuning. We run a baseline on a pre-trained LLM, Qwen3-1.7B (Yang et al. (2025)). We successively improve this baseline by (a) fine-tuning the model on synthetic reasoning traces, and (b) training with reinforcement learning with the GRPO objective with a binary correctness reward and a format correctness reward.

2 Related Work

Financial applications have been an area of increasing interest to reinforcement learning and LLM-related research. As noted recently by Bai et al. (2025), two of the key research challenges remaining to be addressed in RL-based solutions to finance-related tasks are explainability (i.e., interpretability and transparency of predictions) and robustness (i.e., stability under varying conditions).

Some researchers in recent years have attempted to address these challenges. Koa et al. (2024) used tweets to generate explainable 1-day forward predictions for a total of 55 stocks and used GPT-4 as a judge to rate explanation quality. The RL framework they propose uses a verbal self-reflective agent and Proximal Policy Optimization (PPO) to teach the LLM. The authors noted that future research should aim to improve the robustness of the predictions against compounding error to reduce the need for human-in-the-loop, and also to expand on the sophistication of the metrics used to evaluate explanation quality.

Some researchers have found success in utilizing earnings transcripts as an input to an LLM. Jha et al. (2024) used ChatGPT 3.5 to extract insights regarding company managers’ future capital expenditure expectations from earnings transcripts and condense them into a so-called ChatGPT Investment Score. While this work did not utilize RL techniques to further enhance the performance of the LLM, their findings do suggest that LLMs are able to elicit corporate policy information that is relevant to financial forecasting from earnings transcripts that is otherwise not readily apparent or available. Relatedly, Ni et al. (2024) used earnings transcripts as well as other financial data to refine an LLM’s performance on “Long” and “Short” investment recommendations using instruction fine-tuning and quantized low-rank adaptation (QLoRA) compression. While not an RL-based approach, their work exceeded the performance of a GPT-4 benchmark and illustrated the benefits of a rich fine-tuning dataset.

To our knowledge, work done to-date on RL for financial predictions has yet to develop more sophisticated metrics by which to assess the reasoning trace and/or explainability of the model’s prediction, quantify the benefit of efficient tool-calling, or leverage GRPO (Shao et al. (2024)) for this task. These unaddressed challenges motivate the work herein and detailed in the technical outline which follows.

3 Method

Our approach follows a structured pipeline, as illustrated in Figure 1, to develop a model capable of predicting future stock performance from earnings call transcripts. The core of our methodology is to enhance a base Large Language Model (LLM) through several stages of supervised fine-tuning (SFT) and reinforcement learning with verifiable rewards (RLVR).

The key steps in our methodology are as follows:

1. **Data Collection and Processing:** A large dataset of earnings call transcripts and their corresponding stock performance data are collected and processed. To each transcript, we assign a ground truth label which mirrors how professional stock analysts and investors rate stocks. This label can take on one of five values: STRONG BUY, BUY, HOLD, SELL, or

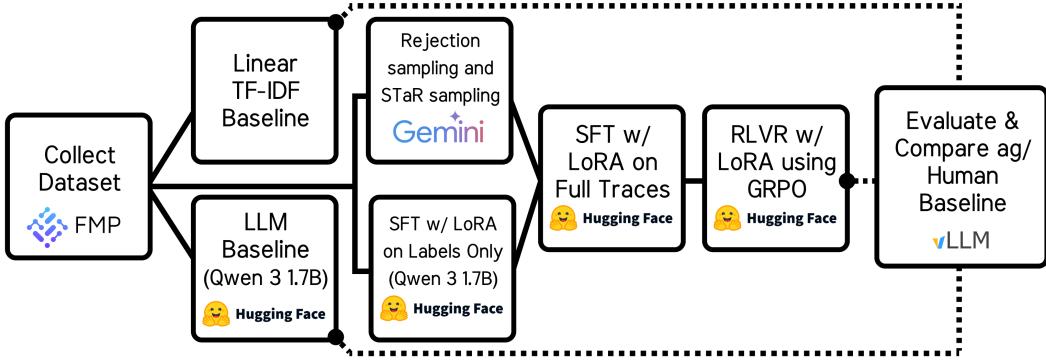


Figure 1: Methodology overview.

STRONG SELL. We also collect human analyst ratings and reference data. Further details on the dataset are presented in Section 4.1 and on the processing/labeling of the dataset in Appendix A.

2. **Baselines:** Two initial baselines are established to serve as our starting point, (1) a traditional linear TF-IDF model and (2) the pre-trained, base LLM chosen for this project without any fine-tuning applied (Qwen 3 1.7B). The TF-IDF baseline serves to gauge the signal strength of the text of the earnings transcripts alone.
3. **Supervised Fine-Tuning (SFT) on Labels (optional step):** In order to help the LLM learn the structure of the task prior to undergoing reinforcement learning, we attempted two stages of supervised fine-tuning (SFT). The first fine-tuning stage involves training the base LLM (Qwen 3 1.7B) on the transcripts with their final performance labels only (e.g., “STRONG BUY”, “HOLD”) as the target. These “ground-truth” labels are readily available from the initial data processing step and thus require minimal effort. This step aims to de-bias the model from its pre-training data and align it with our specific classification task. A total of 5,000 labeled examples were used in this step. Further information on the chosen base model is provided in Section 4.2. This step is labeled “optional” as we tried multiple variations of the full methodology which either included or excluded this step.
4. **Reasoning Trace Generation:** In anticipation of the second stage of fine-tuning, and to enable the model to learn the *reasoning* behind a prediction, we use a more powerful model (Gemini Pro 2.5) to generate reasoning traces. We employ two techniques for this:
 - **Rejection Sampling:** We generate 1,600 rationales and keep only those that lead to the correct ground-truth label.
 - **Self-Taught Reasoning (STaR):** Provide the model with the ground-truth label and prompt it to generate a plausible rationale in hindsight. We prune these traces to only those that (a) show a rationale which builds up to a conclusion (as opposed to a rationale which states the conclusion upfront and then justifies it), and (b) do not indicate that the rationale was generated in hindsight.
5. **SFT on Full Traces:** We conduct a second round of SFT, this time training the model on the full reasoning traces generated in the previous step. This teaches the model not only the final answer but also a step-by-step process by which to arrive at it. A total of 414 examples (combined from both rejection sampling and STaR) were used.
6. **Reinforcement Learning with Verifiable Rewards:** The final stage involves applying reinforcement learning to the fine-tuned model. Reinforcement learning with verifiable rewards (RLVR) utilizes straightforward, objective reward functions to provide feedback to the model. Our task at hand, given the clear label applied to each transcript, is well-suited to leverage RLVR.

Within the RLVR framework, we applied Group Relative Policy Optimization (GRPO), a derivative of the Proximal Policy Optimization (PPO) technique (Shao et al. (2024)). We chose to use the GRPO algorithm because it was designed for enhanced mathematical reasoning and improved memory usage compared to PPO. These particular attributes are

favorable given the mathematical nature of financial analysis and the computational burden involved in working with long text bodies such as earnings transcripts.

Unlike PPO, GRPO does not utilize or train a value model, but rather at each step t uses G samples a group of G outputs $\{o_1, o_2, \dots, o_G\}$ from the old policy $\pi_{\theta_{old}}$, all of which are generated as responses to the same question q , and takes the average reward score across these samples as the baseline. As given in Shao et al. (2024), the objective function of GRPO is:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}[q \sim P_{soft}(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{old}}(O|q)]$$

$$\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left[\frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{old}}(o_{i,t}|q, o_{i,<t})} \hat{A}_{i,t} - \beta \left(\frac{\pi_{ref}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta}(o_{i,t}|q, o_{i,<t})} - \log \frac{\pi_{ref}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta}(o_{i,t}|q, o_{i,<t})} - 1 \right) \right]$$

where ε and β are hyperparameters, and the KL divergence term uses the following unbiased estimator from Schulman (2020):

$$\mathbb{D}_{KL} [\pi_{\theta} || \pi_{ref}] = \frac{\pi_{ref}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta}(o_{i,t}|q, o_{i,<t})} - \log \frac{\pi_{ref}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta}(o_{i,t}|q, o_{i,<t})} - 1$$

As also noted in HuggingFace GRPO documentation, recent studies (e.g., Hu et al. (2025), Yu et al. (2025), Liu et al. (2025)) relating to GRPO have motivated a use of $\beta = 0.0$ as the coefficient on the KL divergence term by showing that it is not required for effective training. We have utilized this choice of β in the experiments documented herein.

We used two reward functions: a binary reward function which gives 1.0 for the exact label match, and a format reward, which gives 0.50 for correctly producing a completion in the form "<think> ...rationale... </think> label".

7. **Evaluation:** The final model is evaluated on the held-out test set and its performance is compared against the initial baselines and a human analyst consensus.

For computational efficiency, we utilize Low-Rank Adaptation (LoRA) for all supervised fine-tuning and reinforcement learning with verifiable rewards stages and the vLLM inference engine for faster model sampling.

4 Experimental Setup

4.1 Dataset and Data Processing

The dataset we used is sourced from Financial Modeling Prep (2025) and consists of approximately 120,000 earnings call transcripts over the past 20 years taken from 3,000 US-listed stocks which themselves span 11 economic sectors. The transcripts are text-heavy, with a median length of around 8,000 tokens.

To each transcript, we assign a ground truth label which mirrors how professional stock analysts and investors rate stocks. This label can take on one of five values: *STRONG BUY*, *BUY*, *HOLD*, *SELL*, or *STRONG SELL*. These labels are constructed based on each stock's 1-month forward return, percentile-ranked against other stocks within the same sector for that day. This within-industry ranking is crucial to remove broad market and sector-wide drifts, isolating company-specific performance. A two-day lag is incorporated in calculating forward returns to account for the timing of earnings releases. Figure 2 illustrates key attributes of the processed dataset. We made a train/test split at January 1, 2023. This split leaves approximately 90k samples in the training set and 18k samples in the test set.

4.2 Models

The baseline LLM which we used and fine-tuned throughout this work was Qwen 3 1.7B (Yang et al. (2025)). We chose this model because, at the time of this report, it represented the state-of-the art for a "small" language model in reasoning benchmark tasks. This model can be trained with LoRA on a single NVIDIA H100 without quantization.

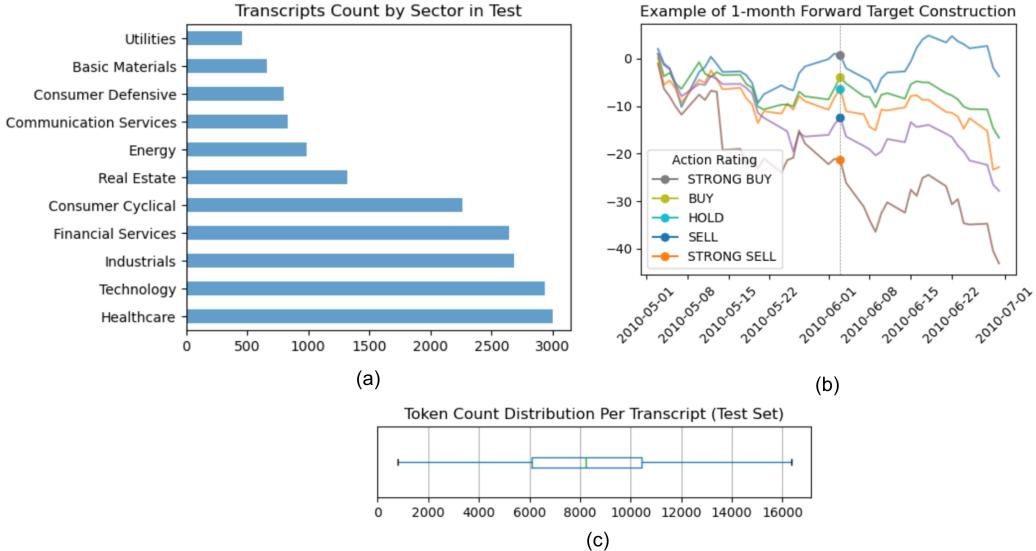


Figure 2: Dataset details: (a) number of transcripts per economic sector; (b) example of 1-month forward target construction; (c) distribution of transcript length (in tokens) in test set.

We used Gemini Pro 2.5 for the task of generating reasoning trace examples on which to fine tune the base model during the supervised fine-tuning stage of the pipeline (see Section 3). This LLM represents a near-state-of-the-art reasoning model at the time of writing.

4.3 Evaluation

The primary metric for evaluation is the macro F1 score, which is suitable for our multi-class classification problem and potential class imbalance. The F1 score is the harmonic mean of precision and recall, providing a single measure that balances both false positives and false negatives. It is particularly useful when the cost of misclassification is high and when classes are imbalanced.

Macro F1 calculates the F1 score for each class independently and then averages them, giving equal weight to each class regardless of its frequency in the dataset. This makes it especially appropriate for our case, where some classes (e.g., “STRONG SELL”) may be underrepresented. Unlike overall accuracy, which can be skewed by dominant classes, the macro F1 score ensures that performance across all classes is fairly represented, providing a more holistic view of the model’s effectiveness.

We compare our model’s performance against the linear TF-IDF baseline, the initial Qwen 3 1.7B baseline, and a human analyst consensus rating. The use of the F1 score is common in evaluating multi-class classification models in machine learning, though it presents a lower-bound on the usefulness of our model. Of note, the F1 score does not account for “near misses” (e.g., rating something a “BUY” when it is in reality a “STRONG BUY”).

A full example of an input to the model, including prompt and earnings transcript, is shown in Appendix E. An example output generated by the model is shown in Appendix D.

5 Results

Our evaluation explored the predictive capabilities of large language models (LLMs) and traditional baselines for forecasting stock returns using earnings call transcripts. We systematically compared raw LLMs, fine-tuned variants, traditional machine learning methods, and human analyst baselines. Our goal was not only to assess raw predictive performance, but also to understand the role of model bias, training strategies, and interpretability in financial forecasting tasks.

5.1 Quantitative Evaluation

We evaluated a variety of modeling strategies on their ability to classify earnings call transcripts into one of five target labels (STRONG BUY, BUY, HOLD, SELL, STRONG SELL) based on 1-month forward sector-relative returns. The primary metric for evaluation is the macro F1 score, chosen for its robustness to class imbalance and its balanced emphasis on precision and recall across all classes.

Our initial baselines illustrate the relative difficulty of this task. The pre-trained Qwen3 1.7B model, without any fine-tuning, performed poorly with a macro F1 score of 0.0897, revealing a strong optimism bias. This result should lead financial researchers to proceed with caution when using LLMs. In contrast, a simple linear TF-IDF model trained on the full dataset achieved a significantly higher F1 score of 0.2597, indicating that important predictive signals are embedded in the raw textual content of earnings transcripts.

We also compared model performance against human analyst consensus ratings derived from Financial Modeling Prep (FMP). These ratings were converted into class predictions mapped to our labels. The resulting macro F1 score was 0.1743, lower than the TF-IDF model but well above the base LLM, providing a meaningful intermediate benchmark.

Performing supervised fine-tuning on Qwen3 using 5,000 examples with only the final label as the target resulted in a large improvement over the base Qwen3 model, reaching an F1 score of 0.2254. This shows that even minimal task-specific supervision substantially improves LLM performance in this domain. However, the model continued to display a positive class skew and this fine-tuning step destroyed the models ability to produce a reasoning trace.

To improve model interpretability and reasoning, we next fine-tuned the model on 414 synthetic reasoning traces generated by Gemini Pro 2.5 using both rejection sampling and Self-Taught Reasoning (STaR). When fine-tuned on traces alone (skipping the label-only SFT step), the model achieved an F1 score of 0.1920.

Fine-tuning in two stages—first on labels, then on traces—resulted in a modest F1 score of 0.2094. This suggests a trade-off: integrating structured reasoning improves alignment with human-like explanations, but may degrade final classification accuracy slightly if the training set is small or imbalanced.

Applying RLVR with Group Relative Policy Optimization (GRPO) to the traces-fine-tuned model yielded a final F1 score of 0.2084, maintaining comparable performance while potentially improving sample efficiency and policy robustness. A second RLVR run that started from the combined labels+traces fine-tuned model reached 0.2061, showing no clear performance gain but supporting the overall viability of reinforcement learning as a refinement step.

These results are summarized in Table 1. While the TF-IDF baseline remains the strongest performer under current compute constraints, our LLM pipeline showed consistent improvements at each stage of fine-tuning and RL. Notably, the final pipeline variants surpassed the human analyst consensus, highlighting the potential of reinforcement learning-enhanced LLMs for financial forecasting. However, larger training datasets, improved balance in synthetic traces, and expanded compute resources will be critical to closing the performance gap with traditional baselines.

We are encouraged that the progressive training pipeline of Qwen3 baseline, supervised fine-tuning on synthetic traces, and finally Reinforcement Learning with Verifiable Rewards, shows **progressive improvement**. Intuitively, the SFT step on the synthetic reasoning traces serves to eliminate Qwen3-1.7B’s “optimism bias”. The RLVR step serves to draw out additional performance from the model.

5.2 Qualitative Analysis

A key theme across models was the prevalence of positive bias. Both the raw Qwen3 and the reasoning traces generated by Gemini 2.5 Pro exhibited consistent optimism, often overestimating favorable outcomes. Despite Gemini’s frontier-level capabilities, its generated traces frequently leaned toward bullish interpretations, reinforcing the observation that state-of-the-art models can be unreliable in financial settings without targeted adaptation. We passed 1,690 transcripts to Gemini 2.5 Pro from a class balanced set. After keeping only correct outcomes (i.e., rejection sampling), we were left with a highly imbalanced dataset of only 355 samples. The class imbalance is shown in Figure 3, “Rejection Sampling.” We took as inspiration the STaR (Zelikman et al. (2022)) paper to generate

Table 1: Performance Comparison

Method	F1 Score
Qwen3 1.7B baseline	0.0897
Human baseline	0.1743
TF-IDF baseline	0.2597
Qwen3 Labels SFT	0.2254
Qwen3 Traces SFT	0.1983
Qwen3 Labels + Traces SFT	0.2094
Qwen3 Traces SFT + RLVR	0.2084
Qwen3 Labels + Traces SFT + RLVR	0.2061

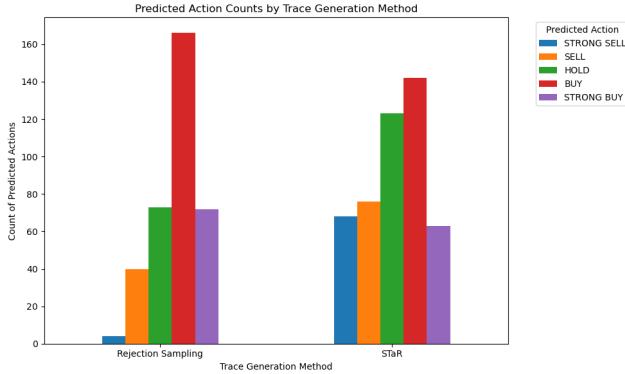


Figure 3: Synthetic Reasoning Traces Generation

additional synthetic data. In this method, we gave Gemini 1,000 transcripts with the ground truth label and asked it to generate a reasoning trace in hindsight. We then separately used Gemini 2.5 Pro as a "LLM-as-judge" to ensure that results had both of the following characteristics: (a) the trace built up reasoning, as opposed to first stating the conclusion and then justifying, and (b) the trace had no indication that this reasoning was done in hindsight. Surprisingly, we were only left with only 472 traces that satisfied the criteria, and these 472 traces as well were imbalanced as shown in Figure 3. Lastly, we had to further prune this data, as we chose a 16,384 token-limit for all transcripts in all experiments.

Notably, human analyst predictions shared this same optimistic skew, suggesting that this bias may reflect broader sentiment tendencies in the financial domain, rather than being purely a model artifact. In contrast, the TF-IDF classifier produced more balanced predictions, likely due to its reliance on empirical word-frequency patterns rather than contextual language modeling.

While reasoning trace-based fine-tuning introduced valuable interpretability, it also posed risks to predictive calibration, as shown by the drop in performance after applying traces. This highlights a critical tension: reasoning fidelity and interpretability may not always align with prediction accuracy, particularly in high-stakes domains like finance.

Lastly, we note that the SFT on labels caused the LLM to lose its ability to provide a thinking trace. The completion target for this run was in the form `<think> </think> label`. We ran a RLVR with GRPO run warm-starting on this LoRA, but the run did not finish in time (estimated 75 hours to train). This training run was especially long we believe because it needed to re-teach the model to provide a reasoning trace.

6 Discussion

The performance of the TF-IDF model confirmed that the transcript text contains valuable predictive signals, while the poor performance of the base reasoning LLM and human analysts highlighted the

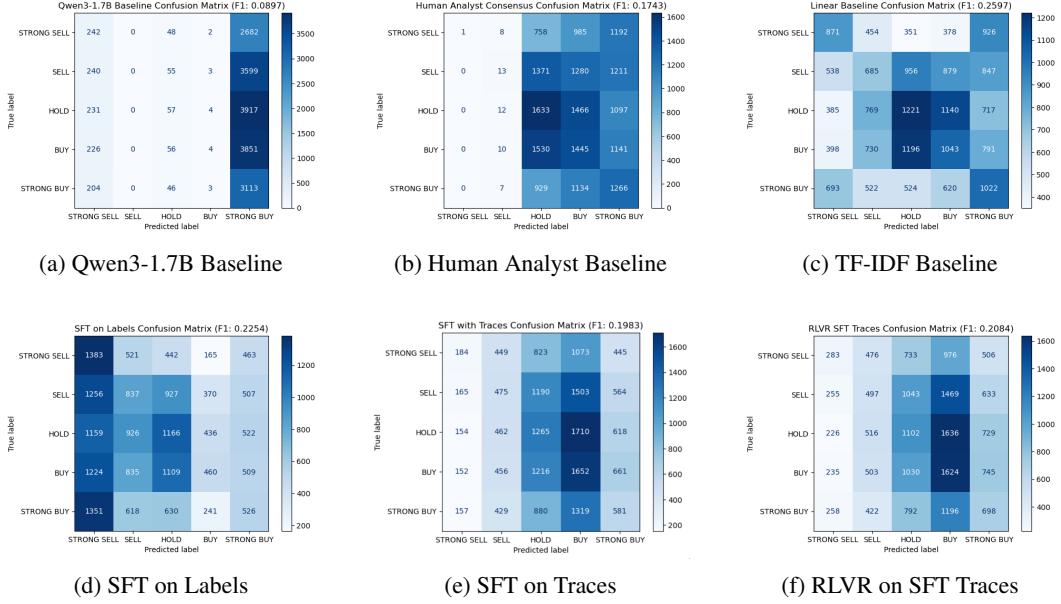


Figure 4: Visualization of results as confusion matrices across methods and metrics.

difficulty of overcoming inherent positive biases. Our results demonstrate that a multi-stage training process with SFT followed by RLVR can improve the predictive capability of a base reasoning LLM.

Despite the success of this approach, this study has several limitations that present avenues for future exploration. Methodologically, there is a potential for data leakage - our 20-year-long dataset may contain historical events and performance data that were part of the LLM’s original pre-training corpus. Future studies could mitigate this by using more recent data that post-dates the model’s knowledge cutoff, or by creating anonymized or synthetic examples to ensure a more rigorous test of the model’s reasoning abilities. A more careful ablation study is needed to precisely quantify the contribution of each stage of our pipeline—from the initial SFT on labels to the final RLVR optimization. More comprehensive hyperparameter exploration of the number of generations, number of iterations, and KL divergence coefficient would also benefit the overall model performance and potentially the efficiency of the modeling pipeline.

The current implementation relies solely on the text of the transcript, but a human analyst would use tools and has access to additional data. Integrating tool use, such as a calculator for financial ratios or an API to pull applicable data, would allow the model to use additional data and quantitative metrics. Future iterations could also employ an LLM-as-a-judge to provide process-based rewards, evaluating the quality of the reasoning steps themselves rather than just the final outcome. Additional evaluation metrics such as hinge loss can provide better insight into the model’s performance improvement at various stages of the pipeline. Finally, while we used a 1.7B parameter model for efficiency, exploring the performance of this pipeline on larger, more capable foundation models is a logical next step, though it would come at a significant computational cost.

7 Conclusion

In summary, our work explores the intersection of large language models and financial forecasting, introducing a novel pipeline that combines synthetic reasoning supervision with reinforcement learning to enhance predictive performance on earnings call transcripts. While constrained by limited compute and data volume, we demonstrate that each stage of our pipeline (baseline, supervised fine-tuning, and RL with verifiable rewards) provides measurable performance gains. Although our approach, trained on less than 2% of the training data, did not surpass a simple TF-IDF baseline trained on the *full training data*, it outperformed human analysts and highlighted the potential of structured reasoning in financial LLMs. These findings point toward a promising future for interpretable, LLM-driven financial analysis, especially as compute, data quality, and training strategies continue to

improve. Future work will focus on scaling the methodology, integrating external tools, and refining reward mechanisms (including dense progress rewards) to close the gap between reasoning-rich models and strong baseline performance.

8 Team Contributions

- **Jonathan:** Data acquisition; target construction; linear baseline; code implementation; project strategy and various training and evaluation runs; report writing; presentation and poster.
- **Ram:** Code implementation; project strategy and various training and evaluation runs; report writing; presentation and poster.
- **Tamika:** Code implementation; project strategy and various training and evaluation runs; report writing; presentation and poster.

Changes from Proposal All team members contributed to core project components, as anticipated in our original plan.

References

Yahui Bai, Yuhe Gao, Runzhe Wan, Sheng Zhang, and Rui Song. 2025. A review of reinforcement learning in financial applications. *Annual Review of Statistics and Its Application* 12, 1 (2025), 209–232.

Google DeepMind. 2023. Gemini: A Family of Highly Capable Multimodal Models. *arXiv preprint arXiv:2312.11805* (Dec 2023). <https://doi.org/10.48550/arXiv.2312.11805>

Financial Modeling Prep. 2025. Financial Data API. <https://site.financialmodelingprep.com/>. Accessed: 2025-04-24.

Jingcheng Hu, Yinmin Zhang, Qi Han, Dixin Jiang, Xiangyu Zhang, and Heung-Yeung Shum. 2025. Open-Reasoner-Zero: An Open Source Approach to Scaling Up Reinforcement Learning on the Base Model. *arXiv:2503.24290 [cs.LG]* <https://arxiv.org/abs/2503.24290>

Manish Jha, Jialin Qian, Michael Weber, and Baozhong Yang. 2024. *ChatGPT and corporate policies*. Technical Report. National Bureau of Economic Research.

Kelvin JL Koa, Yunshan Ma, Ritchie Ng, and Tat-Seng Chua. 2024. Learning to generate explainable stock predictions using self-reflective large language models. In *Proceedings of the ACM Web Conference 2024*. 4304–4315.

Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. 2025. Understanding R1-Zero-Like Training: A Critical Perspective. *arXiv:2503.20783 [cs.LG]* <https://arxiv.org/abs/2503.20783>

Haowei Ni, Shuchen Meng, Xupeng Chen, Ziqing Zhao, Andi Chen, Panfeng Li, Shiyao Zhang, Qifu Yin, Yuanqing Wang, and Yuxi Chan. 2024. Harnessing earnings reports for stock predictions: A qrlora-enhanced llm approach. In *2024 6th International Conference on Data-driven Optimization of Complex Systems (DOCS)*. IEEE, 909–915.

John Schulman. 2020. Approximating KL Divergence. <http://joschu.net/blog/kl-approx.html>

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300* (2024).

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men,

Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. Qwen3 Technical Report. arXiv:2505.09388 [cs.CL] <https://arxiv.org/abs/2505.09388>

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. 2025. DAPO: An Open-Source LLM Reinforcement Learning System at Scale. arXiv:2503.14476 [cs.LG] <https://arxiv.org/abs/2503.14476>

Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. STaR: Bootstrapping Reasoning With Reasoning. arXiv:2203.14465 [cs.LG] <https://arxiv.org/abs/2203.14465>

A Target Calculation Details

From the FMP dataset, we pull all transcripts for US stocks in the past 15 years for which the company has at least 4 transcripts. This avoids nuances with newly IPO'd stocks or companies undergoing some kind of distress or corporate action. This is approximately 3,000 stocks and 120,000 transcripts in total. Note that the data does not include dead stocks; this is a limitation with our analysis, which implies we exclude situations where the forecast was to BUY and the stock went bankrupt, or when the forecast was SELL and the stock was taken over for cash.

For each of these roughly 3,000 stocks, we get adjusted closing price time series. From these, we calculate daily log returns. These returns are then summed over subsequent forward windows (1m, 3m, 6m). We add a delay of 2 days at the inception of the forward windows. This is to account for the uncertainty of whether earnings conference calls were post-market close or pre-market and also to account for the convention of trading at the closing price on the following day.

This produces a panel dataset: for each stock, for each day, we have forward cumulative log returns for 1m, 3m, and 6m. These forward returns are then percentile ranked within industry classification per day. This ranking is critical and eliminates market and industry drift across related stocks. Classification targets are then formed: +1 if the forward return is in the top third within industry, 0 if it is in the middle third, and -1 if it is in the bottom third. We join the transcript dataset with only the relevant rows from the returns and targets panel.

This results in a final dataset in which each sample contains a ticker symbol, transcript date, industry classification, transcript text, 1m/3m/6m within-industry return percentiles, and 1m/3m/6m target quantiles. We split the dataset into train and test using a date cutoff of January 1, 2023 which results in approximately 90,000 samples in the training set and 20,000 samples in the test set.

B Additional Experiments

We ran three additional experiments with reinforcement learning with verifiable rewards and GRPO on the baseline Qwen 3 1.7B model without initially performing supervised fine-tuning (SFT). These results provided useful insights regarding the utility of various reward functions and of performing SFT. All three experiments used the format reward previously described in Section 3. Other choices of hyperparameters for the GRPO algorithm largely aligned with those of the results described in the main report.

- The first experiment used the simple binary reward described in Section 3 and achieved an F1 of 0.07, which is worse than the baseline Qwen 3 1.7B result.
- The second experiment used a more complex reward function which more strongly penalized incorrect extreme (“STRONG SELL” or “STRONG BUY”) predictions. It achieved an F1 score of 0.13, which represents a slight improvement over the baseline Qwen 3 1.7B result.
- The third experiment utilized the complex reward function of the previous experiment as well as an additional “consistency reward”. This reward function was motivated by manual

inspection of model outputs of the previous experiment, which exhibited that the LLM routinely was returning a final prediction which disagreed with the actual reasoning of the thinking trace. This reward function essentially penalized the model when such an inconsistency occurred. The F1 score of this experiment was 0.12. Manual inspection of the outputs of the SFT+RLVR pipeline detailed in the main report indicated that the model appeared to have learned to correct this inconsistency error via SFT.

Figure 5 below highlights the total reward curves during training for the three additional experiments.

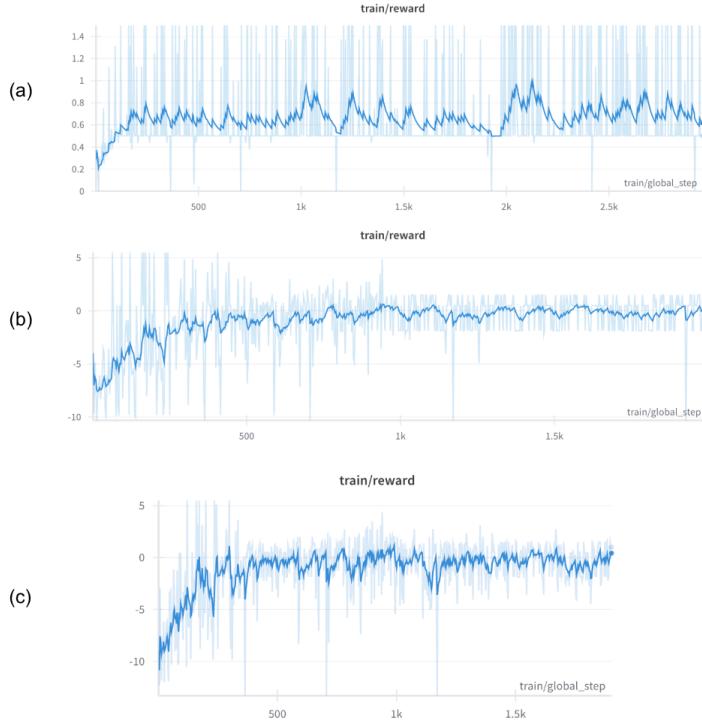


Figure 5: Total reward curves during training of the three additional experiments: (a) binary reward, (b) complex reward; (c) complex reward + consistency reward.

C Implementation Details

C.1 Libraries and Compute

Training and testing implementations utilized Hugging Face libraries, in particular the `tr1`, `datasets`, `transformers`, and `peft` libraries. Training and testing was performed on RunPod using NVIDIA H100 GPUs. Typical production runs of the training+testing pipeline with RLVR required approximately 20 hours on a single NVIDIA H100 GPU. We logged all SFT and RLVR training runs to the Weights & Biases experiment tracking website.

C.2 GRPOConfig Hyperparameters

The table below summarizes key hyperparameters used for the GRPO algorithm (as implemented with the HuggingFace `GRPOConfig` object) for the training runs detailed in the main report. The choice of ‘`num_iterations`’ greater than 1 allows multiple gradient steps to be taken with each training batch, enabling more data-efficient training. The other four hyperparameters listed in Table 2 were chosen so as to encourage sufficient variability during training among the multiple generations sampled by the model for each output.

Hyperparameter	Column 2
temperature	0.95
top_p	1.0
min_p	0.1
num_generations	8
num_iterations	3

Table 2: Hyperparameters for GRPO algorithm.

D Example Rollout (i.e., Generation)

<think>

Okay, let's tackle this question. The user wants to predict the stock's relative performance to stocks in the same sector over the next month, rating it as STRONG BUY, BUY, HOLD, SELL, or STRONG SELL.

First, I need to understand the sector. The company mentioned is Oil-Dri Corporation of America, which is involved in cat litter, animal health products, and other related areas. The sector for this would be the consumer products or retail sector, specifically focusing on pet care and cat litter.

Looking at the earnings call, the company reported strong financial results, with a 11% increase in retail and wholesale product sales, 8% increase in net sales, and a 25% increase in animal health products. They also mentioned improvements in gross profit, lower freight and natural gas costs, and a strong income from operations. The CEO emphasized the team's resilience during the pandemic and the company's ability to meet demand.

Now, the question is about relative performance compared to sector peers. The key factors here are the company's financial performance, market trends, and sector-specific factors. The company is in the consumer goods sector, specifically pet care.

Considering the sector's general performance, the consumer products sector has been affected by the pandemic, with a shift towards e-commerce and online sales. However, the company's performance in retail and wholesale products was strong, with a 11% increase. The animal health segment showed growth, particularly in the Americas, with a 25% increase in sales.

The sector as a whole might be experiencing some challenges, like supply chain issues and market shifts, but the company's strong financials and focus on value-driven products (like cat litter) could give them an edge. The CEO mentioned that the company is well-positioned to compete in the value segment, which is a key factor in the consumer products sector.

However, there are risks. The company is in a sector that's heavily influenced by consumer behavior and economic conditions. If the pandemic's impact subsides, there might be competition from other brands, and the company's ability to maintain pricing and quality could be tested. Also, the sector might face challenges with supply chain disruptions and the ongoing impact of the pandemic.

Given the company's strong financial results, improved gross margins, and focus on value, it's possible that the stock could perform well relative to peers. However, the sector's overall performance and market conditions could affect this. The user wants a relative performance prediction, so I need to consider if the stock is likely to outperform, underperform, or hold compared to others in the same sector.

The answer should be based on the company's strong performance, market position, and sector trends. Since the company is in a sector that's adapting to digital trends and value strategies, and they have shown resilience, the relative performance might be positive. However, there's a risk of sector-wide challenges, so the rating could be HOLD or BUY. But given the strong financial results and the company's focus on value, a BUY might be appropriate.

</think>

BUY

E Example State (i.e., Input)

```
<|im_start|>user
You are an expert institutional equity analyst.
Given the following text, predict the stock's relative performance to stocks in the
same sector over the next month.
You may rate a stock STRONG BUY, BUY, HOLD, SELL, or STRONG SELL.
Think carefully. You must end your output with one of: STRONG BUY, BUY, SELL,
STRONG SELL, or HOLD.
```

Input:

"""Operator: Ladies and gentlemen, thank you for standing by and welcome to the Third Quarter Oil-Dri Corporation of America Earnings Conference Call. At this time, all participants are in a listen-only mode. After the speakers presentation, there will be a question-and-answer session. [Operator Instructions] I would now like to hand the call over to Dan Jaffee, President and Chief Executive Officer. Please go ahead.

Dan Jaffee: Thank you. Welcome everybody to the Oil-Dri third quarter investor teleconference joining me both physically and virtually. So we are in this COVID reality, Susan Kreh, our Chief Financial Officer; Molly VandenHeuvel, our Chief Operating Officer; Jessica Moskowitz, Vice President and General Manager of the Consumer Products Division; Flemming Mads, President of Amlan International; Laura Scheland, General Counsel; and Leslie Garber, Manager of Investor Relations.

Leslie, will you walk us through the Safe Harbor?

Leslie Garber: Thank you, Dan. Welcome, everyone. On todays call, comments may contain forward-looking statements regarding the companys performance in future periods. Actual results in those periods may materially differ. In our press release and our SEC filings, we highlight a number of important risk factors, trends, and uncertainties that may affect our future performance. We ask that you review and consider those factors in evaluating the companys comments and in evaluating any investment in Oil-Dri stock. Thank you for joining us. Dan?

Dan Jaffee: Great. And before I turn it over to Susan to walk us through the financial results, just some high-level comments which are very proud of the Oil-Dri team. This is a very, very challenging environment. We had to go virtual, we had to figure out what essential meant, and we had people pantry loading cat litter. So, we had to meet incredible demand, and our entire Oil-Dri team did a fantastic job staying healthy both physically and financially and where it enabled us to deliver an incredible quarter in an incredibly challenging time reflecting -- I will tell you those investors who have been with us 2 years or longer, I had to live through this go-live pain of our ERP system when we launched that in August a year and a half ago, almost 2 years ago now, and we would not have made it without that system. If we were on the old system, we would not have made it. We would have been worse than a Lucille Ball skit. And so, I just feel very fortunate. Timing is everything. When this hit, we had the right team in place, we had the right systems in place, we have been far enough and along on the implementation of our S&OP process so that we could get out in front of the demand and more than meet it, we have 99% and 100% fill rates from really all of our major customers and received a lot of kudos from them. So just recognize that your Oil-Dri team really delivered during this quarter. Susan, I will let you go through the specifics from a financial standpoint.

Susan Kreh: Thanks, Dan. And although I usually jump right into the numbers, I want to reiterate a few things that you just said. I mean this was truly an unprecedented quarter and I just want to thank all our team members who enabled us to achieve such good financial results, and we are going to talk about those shortly. As Dan said, in mid-March, our teammates were asked to work from home and they pivoted very quickly. They embraced new technologies practically overnight in order to deliver the results required to support our business and our customers and not just at ordinary levels, but at an all-time record levels of net sales. I particularly want to thank our frontline workers who maintained safe practices to keep each other healthy and we put in a lot of extra effort to keep our customers

shelves stocked. And with that, now I will shift to our financial results. Consolidated net sales for our fiscal third quarter 2020 were as I just mentioned a minute ago an all-time record of \$76.3 million, an 8% increase compared to net sales in the same quarter of fiscal 2019. We saw strength during the quarter in both of the products, where we are focusing our growth investments in those products or cat litter, which is in our retail and wholesale products group, and animal health in our business-to-business products group. In our B2B group, net sales for the third quarter of fiscal year 2020 were \$26.7 million, an increase of 2% over the prior year. And within that, there were some ups and downs by product line, but we were pleased to see that our increased focus on our animal health products is paying off. Net sales of our animal health and nutrition products increased 25% year-over-year during the third quarter driven by increases in many of our markets for our animal feed additives, particularly in Latin America, Mexico, Africa, the Middle East and Asia outside of China. Additionally, net sales within B2B of our agricultural and horticultural chemical carrier products increased 11%, and offsetting this growth was year-over-year decrease of 8% for our fluids purification products where sales were impacted in part due to decreases in edible oil sales resulting from closures of restaurants and schools due to the outbreak of COVID-19. Sales were also unfavourably impacted by the closure of one of our customers bio-diesel processing plant. Now, switching to our retail and wholesale products group. Net sales for the third quarter were \$49.6 million, an increase of 11% over the same quarter in the prior year. This increase was driven by 20% year-over-year growth in net sales of cat litter in both private label and branded litters. In addition to the organic growth that we have been experiencing throughout the previous quarters here in fiscal 2020, as Dan mentioned, we had incremental increases in cat litter sales during the third quarter as customers purchased more cat litter and related products in anticipation of future potential shortages or store closures caused by COVID-19. Also included in our retail and wholesale product groups results were lower sales of our industrial and sports products compared to the third quarter of fiscal year 2019. Sales of industrial and sports products decreased 22% or \$2.2 million, primarily driven by the impact of businesses and sports fields shutting down beginning in March due to COVID-19. If we take look at our consolidated gross profit for the 3 months ended April 30, it was \$21.4 million, which was an increase of 27% over the third quarter of fiscal 2019. This improvement had two primary drivers being a decrease in the costs of freight and natural gas, which on a per-ton basis, which is the way we look at it declined 20% and 32% year-over-year respectively. Thats 20% for freight on a per ton basis and 32% for the natural gas on a per ton basis year-over-year. During the quarter, we did incur some additional employee compensation costs to meet increased customer demand as well as some incremental cleaning and sanitation process due to COVID-19 where reduced costs didnt have a significant impact on our consolidated gross profit, and in fact were basically offset by the reductions we saw in travel and expense in our SG&A expenses for the quarter. Our third quarter income from operations of \$5.7 million is more than double our income from operations of \$2.3 million during the same period in fiscal 2019 driven by the stronger sales and the improved gross profit resulting from lower freight and natural gas. The third quarter net income attributable to Oil-Dri of \$4.6 million compares to \$5.6 million during the third quarter of fiscal 19 and that result included a material one-time benefit of proceeds under a confidential agreement resolving legal proceedings that was included in other income during that period. Net income per diluted common share for the third quarter of fiscal 2020 was \$0.61 compared to \$0.74 in the third quarter of fiscal 2019, which again included the one-time material legal settlement. If you look at it on a 9-month year-to-date basis, net income per diluted common share of \$1.69 for 2020 year-to-date compares to \$1.17 for the same fiscal year-to-date period in fiscal 2019, thats a 44% increase year-over-year. So that is a really good news. And shifting from that good news I would like to highlight two subsequent events that have substantially improved our liquidity during these uncertain times. Because our final debt payment of \$3.1 million is coming due on August 1 of this year we opportunistically amended our note agreement with Prudential details of which you can find in our 10Q that was filed with the SEC this past Monday. Under the amended agreement Oil-Dri issued 10 million in new notes further 10 year tenor and that cash is on hand today. Another event that occurred subsequent to the end of the quarter was that Oil-Dri entered into a confidential agreement to grants and non exclusive perpetual license for 13 million. This amount has been received by Oil-Dri. Both of the aforementioned items

will be included in our fourth quarter financial results which we expect to release on October 13. So all in all really strong quarter and really good liquidity position and with that Dan I will turn it back over to you.

Dan Jaffee: Great thank you Susan obviously great quarter great year to date results so we are going to open up for questions as always ask your most important question first and then go back to the end of the queue and so forth and so on this allows everybody the chance to at least get one question in and maybe two or three so let's please open it up.

Operator: [Operator Instructions] Our first question comes from Ethan Starr. Your line is open.

Ethan Starr: Good morning. Nice quarter considering all the challenges.

Dan Jaffee: Yes thank you.

Ethan Starr: I understand from a recent article that Oil-Dri has improved results in customer service leading to significant improvements in cost reduction. Can you please explain this in more detail and give some idea of maybe how much money has been saved please?

Dan Jaffee: Sure, I am going to let Molly answer the question.

Molly VandenHeuvel: Sure, thanks Dan, so as the article said, we really focused on customer service first, and once we get our customer service levels at benchmark levels, we focused on cost reduction. You really can see that in the operating income, so I think the amount becomes evident in our cost of goods sold. We focused on a lot of the basics making things more efficient, making sure we have the right pricing and contracts in place, and just overall operational effectiveness.

Dan Jaffee: Yes. And I'll add because she doesn't want to toot her own horn, but Molly has led our team, and increased quality leads to lower costs. It always does because you have less rework. You have greater customer satisfaction which leads to more sales which also can lower your costs if you could get to spread your fixed costs over a greater base, so the results you are seeing, you are seeing the topline growth led by our general managers and then you are seeing the bottom line growth led by our supply chain who's just functioning very, very well. So, it's been a win-win, but thank you for your question.

Ethan Starr: Okay, I will get back in the queue. Thanks.

Dan Jaffee: Thanks

Operator: Our next question comes from John Bair of Ascend Wealth Advisors. Your line is open.

John Bair: Thank you. Yes, I will echo nice quarter. I also appreciate more detailed information in the press release and commend you for the aspect of helping the frontline workers, so I'm wondering to what sense or to what degree do you have a sense that those extra orders and so forth in March and April may have pulled forward orders and kind of parallel to that is, did you see an increase in orders through online vendors like a Chewy or somewhere like that that you think might continue on in the future. Thanks.

Dan Jaffee: Alright. Well John, thank you; and Jessica who runs our Consumer Products division, she will answer that question.

Jessica Moskowitz: Hi good morning.

John Bair: Good morning.

Jessica Moskowitz: So in terms of the shift from Q3 -- end of Q3 from Q4, if we look at kind of the first half, I would say we were benchmarked up 10% and then Q3 was up around closer to 19%. So I would say that 9% to 10% difference is likely what shifted up into Q3 from Q4, and year-to-date still we had had strong sales year to date, so we could expect that to continue, but again the shift would likely be 9% to 10%. In terms of Chewy and Amazon, we continue to focus on e-commerce. We have seen strong sales in e-commerce as a result of COVID, and we do expect it to continue going forward. To what extent, I don't think we can disclose that at this point, but it is something we are investing in and obviously have seen a consumer momentum into those channels, so intend to continue to focus on it.

John Bair: Well, I would think that sometimes when you are forced into a habit change that you weren't expecting to do like this pandemic has caused a lot of people whether it's work or shopping or whatever that you might have a sense of whether those trends are continuing, in other words as online ordering as opposed to going into a store to pick it up. And just curious whether you are seeing some momentum continue in that area?

Jessica Moskowitz: Yes. I think the macro level trends would suggest that consumers are shifting into those channels, and our results would continue to support that kind of ongoing shift into e-com.

John Bair: Alright. Great, thank you. I will get back in queue.

Dan Jaffee: Thank you, John.

Operator: Our next question comes from Robert Smith of Center for Performance. Your line is open.

Robert Smith: Yes, good morning and congratulations on the good quarter and also on the dividend increase. My question focuses on animal health. Its a two-part question. First part Dan if animal health were a standalone unit, at what revenue level would it become profitable, at what kind of a ballpark number?

Dan Jaffee: I mean, I will take that, and I can tell you I dont know. So I have no idea. I don't look at it that way. I mean everything is so comingled . I mean if it was a standalone, it couldnt be, I mean, the products are co-generated and we make money by being in all aspects of our business. I once used this analogy to Clorox. We have been supplying them Fresh Step for over 40 years now. And I said, we cant just sell the fillet to make money. We have to sell the [indiscernible] the tail, the whole thing and then the fillet becomes very profitable, but if all you did was sell the fillet, you wouldnt make any money in the steak business at Oil-Dri. So honestly, its an unanswerable question. So, I don't know.

Robert Smith: The second part of the question is with COVID entering the Southern Hemisphere more vigorously now, how do you see the animal health possibilities in the remaining quarter of the year?

Dan Jaffee: Great. Fleming, who is the President of that division will field that question.

Fleming Mods: Yes. Hi, great question. Thank you. Yes. So right now, we are continuing to see good support for our products globally in all the regions. The COVID-19 is of course with the shutdown affecting more of the industrial kitchens, the schools are starting to come back. So, at this point, we are not anticipating any huge impact to the business.

Robert Smith: Well, I meant in particular, in the animal health, I think Brazil is an important market, so they are having quite a bit of difficulty down there.

Fleming Mods: Yes. I mean, they are seeing challenges, but they are also seeing additional improved exports into China. So, theres different resources and sources pulling in different directions that we would normally see.

Robert Smith: Okay, thank you. I will get back in the queue.

Dan Jaffee: And before we go for the next question, I will just I think we are all recognizing that there is a new world now, we used to divide equities into large cap and small cap and micro cap and growth versus value and all these different cuts. I think what we are going to be seeing forever more now is essential versus non-essential. No one ever thought about this before COVID, but boy, am I better to be lucky than good. I feel very sorry for fellow CEOs and businesses who were run every bit as well as Oil-Dri, but there are non-essential businesses and nothing they can do about it. And so most of our sales are essential either tied to food production or animal, care of pets and so we got very lucky there. So all I can tell you that Bob is that yes, there is COVID going on throughout the world, but you may or may not die from COVID, you will die if you dont eat. So, food is absolutely essential and we feel very lucky and fortunate that our core businesses are tied to things that will were not and really wont be impacted too much by global pandemic. Next question?

Operator: Our next question is a follow-up from Ethan Starr. Your line is open.

Ethan Starr: Yes. First, the quick question is the license fee, is that related to cat litter?

Dan Jaffee: Laura, you wanted to answer?

Laura Scheland: I am sorry, we cant provide any detail.

Ethan Starr: Okay, thats fine. My meal question is this, what are your biggest challenges in getting more consumers to purchase lightweight scoopable litter or the lowest cost litter, which is Cats Pride among the major brands and how do you plan to address those challenges?

Dan Jaffee: I mean, Jessica, I am happy to you may answer it, I am happy to jump in, give me a thumbs up or thumbs down, she is not on the screen. If you want to take it give me a thumbs up. Okay, Jessica is going to field that.

Jessica Moskowitz: Yes. I can take it, Dan and feel free to build on it. As you would like, I think I wouldnt really view it as a challenge. I think its an opportunity. Lightweight cat litter continue to we continue to improve quality, while lowering the price and I think that our brands and our private label and branded is well positioned as a value player to really continue to dominate the market. So, I feel optimistic and our goal is to continue to delight our consumers

and that's through offering them great products at a great price. And as we kind of increase the quality and are able to lower prices or offer at the same value, that value continues to delight consumers even more.

Dan Jaffee: Well said. And I think the only thing I would add is we all we are in a recession. They have announced that and actually they think now it started before COVID even hit and COVID put it into hyper gear, but clearly, when you are in that position, people will be more incentivized to save money in their weekly or monthly grocery bill. And so competing in the value segment whether its through private label or through OPP opening price point quality brands is a good place to be. So, having cats lightweight scoopable at \$5.98 at Walmart is a good place to be. Being the OPP well Lehmann market is a bad example because they want bankrupt. But you get the idea, you don't want to be the cheap guy at a premium retailer, but clearly, when its price related and we are leaning heavily into and Jessica's team was doing a phenomenal job, Walmart, Target, Albertsons, Safeway, Dollar General, Family Dollar, the e-commerce, all these areas where people can shop and compare on price, we believe we are well positioned to supply them with high-quality products at affordable prices.

Ethan Starr: Okay, are you taking share?

Dan Jaffee: Well, you can see our growth. I mean we are growing faster than the category. So by definition we are taking share.

Ethan Starr: Okay, I will get back in queue.

Operator: Our next question is a follow-up from John Bair. Your line is open.

John Bair: Thank you. Kind of a simple question, what is the most exciting part of your business, I think part of the business that you are most excited about and offers what you think will be best opportunities going forward and conversely, where do you see the biggest challenges right now? Thanks.

Dan Jaffee: But that's like you asked me which of my children do I love the best, I mean that's a telephone, but okay.

John Bair: Simple question, yes.

Dan Jaffee: I love our businesses that are around creating value from sorbent minerals, how about that one? No, I am excited about all of our businesses. We have got great people in place. We have got great strategy and we play various ones, I mean, you said which ones are we targeting for the highest growth, well that's clearly on a cat litter side, its with our high-quality lightweight litters and then its on the animal health side, where as you well know, there is a global push away from antibiotics in the human food chain. So they are pulling them out of raising the animals. And so now they have a real unmet need they need to figure out how to maximize production and we have a fantastic antibiotic free solution. So we are very excited about those in terms of the growth potential. But I will tell you very excited and thankful for our Oil-Dri floor absorbent, industrial absorbents business, our agricultural business, our fluids purification business, these are all solid differences where we have a real reason to be and a real reason to compete and are doing well. So excited about all of them biggest challenge that is staying disciplined I would say we all when we started doing really well and we have done really well for a long time but you can see the momentum is gaining is staying disciplined not just chasing every hand but knowing when to fold for those of you who play poker so we are going to stay disciplined we are going to focus our resources on the businesses we think can grow most rapidly and those are the consumer products and the animal health and we will support the other division to but they are fighting a different war and so I would say that's our challenge is just as a team to stay as hungry as you are when things are going poorly as you need to be when things are going well.

John Bair: Okay, okay I will get back to the queue.

Dan Jaffee: Okay

Operator: Our next question is a follow up from Robert Smith. Your line is open.

Robert Smith: Hi Dan. Can you give us some additional color on your approach to digital advertising and within the context of the advertising and promotion budget?

Dan Jaffee: Sure. And Jessica take that one?

Jessica Moskowitz: Sure. So our approach is that going to lead to be hyper targeted I think digital offers us an opportunity to do that and to really understand who are consumers and reaching our consumers where she already is where testing we continue to do kind of the main digital advertising social search kind of typical digital tactics that has been also testing new tactics like video

Robert Smith: Here is an idea for you carry on carry less.

Jessica Moskowitz: Got it.

Operator: Our next question comes from Ethan Starr. Your line is open.

Ethan Starr: Yes I am wondering I guess my impression is that Oil-Dri saved much more money due to COVID on like things lower energy prices and travel expenses that you actually spent extra for COVID and I am just if you can expand on that and may be go into some number of possible and also to what extent when energy prices drop does that benefit Oil-Dri or does it benefit your customers or how does that where does that when did the shift prices stopped to reflect that thank you.

Ethan Starr: Yes. And Molly, just so that I can earn my keep, I'll do the play by play. You can do the color if I need anything else but I'm proud of the fact that yes we save some money on travel and entertainment as Susan said but during

Operator: We have 2 minutes remaining for this conference.

Dan Jaffee: Okay, thank you. Yes, during the peak of the demand when we were pushing our people really hard we gave the equivalent of a \$2 an hour shift premium to all of our frontline workers and that more than equips the savings now when things started to dive down we then went to \$1 an hour shift premium and that was a break even to a slight benefit none of this was material to the bottom line so I'm not really going to give you the dollars but I'm just telling you we took the money away that we were saving and gave it to those people that were what I called heavy combat pay and so proud that we did that I now don't remember the second part of your question

Ethan Starr: Well just to what extent this energy prices drop how much is the benefit for Oil-Dri and what extent does it shift at some point you said that was customers.

Dan Jaffee: Okay Molly will take this one although I can take the pricing part because I was telling you to the board yesterday look we are in a rational industry and when gas prices go crazy one way or the other if our competition moves we move with them and vice versa so we are in a rational business but Molly I will let you put some more

Molly VandenHeuvel: Right and it is not a straight answer for energy and fuel prices for transportation we do see immediate savings where we pay for freight as part of our product costs as part of our fuel surcharges part of our contracts typically that's aligned with the fuel index. And sometimes, we get that when we pay for it, but we have a lot of customers who pay for freight and then they would see the benefit of that for natural gas we have seen some benefit in the third quarter so that was pretty immediate and then for other material costs savings there is a bit of a lag that's really due to commodity usage and lagged timing so it is not really straight forward answer some of immediate some is longer term and the savings that we saw in the third quarter are more than just commodity we really have built a good processes in place to drive sustainable savings within the operation

Ethan Starr: Okay, great. Appreciate that. And I'm assuming that the lag is probably mostly in the resin price plastic packaging package prices. Well thank you very much again.

Molly VandenHeuvel: [indiscernible]

Ethan Starr: Yes, thanks for the dividend increase also Dan.

Dan Jaffee: Great. And I hope you also spotted the share repurchase which again was a way of trying to deliver value back to our shareholders and we like we saw we are opportunistic so that we had a chance to buyback shares with higher dividend and our cash is earning 0.2% so and we have plenty of opportunity that we can deploy that cash on but we are going to still stay opportunistic on the share repurchase program so hopefully spotted that in the queue as well thank you guys its been a half an hour and we look forward to talking to you it will be our end of the year that will be our fourth quarter and fiscal year end I would tell you on one of our major metrics which is we look at pre tax pre bonus income and then we divvy up that income between the shareholders and then the teammates. We already, through nine months have made more than we have ever made and any fiscal year we have ever had so as I joke we are playing with a house of money at this point but we are having a record year it feels really good it feels really good that all the investments we have made in the last two and half years on people and infrastructure are why this is happening and its just it feels good it is predictable and we are up in front of our businesses and we just we have always taken a long term approach and the seeds we planted years ago are sprouting now and seeds we plant today will start in a few years so thank you for the long-time holders and we are happy to reward you with I think will results 17 year in a row

of dividend increases which is fantastic. So thanks everybody. We will talk to you again in a quarter.

Operator: Ladies and gentlemen, this concludes todays conference call. Thank you for participating. You may now disconnect. Everyone have a great day."""

Answer:

<|im_end|>

<|im_start|>assistant