

Compute

Institutional Equity Research
December 1, 2025

DeepSeek Continues to Prove That Post-Training Just Works

Earlier this morning, DeepSeek released DeepSeek-V3.2 and DeepSeek-V3.2-Speciale, which are incremental updates to DeepSeek's previous V3.2-Experimental model from a few months ago. While we would still characterize this as an incremental release, there is plenty in the models and technical paper to get excited about, which only heightens our anticipation for a DeepSeek-4 release in the near future.

DeepSeek's V3.2 release is a reminder that sometimes old is gold and all you really need is more post-training. The model keeps the same backbone as V3.2-Exp and V3.1-Terminus, with the only architectural change being the addition of DeepSeek Sparse Attention (DSA). That being said, the majority of the focus was not on the base model itself but what was done with it in post-training. DeepSeek takes the existing base and scales reinforcement learning to a regime where post-training compute now exceeds 10% of the pre-training budget compared to DeepSeek-R1 where RL compute was around 1.5% of pre-training, a relatively large jump that is enough to pull the models from "strong open base model" into GPT-5 and Sonnet 4.5 territory across a wide range of reasoning and agentic benchmarks.

Crucially, DeepSeek is doing this without inventing a new objective. Group Relative Policy Optimization (GRPO) stays in place and the gains come from what we would characterize as a focused set of fixes that make RL scale in practice. This includes an unbiased KL estimator that corrects the original K3-style estimate so gradients are not biased toward extreme low probability tokens and off-policy sequence masking so negative, highly off-policy rollouts do not destabilize training. DeepSeek also locks MoE routing paths between sampling and optimization and reuses the same top-p masks across old and current policies so importance sampling does not break on a moving action space. In parallel, DeepSeek Sparse Attention wraps MLA with a tiny FP8 indexer that picks a few thousand relevant key-value entries per token, which makes attention behave closer to linear than quadratic at 128k context while preserving the base model's behavior. Once you have this stack, it becomes rational to embrace token inefficiency rather than fight it, which is exactly what we see V3.2-Speciale do with a softer length penalty and heavier reasoning RL, including DeepSeekMath-V2 prover rewards ([see our breakdown of Math-V2](#)). Speciale spends 2-3x more output tokens than rival models yet remains cheaper end-to-end and reaches Gemini 3 Pro reasoning levels while also hitting gold medal performance on the 2025 IMO, CMO, IOI, and ICPC finals without introducing a new architecture.

These models demonstrate that agentic capabilities may depend more on environment design than on model class. DeepSeek builds a synthetic universe of about 1,800 task-oriented environments and roughly 85,000 prompts that cover code agents, search agents, a Jupyter-style interpreter, and general agents. Many of these tasks are real-tool environments with synthetic prompts mined or generated from the web, while others are fully synthesized but easy to verify. RL on this synthetic corpus alone is enough to move benchmark scores on V3.2 in ways that look more like environment quality and coverage than like a mysterious agent-architecture edge. The lesson is that if you can build large, diverse, verifiable environments, you can push the RL stack a very long way.

We need more compute, not less. The technical paper frames the open-source deficit in complex tasks primarily as a compute problem rather than a gap in ideas or data. DeepSeek argues that open models have been hamstrung by inefficient attention for long contexts, underinvestment in post-training compute, and a lack of rich, scalable agent environments. V3.2 largely addresses the latter two and DSA handles the first, but the compute gap still remains despite these improvements. Frontier labs no longer have a structural data or knowledge advantage. What they have are larger, more continuous compute budgets that let them run this kind of post-training at even greater scale. In DeepSeek's own framing, the bottleneck is no longer algorithmic ingenuity, but rather who can afford to run enough post-training on top of a strong base.

INDUSTRY UPDATE

Price (12/1/25)

Industry:

TECHNOLOGY

Alexander Platt

(503) 603-3045

AJPlatt@dadco.com

DaVinci Overview

D.A. Davidson's DaVinci initiative focuses our technical-oriented research, data-driven insights, and prescient think pieces under one unified framework. We note that for our DaVinci coverage of deep tech businesses, we employ an early-stage venture approach focusing on technical foundations, disruptive potential, and long-term strategic value, rather than near-term financial and valuation metrics given the unique growth trajectories of pre-inflection markets.

This report is intended for AJPlatt@dadco.com. Unauthorized distribution prohibited.



Copyright D.A. Davidson & Co., 2025. All rights reserved.

Potential Risks

Required Disclosures

Best-of-Breed: Expected to outperform on a risk adjusted basis over a five-year time horizon, but may be fully valued over a 12-18 month time horizon.

D.A. Davidson & Co. makes a market in Apple Inc., Amazon.com, Inc., CoreWeave, Inc., Alphabet Inc., Microsoft Corporation, Nebius Group N.V., NVIDIA Corporation and Oracle Corporation.

D.A. Davidson & Co. or any of its affiliates, does or seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.

D.A. Davidson & Co. is a full service investment firm that provides both brokerage and investment banking services. Alexander Platt, the research analyst principally responsible for the preparation of this report has received and is eligible to receive compensation, including bonus compensation, based on D.A. Davidson's overall operating revenues, including revenues generated by its investment banking and institutional equities activities. D.A. Davidson & Co.'s analysts, however, are not directly compensated for involvement in specific investment banking transactions.

I, Alexander Platt, attest that (i) all the views expressed in this research report accurately reflect my personal views about the common stock of the subject company, and (ii) no part of my compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views expressed in this report.

Rating Information

D.A. Davidson & Co.'s Institutional Research Rating Scale Definitions (maintained since October 10, 2017); information regarding our previous definitions is available upon request:

BUY: Expected to produce a total return of over 15% on a risk adjusted basis over the next 12-18 months

NEUTRAL: Expected to produce a total return of -15% to +15% on a risk adjusted basis over the next 12-18 months

UNDERPERFORM: Expected to lose value of over 15% on a risk adjusted basis over the next 12-18 months

Rating Distribution (as of 9/30/25)	Coverage Universe Distribution			Investment Banking Distribution		
	IR	WMR	Combined	IR	WMR	Combined
BUY (Buy)	59%	85%	62%	8%	0%	7%
NEUTRAL (Hold)	40%	13%	37%	4%	0%	3%
UNDERPERFORM (Sell)	1%	2%	1%	0%	0%	0%

IR denotes Institutional Research; WMR denotes Wealth Management Research whose rating scale is Buy/Add, Neutral, Sell/Reduce. Investment Banking Distribution denotes companies from whom D.A. Davidson & Co. has received compensation in the last 12 months. Best-of-Breed: Expected to outperform on a risk adjusted basis over a five-year time horizon.

Target prices are our Institutional Research Department's evaluation of price potential over the next 12 months, based upon our assessment of future earnings and cash flow, comparable company valuations, growth prospects and other financial criteria. Certain risks may impede achievement of these price targets including, but not limited to, broader market and macroeconomic fluctuations and unforeseen changes in the subject company's fundamentals or business trends.

While the Best-of-Breed designation does not contain a separate rating and/or price target from that of the standard ratings system referenced above, the expectation is that the security, based on the 12 criteria utilized in assessing the "Best-of-Breed" designation, will outperform over a five-year time horizon, not the standard 12-18 month time horizon.

For a copy of the most recent reports containing all required disclosure information for covered companies referenced in this report, please contact your D.A. Davidson & Co. representative or call 1-800-755-7848.

Other Disclosures

Information contained herein has been obtained by sources we consider reliable, but is not guaranteed and we are not soliciting any action based upon it. Any opinions expressed are based on our interpretation of data available to us at the time of the original publication of the report. These opinions are subject to change at any time without notice. Investors must bear in mind that inherent in investments are the risks of fluctuating prices and the uncertainties of dividends, rates of return and yield. Investors should also remember that past performance is not necessarily an indicator of future performance and D.A. Davidson & Co. makes no guarantee, express or implied, as to future performance. Investors should note this report was prepared by D.A. Davidson & Co.'s Institutional Research Department for distribution to D.A. Davidson & Co.'s institutional investor clients and assumes a certain level of investment sophistication on the part of the recipient. Readers, who are not institutional investors or other market professionals, should seek the advice of their individual investment advisor for an explanation of this report's contents, and should always seek such advisor's advice before making any investment decisions. Consensus estimates are obtained from Capital IQ. Further information and elaboration will be furnished upon request.

Other Companies Mentioned in this Report

Company Name	Ticker	Rating	Price
Apple Inc.	AAPL	NEUTRAL	\$278.85

This report is intended for AJPlatt@dacdo.com. Unauthorized distribution prohibited.



Company Name	Ticker	Rating	Price
Amazon.com, Inc.	AMZN	BUY	\$233.22
CoreWeave, Inc.	CRWV	UNDERPERFORM	\$73.12
Nebius Group N.V.	NBIS	BUY	\$94.87
Alphabet Inc.	GOOGL	NEUTRAL	\$320.18
Oracle Corporation	ORCL	NEUTRAL	\$201.95
Meta Platforms, Inc.	META	BUY	\$647.95
Microsoft Corporation	MSFT	BUY	\$492.01
NVIDIA Corporation	NVDA	BUY	\$177.00