Dynamic Model of Circular Financing in the AI Compute Market: Stability with Neobanks in the Loop

The Case of Revolut in the NVIDIA Ecosystem

Felix Desyatirikov

ITMO University, St. Petersburg, Russia

felixdesyatirikov@gmail.com

Abstract. Neobanks have become a new kind of participant in the AI compute market: they own large behavioural datasets and train proprietary models, yet own no hardware and increasingly receive capital from the chip vendor itself. We ask how such a participant changes the stability of the market. The market is modelled in control-theory terms as a ninth-order block diagram linking NVIDIA with hyperscalers (Microsoft, Meta, Alphabet, Amazon, Oracle), AI labs connected to NVIDIA by cross-investments (OpenAI, Anthropic, xAI, Mistral), GPU-rental neoclouds (CoreWeave, Nebius, Nscale) and end consumers of compute, with Revolut as the neobank case. The model contains a positive feedback loop (circular vendor financing) and a negative one (substitution by in-house chips). Its fast market link is identified from NVIDIA's share-price decline of October 6 – November 10, 2026 (−29.4%), driven by an October background shift, the November 3 US midterm election and a hypothetical risk specific to NVIDIA (from November 9): $T = 0.70$ days, $\xi = 0.40$, RMS error 1.8 pp, $R^2 = 0.96$. A coupling parameter θ describes how much of the vendor's financing reaches the neobank channel. With $\theta = 0$, Revolut's demand is paid for by its own customers: its pole stays outside the loops and the stability margin with respect to circular financing is 8.8. In the hypothetical scenario $\theta = 1$, the neobank becomes part of the positive loop: the critical financing gain falls from 1.31 to 0.79, the margin from 8.8 to 5.2 (−40%), shock amplification rises by 9% and the response slows by 22%. Only the four lowest-order coefficients of the characteristic polynomial change, so the coupling affects the slow fundamental dynamics while the market reaction itself remains intact. Data on Revolut published by Pavel Nesterov add a second effect: its blended LLM inference cost fell from \$1.49 to \$0.49 per million tokens between April and November 2026, and its AI coding inference costs fell by about 90% after it moved from commercial APIs to open-weight models self-hosted on 56 NVIDIA B200 GPUs. Treated as a slow drift of the neobank channel gain, falling token cost leaves the margin unchanged when compute is customer-paid; when compute is vendor-financed, the margin ranges from 7.2 to 2.9 depending on the price elasticity of token demand. Neobanks thus act either as a stabilising feedforward channel or as an additional amplifier, depending on who pays for their compute. The bank's own compute channel remains stable for any coupling, but vendor financing transmits NVIDIA-specific shocks to it: together, the election shock and the hypothetical risk cut its compute purchases by about 3.7%, with the effect building up over about half a year.

Keywords: neobanks; Revolut; AI compute market; stability margin; positive feedback; circular financing; characteristic equation; root locus; block diagram; system identification; token cost; price elasticity; parametric drift; robust stability; NVIDIA.

1. Introduction

Until recently the demand for AI hardware was described by a few groups: hyperscalers that buy GPUs, AI labs that consume compute, and neoclouds that rent it out. A neobank such as Revolut does not fit any of them. It holds data on about 26 million users in 111 countries and has trained its own foundation model PRAGMA, so in one sense it competes with AI labs. It owns no data centres of its own and rents compute from a neocloud (Nebius), yet since April 2026 it has also served open-weight coding models on 56 NVIDIA B200 GPUs, cutting its AI coding inference costs by about 90%. It is private, so market shocks do not move its valuation daily. And the chip

vendor is among its shareholders: NVIDIA's venture arm NVentures joined the round valuing Revolut at $75 bn.

This combination raises a control question: does such a participant stabilise the AI compute market or destabilise it? The answer depends on the structure of the feedback loops, so we build an explicit block diagram, identify its fast part from real data and then compare the closed-loop system with and without the neobank inside the vendor-financing loop.

Describing economic systems with block diagrams and transfer functions has a long tradition. Tustin (1953) represented macroeconomic models as block diagrams borrowed from electrical-circuit and servomechanism theory, and Phillips (1954) used the same tools to analyse stabilisation policy; Neck (2009) reviews the subsequent development. Block-diagram models of macroeconomic dynamics are still in use (Andresen), and transfer-function models with two feedback loops are a standard tool for analysing the stability of vendor–customer supply chains (Disney and Towill, 2002). This paper applies the same approach to the AI compute market.

The paper makes five contributions:

1. A block-diagram and state-space model of the AI compute market around NVIDIA, with an explicit place for the neobank channel (Sections 3–5).
2. Identification of the market dynamics from NVIDIA's October–November 2026 share-price decline, used as a natural step-response experiment (Section 6).
3. A quantitative analysis of how the stability margin, the roots of the characteristic equation and shock amplification change as the neobank is drawn into the circular-financing loop (Sections 7.1–7.6).
4. An analysis of falling token costs at Revolut, as reported by Pavel Nesterov, as a parametric drift of the neobank channel gain and of its effect on the stability margin (Section 7.7).
5. An analysis of the stability of the bank's own compute channel and of how vendor financing transmits NVIDIA-specific shocks to it (Section 7.8).

2. Data: the 2026 NVIDIA decline as an identification experiment

Between October 6 and November 10, 2026 (25 trading days), NVIDIA's share price fell from $239.24 to $168.89 (−29.4%, roughly −$1.7 trillion of market capitalisation). The output variable is the logarithmic deviation from the peak:

$$y(t) = 100 \ln \frac{P(t)}{P_0}, \qquad P_0 = 239.24 \text{ \$ (close of Oct 6, 2026)}, \qquad [y] = \%$$

Time t is measured in trading days: t = 0 is the October 6 close, and the session of day n runs from t = n − 1 to t = n.

Date (2026)	t, d	Event	Close, $	Daily change
Nov 10	25	Continued reaction to the hypothetical risk f_2	168.89	−5.5%
Nov 9	24	Hypothetical risk f_2, specific to NVIDIA; the market rose (Nasdaq +2.74%)	178.67	−4.27%
Nov 6	23	Partial rebound of chipmakers	186.64	+2.78%
Nov 5	22	AI-sector sell-off after the midterm elections took place: Nasdaq −5.52%, SOX −12%	181.60	−9.26%
Nov 4	21	First reaction to the election results (Nasdaq +1.44%, NVDA down)	200.13	−3.17%
Nov 3	20	US midterm election day	206.69	+2.29%
Oct 30	18	Weak Intel outlook, semiconductor sell-off	204.64	−4.22%
Oct 22	12	Doubts about AI spending returns and the cost of capital	212.21	−4.40%
Oct 6	0	Record close (peak)	239.24	—

The calculations use 22 closes from October 5 to November 10; four of them (October 12, 21, 26 and 28) were reconstructed from the reported change on the following day. NVIDIA's fiscal Q3 2027 report (November 17) is

outside the sample.

Modelling assumptions. (1) All signals are small deviations from the early-October steady state, so the links are linear. (2) Fundamental links have time constants of weeks to quarters, while the market reacts within 1–2 days; over five weeks the slow loops are almost frozen. (3) Disturbances are steps with unknown amplitudes. (4) Besides the two specified disturbances (the election and the hypothetical risk), a background disturbance f_0 is needed: the price had already fallen 15% by November 2 amid doubts about AI capex returns, neocloud debt and the Intel outlook. (5) "AI-model investors" are model developers linked to NVIDIA by cross-investments: OpenAI, Anthropic, xAI (SpaceX) and Mistral.

3. Block diagram of the AI compute market

NVIDIA's price is the output of a closed-loop system. Four stakeholder groups generate orders, orders shape revenue expectations, and expectations drive the market valuation; through two feedback loops the valuation returns to the stakeholders (Figure 1).

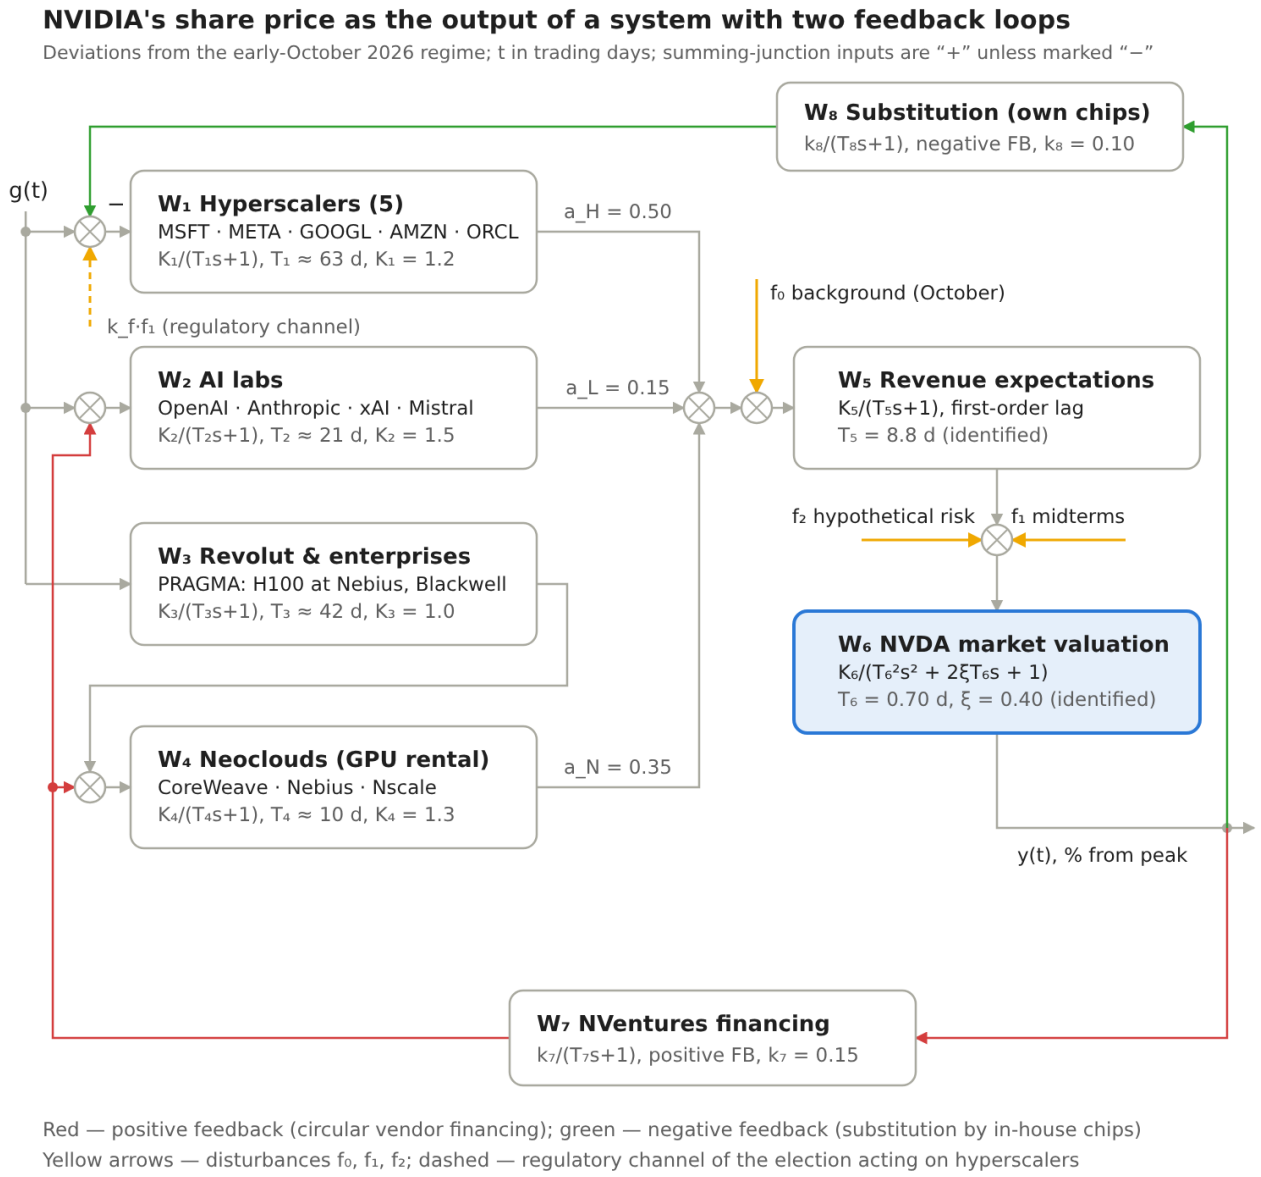


Figure 1. Block diagram of the AI compute market around NVIDIA. The neobank channel is $W_3 \to W_4$

Link	Companies	Relationship with NVIDIA	Order weight
W_1 Hyperscalers, the five largest chip buyers	Microsoft, Meta, Alphabet, Amazon, Oracle	About half of data-center revenue; in fiscal Q1 2027, $38 bn of $75.2 bn	a_H = 0.50
W_2 AI labs, the investment loop	OpenAI, Anthropic, xAI (SpaceX), Mistral	NVIDIA is a shareholder ($30 bn in OpenAI) and a compute supplier	a_L = 0.15
W_3 Neobanks and enterprises	Revolut, enterprise clients	Revolut trains PRAGMA on H100 GPUs in Nebius AI Cloud and scales it on Blackwell Ultra; NVentures is a shareholder	via W_4
W_4 Neoclouds	CoreWeave, Nebius, Nscale	Rent out capacity built on NVIDIA GPUs; NVIDIA holds stakes and provides guarantees (up to $6.3 bn of unused CoreWeave capacity)	a_N = 0.35

The weights a are simplified revenue shares; NVIDIA's filings show its clients only as "Customer A, B, C…".

Two feedback loops.

1. **Positive feedback (W_7), circular vendor financing.** A high share price gives NVIDIA money to invest in its clients; they buy NVIDIA GPUs, orders grow and the price rises further. On the way down the loop works in the same direction.
2. **Negative feedback (W_8), substitution.** Expensive GPUs and a gross margin of about 75% push hyperscalers towards Google TPU, Amazon Trainium, Meta MTIA, Microsoft Maia. The loop is slow ($T_8 \approx$ one year) and stabilising.

4. The neobank in the system

Revolut occupies five positions in the diagram at once.

Role	Fact	Place in the block diagram
Compute consumer and developer of its own model	PRAGMA trained on data from about 26 million users, 207 billion tokens; in production it serves 80 million customers on Blackwell Ultra	Link W_3, driven by the reference input g
Client of a hardware holder	Training ran in Nebius AI Cloud; Nebius is a neocloud in which NVIDIA has also invested	Series connection W_3W_4
Vendor investee	NVentures joined the round valuing Revolut at $75 bn; Revolut is part of NVIDIA's £2 bn UK programme	Potential entry of the financing signal x_8 into W_3
Alternative to AI labs	Instead of OpenAI or Anthropic APIs the bank builds its own model with NVIDIA's open tools	Branch parallel to W_2 at the order junction Σ_E
Self-hoster of open-weight models	Since April 2026 it serves GLM coding models to about 2,000 engineers on 56 NVIDIA B200 GPUs, at about a tenth of the cost of Claude Code or Codex	Direct purchase of NVIDIA hardware; signal moves from W_2 to W_3

Revolut's unit economics are public too. According to a post by Pavel Nesterov (November 2026), its blended inference cost across all LLM use cases (tokens bought from external vendors, self-hosted GPUs including idle time, tool costs of agentic workflows; training excluded) fell from $1.49 per million tokens in April 2026 to $0.49 in November, three times cheaper in seven months, while about half of its tokens were still generated by external proprietary models. This is a different metric from the 90% figure, which compares coding workloads before and after the move from commercial APIs. Section 7.7 uses both. Turning such public data into a block diagram with transfer functions follows the approach of Desyatirikov (2024), where a non-technical system — the hype cycle of the Java ecosystem — is modelled in control-theory terms.

The coupling parameter θ. The third role is the one that matters for stability. We introduce $\theta \in [0, 1]$, the share of the vendor-financing signal x_8 that reaches the neobank channel:

- $\theta = 0$ — the neobank's compute is paid for by its own customers; the channel W_3W_4 is a pure feedforward path from final demand g to NVIDIA orders.
- $\theta = 1$ — hypothetical scenario: vendor money invested in the neobank is fully converted into compute purchases; the neobank becomes part of the positive loop.

In equations this changes only the neobank link input: $T_3\dot{x}_3 = -x_3 + K_3(g + \theta x_8)$. Analysts point to the same distinction in plain words: demand financed by NVIDIA is worth less than demand from a paying end customer.

The $\theta = 1$ case is a hypothetical scenario. The paper does not claim that NVIDIA's investment in Revolut is used to pay for the bank's compute: the NVentures stake is an equity investment, and no such arrangement is publicly known. $\theta > 0$ is used only to show how the system would behave if vendor money did turn into compute purchases.

5. Elementary links and state-space model

The system consists of seven first-order lags, one oscillatory link and static summing junctions; its order is $n = 9$. Parameters of W_5 and W_6 are identified in Section 6; the others are expert estimates.

Link	Transfer function	Type	Parameters (t in trading days)	Rationale
W_1 Hyperscalers	$K_1/(T_1s + 1)$	First-order lag	$K_1 = 1.2$; $T_1 = 63$	Quarterly capex budgets; accelerator effect $K > 1$
W_2 AI labs	$K_2/(T_2s + 1)$	First-order lag	$K_2 = 1.5$; $T_2 = 21$	Purchases follow funding rounds; leverage
W_3 Neobanks and enterprises	$K_3/(T_3s + 1)$	First-order lag	$K_3 = 1.0$; $T_3 = 42$	Corporate IT budgets are inert
W_4 Neoclouds	$K_4/(T_4s + 1)$	First-order lag (with dead time in a finer model)	$K_4 = 1.3$; $T_4 = 10$	Fast switching of purchases
Σ_E NVIDIA orders	a_H, a_L, a_N	Static gains and summing junction	0.50; 0.15; 0.35	Data-center revenue shares
W_5 Revenue expectations	$K_5/(T_5s + 1)$	First-order lag	$K_5 = 1.0$; **$T_5 = 8.8 \pm 2.9$**	Analysts revise forecasts within 1–2 weeks
W_6 Market valuation	$K_6/(T_6^2s^2 + 2\xi T_6s + 1)$	Oscillatory	$K_6 = 1.2$; **$T_6 = 0.70 \pm 0.13$; $\xi = 0.40$**	Overreaction and partial reversal
W_7 Financing (positive FB)	$k_7/(T_7s + 1)$	First-order lag in a positive feedback path	$k_7 = 0.15$; $T_7 = 63$	NVentures deals take about a quarter
W_8 Substitution (negative FB)	$k_8/(T_8s + 1)$	First-order lag in a negative feedback path	$k_8 = 0.10$; $T_8 = 250$	In-house chip cycle of a year or more

Bold values are identified (± is the least-squares standard error). The time constants span more than two orders of magnitude, so the system is stiff.

$$
\begin{aligned}
T_1\dot{x}_1 &= -x_1 + K_1\,(g - x_9 + k_f f_1) && \text{hyperscalers}\\
T_2\dot{x}_2 &= -x_2 + K_2\,(g + x_8) && \text{AI labs}\\
T_3\dot{x}_3 &= -x_3 + K_3\,(g + \theta\, x_8) && \text{neobank channel}\\
T_4\dot{x}_4 &= -x_4 + K_4\,(x_3 + x_8) && \text{neoclouds}\\
T_5\dot{x}_5 &= -x_5 + K_5\,(a_H x_1 + a_L x_2 + a_N x_4 + f_0) && \text{expectations}\\
T_6^2\ddot{y} + 2\xi T_6\dot{y} + y &= K_6\,x_5 + f_1 + f_2 && \text{market}\\
T_7\dot{x}_8 &= -x_8 + k_7\,y && \text{financing (positive FB)}\\
T_8\dot{x}_9 &= -x_9 + k_8\,y && \text{substitution (negative FB)}
\end{aligned}
$$

With $x_6 = y$, $x_7 = \dot{y}$ the state-space form is $\dot{x} = Ax + Bg + Ef$, $y = Cx$, where

$$A = \begin{pmatrix} -\frac{1}{T_1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{K_1}{T_1} \\ 0 & -\frac{1}{T_2} & 0 & 0 & 0 & 0 & 0 & \frac{K_2}{T_2} & 0 \\ 0 & 0 & -\frac{1}{T_3} & 0 & 0 & 0 & 0 & \frac{\theta K_3}{T_3} & 0 \\ 0 & 0 & \frac{K_4}{T_4} & -\frac{1}{T_4} & 0 & 0 & 0 & \frac{K_4}{T_4} & 0 \\ \frac{K_5 a_H}{T_5} & \frac{K_5 a_L}{T_5} & 0 & \frac{K_5 a_N}{T_5} & -\frac{1}{T_5} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{K_6}{T_6^2} & -\frac{1}{T_6^2} & -\frac{2\xi}{T_6} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{k_7}{T_7} & 0 & -\frac{1}{T_7} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{k_8}{T_8} & 0 & 0 & -\frac{1}{T_8} \end{pmatrix}$$

$$B = \left(\frac{K_1}{T_1},\ \frac{K_2}{T_2},\ \frac{K_3}{T_3},\ 0,\ 0,\ 0,\ 0,\ 0,\ 0\right)^T, \qquad E = \begin{pmatrix} 0 & \frac{K_1 k_f}{T_1} & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ \frac{K_5}{T_5} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & \frac{1}{T_6^2} & \frac{1}{T_6^2} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \qquad C = (0\,0\,0\,0\,0\,1\,0\,0\,0)$$

The neobank coupling is a single element of A (row 3, column 8, in bold). At $\theta = 0$ row 3 contains only the diagonal term, so the neobank pole $-1/T_3$ is decoupled from the loops.

The closed-loop transfer function from final demand to the price is

$$\Phi_{yg}(s) = \frac{W_5 W_6 (a_H W_1 + a_L W_2 + a_N W_3 W_4)}{1 - W_5 W_6 W_7 (a_L W_2 + a_N W_4 + \theta\, a_N W_3 W_4) + W_5 W_6 W_8\, a_H W_1}$$

At $\theta = 0$ the neobank appears only in the numerator; at $\theta > 0$ it enters the denominator as well, i.e. it becomes part of the system's dynamics rather than just a signal path.

6. Disturbances and identification of the market dynamics

Disturbances. The election f_1 and the hypothetical risk f_2 act at the input of the market link W_6, the background f_0 at the input of W_5.

Disturbance	Timing	Point of application	Nature	Evidence
f_1 — midterm election	Reaction from the Nov 4 session: $t_1 = 20$	Input of W_6 (sector risk premium) plus a slow channel $k_f f_1$ into W_1	Common (systematic)	The elections took place, including in Texas and Georgia (about a third of AI data-center spending); Nov 5: Nasdaq −5.52%, SOX −12%, NVDA −9.26%
f_2 — hypothetical risk	Reaction from the Nov 9 session: $t_2 = 23$	Input of W_6, bypassing the stakeholders	Specific (idiosyncratic)	Nov 9: NVDA −4.27% with Nasdaq +2.74%; NVIDIA was the only SOX stock to fall
f_0 — October background	$t_0 = 0$	Input of W_5	Common, fundamental	AI-return doubts, CoreWeave debt, data-center delays due to power constraints

The specific component of f_2 is confirmed by the abnormal return with NVIDIA's beta ≈ 2.2:

$$\varepsilon_{\text{Nov 9}} = r_{\text{NVDA}} - \beta\, r_{\text{Nasdaq}} = -4.27\% - 2.21 \cdot 2.74\% \approx -10.3\%$$

Short-horizon model. Over 25 days the slow feedback links barely move ($|W_7|$, $|W_8| \approx k \cdot t/T \leq 0.06$), so the loop denominator is close to 1 and

$$Y(s) \approx \frac{K_0}{(T_5 s + 1)(T_6^2 s^2 + 2\xi T_6 s + 1)} \cdot \frac{1}{s} + \frac{F_1 e^{-t_1 s} + F_2 e^{-t_2 s}}{T_6^2 s^2 + 2\xi T_6 s + 1} \cdot \frac{1}{s}, \qquad K_0 = K_5 K_6 F_0$$

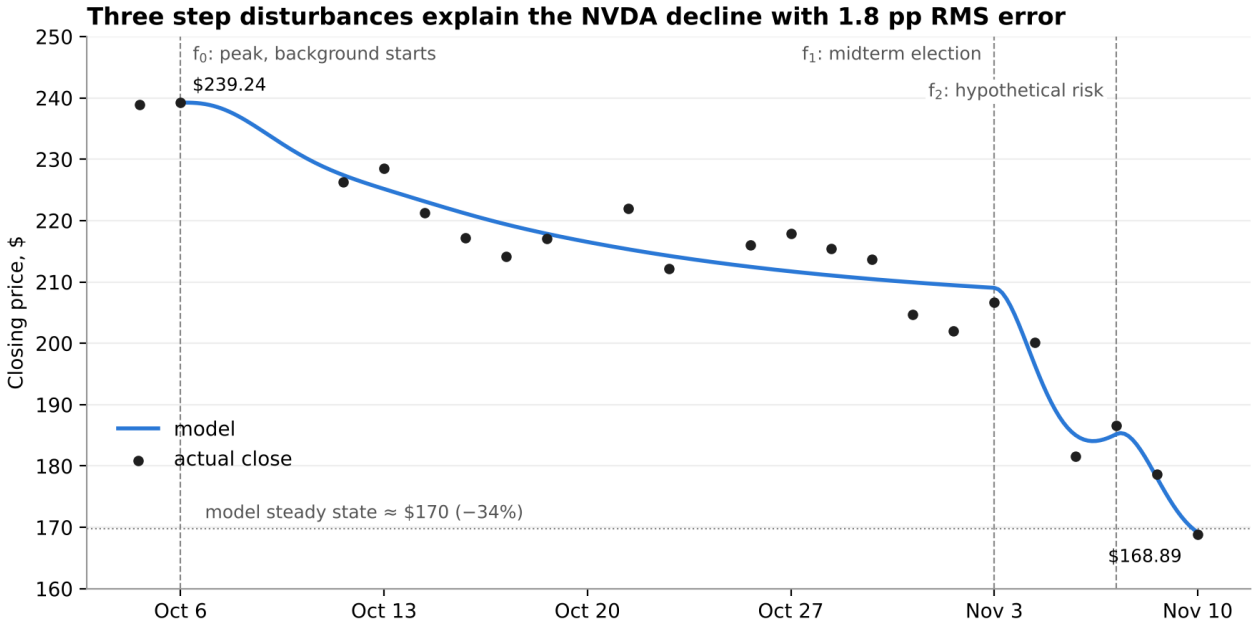


Figure 2. Actual NVDA closes (dots) and the identified model (line). Data: Gotrade/Yahoo Finance, StockTitan

Identification. The engineering method on the election response gives an overshoot σ = 26.8% (from Δmax = −12.94 pp and Δss = −10.20 pp), $\xi = -\ln\sigma/\sqrt{(\pi^2 + \ln^2\sigma)} = 0.386$ and, with peak time t_p = 2.0 days, T_6 = 0.59 days. Least squares with ξ fixed at 0.40 (the cost function is almost flat in ξ: RMS 1.66–1.91 pp for ξ from 0.05 to 0.95) gives the final estimates:

Quantity	Value
Market time constant T_6; damping ratio ξ	0.70 ± 0.13 d; 0.40
Expectations time constant T_5	8.8 ± 2.9 d
Amplitudes K_0 (background), F_1 (election), F_2 (hypothetical risk)	−15.2%, −9.8%, −9.3%
Natural / damped frequency	1.43 / 1.31 rad/d
Overshoot; peak time; settling time (5% / 2%)	25.4%; 2.4 d; 5.3 / 7.0 d
Decay ratio ψ; logarithmic decrement δ; resonance peak M	0.94; 2.74; 1.36
Fit quality	RMS 1.76 pp; R^2 = 0.96

The ratio F_1 / ΔNasdaq = 9.8 / 4.16 = 2.36 matches β: the election shock reached NVIDIA as a purely market-wide disturbance, while the hypothetical risk was idiosyncratic. Only the fast market part is identified from these data; the parameters of the slow links remain expert estimates, which is why the stability results below are reported as functions of θ, k_7 and ε rather than as single numbers.

7. How the neobank changes the stability of the system

This section compares two configurations: $\theta = 0$ (neobank outside the positive loop) and $\theta = 1$ (neobank fully inside it), with intermediate values where useful.

7.1. Characteristic equation. Multiplying the loop denominator by all link denominators gives

$$\begin{aligned} D(s) =& (T_1 s + 1)(T_2 s + 1)(T_3 s + 1)(T_4 s + 1)(T_5 s + 1)(T_6^2 s^2 + 2\xi T_6 s + 1)(T_7 s + 1)(T_8 s + 1) \\ & - K_5 K_6 k_7 (T_1 s + 1)(T_8 s + 1)[a_L K_2 (T_3 s + 1)(T_4 s + 1) + a_N K_4 (T_2 s + 1)(T_3 s + 1) + \theta\, a_N K_3 K_4 (T_2 s + \\ & + K_5 K_6 k_8\, a_H K_1\, (T_2 s + 1)(T_3 s + 1)(T_4 s + 1)(T_7 s + 1) = 0 \end{aligned}$$

At $\theta = 0$ every term contains the factor $(T_3 s + 1)$, which can be taken out: the neobank root $s = -1/T_3 = -0.0238$ stays fixed, and the neobank cannot affect stability. The bold term breaks this factorisation. From $\theta > 0$ the neobank pole moves and becomes one of the roots shaped by the loops.

7.2. Which coefficients change. In non-normalised form (multiplied by the product of all time constants):

Coefficient	a_9	a_8	a_7	a_6	a_5	a_4	a_3	a_2	a_1	a_0
$\theta = 0$	$3.795 \cdot 10^{10}$	$5.551 \cdot 10^{10}$	$9.274 \cdot 10^{10}$	$2.653 \cdot 10^{10}$	$3.053 \cdot 10^{9}$	$1.675 \cdot 10^{8}$	$4.632 \cdot 10^{6}$	$6.481 \cdot 10^{4}$	422.6	0.950
$\theta = 1$	$3.795 \cdot 10^{10}$	$5.551 \cdot 10^{10}$	$9.274 \cdot 10^{10}$	$2.653 \cdot 10^{10}$	$3.053 \cdot 10^{9}$	$1.675 \cdot 10^{8}$	$4.605 \cdot 10^{6}$	$6.298 \cdot 10^{4}$	395.3	0.868

Only the four lowest-order coefficients change, by 0.6% to 9%. High-order coefficients govern fast motion and low-order ones govern slow motion, so the neobank does not alter the market's daily reaction. It changes the slow fundamental dynamics and the steady state. The constant term is the key one:

$$a_0 = 1 - K_5 K_6 k_7 (a_L K_2 + a_N K_4 + \theta\, a_N K_3 K_4) + K_5 K_6 k_8\, a_H K_1 = \begin{cases} 0.950, & \theta = 0 \\ 0.868, & \theta = 1 \end{cases}$$

7.3. Roots.

Root ($\theta = 0$), 1/d	Root ($\theta = 1$), 1/d	Origin and change
$-0.5713 \pm j1.3089$	$-0.5713 \pm j1.3089$	Market link W_6, unchanged
$-0.1072 \pm j0.0115$	$-0.1066 \pm j0.0092$	Expectations and neoclouds, slightly slower
−0.0491	−0.0493	AI labs, practically unchanged
−0.0238	−0.0273	Neobank pole: fixed at $-1/T_3$ for $\theta = 0$, drawn into the loop for $\theta = 1$
−0.0159	−0.0159	Exactly −1/63 (common factor because $T_1 = T_7$)
−0.0126	−0.0098	Hyperscalers and financing: time constant grows from 79 to 102 days
−0.0045	−0.0047	Substitution, practically unchanged

At the current financing gain all roots stay in the left half-plane in both configurations. The Routh table confirms this: its first column for $\theta = 1$ ($3.80 \cdot 10^{10}$; $5.55 \cdot 10^{10}$; $7.46 \cdot 10^{10}$; $2.43 \cdot 10^{10}$; $2.44 \cdot 10^{9}$; $1.20 \cdot 10^{8}$; $3.18 \cdot 10^{6}$; $4.46 \cdot 10^{4}$; 313; 0.868) has no sign changes. The degree of stability, set by the slowest root, is almost the same (−0.0045 versus −0.0047). The coupling's effect shows not in today's transient but in how close the system is to its boundary.

7.4. Stability boundary and margin. Positive feedback loses stability aperiodically: a real root crosses zero when $a_0 = 0$. Solving for the critical financing gain:

$$k_7^{\mathrm{cr}}(\theta) = \frac{1 + K_5 K_6 k_8\, a_H K_1}{K_5 K_6\, (a_L K_2 + a_N K_4 + \theta\, a_N K_3 K_4)} = \frac{1.072}{0.816 + 0.546\,\theta}$$

θ	Open-loop positive-feedback gain L	Critical k_7	Margin k_7^{cr} / k_7 (k_7 = 0.15)	Static gain g → y	Election-shock gain f_1 → y
0	0.122	1.31	8.8	1.62	1.05
0.25	0.143	1.13	7.5	1.65	1.08
0.50	0.163	0.98	6.6	1.69	1.10
0.75	0.184	0.88	5.8	1.73	1.13
1	0.204	0.79	5.2	1.77	1.15

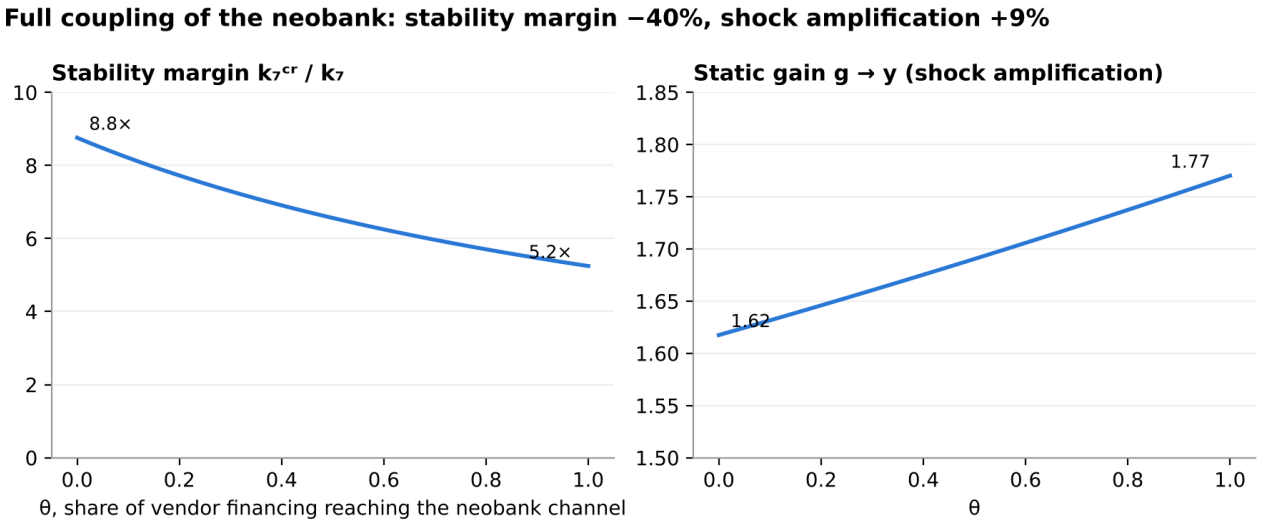


Figure 3. Stability margin and shock amplification as functions of the neobank coupling θ

Full coupling of the neobank reduces the stability margin by 40% (from 8.8 to 5.2) and raises the closed-loop gains by 9.4% for demand shocks and for every disturbance entering the market, including elections and the hypothetical risk. A neobank coupled to the vendor makes the AI compute market both closer to instability and more sensitive to shocks.

7.5. Root locus in the financing gain. Figure 4 shows the largest real part of the roots as k_7 grows.

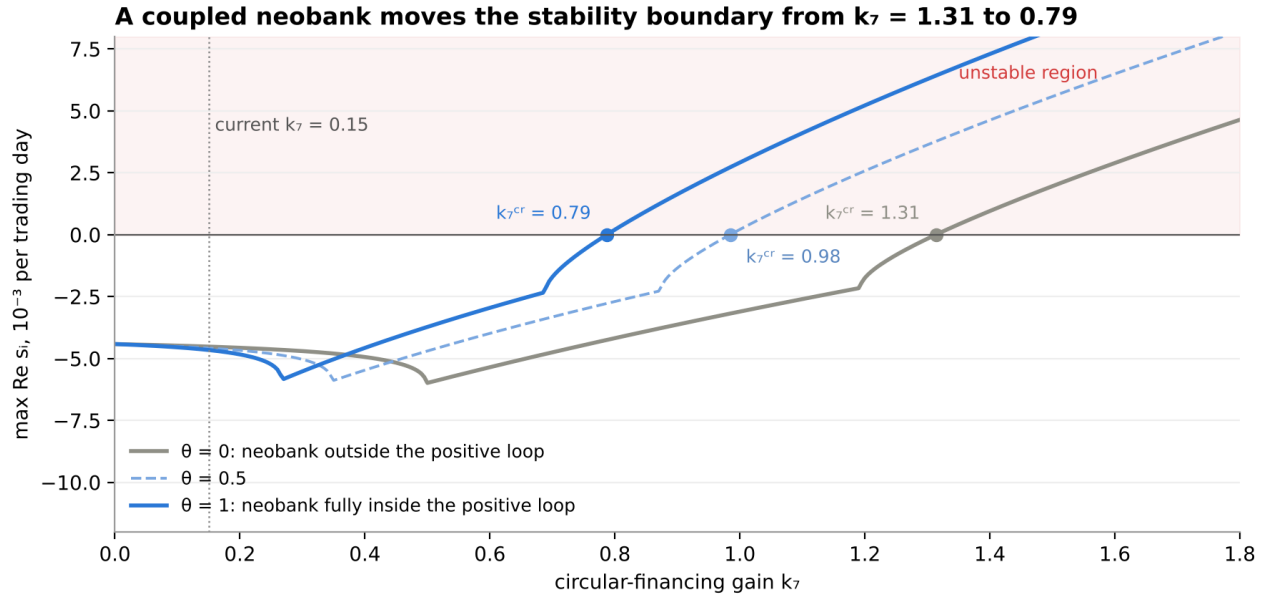


Figure 4. Largest real part of the characteristic roots versus the circular-financing gain k_7 for θ = 0, 0.5 and 1

The curves are non-monotonic. Moderate vendor financing first makes the slowest mode decay faster: the financing root and the substitution root approach each other, merge into a complex pair and move left. The degree of stability is best at $k_7^* = 0.50$ for $\theta = 0$ ($\eta^* = 0.0060$ 1/d). Beyond this point a real root moves right and crosses zero at k_7^{cr}. Coupling the neobank compresses the whole picture: the optimum shifts to $k_7^* = 0.27$ and the boundary to 0.79. The current $k_7 = 0.15$ is below the optimum in both configurations, but with $\theta = 1$ the distance to the optimum shrinks from 0.35 to 0.12. The range in which more vendor financing still helps stability is roughly halved.

7.6. Dynamic consequences at the current operating point. Coupling does not destabilise the system today, but it changes the transient response to a sustained change in final demand:

Indicator	$\theta = 0$	$\theta = 1$	Change
Static gain g → y	1.62	1.77	+9.4%
Time to 63% of the steady value	64 d	72 d	+13%
Time to 95% of the steady value	152 d	185 d	+22%
Peak of the step response	1.67	1.81	+9%

The system becomes more inert and amplifies more: a fall in AI demand takes longer to play out and ends lower. For the October–November 2026 episode the effect would be second-order, since over 25 days the slow loops are frozen (Section 6). It matters over quarters, the horizon of capex and financing decisions.

7.7. Falling token cost as a parametric drift of the neobank channel. The threefold fall in Revolut's blended cost per token is well described by an exponential decay

$$c(t) = c_0\, e^{-t/T_c}, \qquad T_c = \frac{7 \text{ months}}{\ln 3} \approx 6.4 \text{ months} \approx 134 \text{ trading days}, \qquad t_{1/2} \approx 4.4 \text{ months}$$

NVIDIA orders generated by the neobank channel are proportional to token volume V times GPU time per token. If the cost per token falls mainly because GPU time per token falls, and token volume responds to price with elasticity ε ($V \propto c^{-\varepsilon}$), the effective gain of the channel drifts as

$$K_{3,\text{eff}}(t) = K_3 \left(\frac{c(t)}{c_0}\right)^{1-\varepsilon}, \qquad K_{3,\text{eff}}|_{\text{Apr}\to\text{Nov}} = K_3 \cdot 3^{\varepsilon-1}$$

For $\varepsilon < 1$ efficiency wins and the channel orders fewer GPUs; for $\varepsilon = 1$ the cost cut is fully offset by volume; for $\varepsilon > 1$ (the Jevons effect) cheaper tokens raise GPU demand. Revolut's experience points to a clearly positive elasticity: once tokens became cheap, engineers stopped rationing them and applied AI to debugging and ticket work, not only to writing code.

In the characteristic equation K_3 appears only in the term $\theta\, a_N K_3 K_4$ (Section 7.1). Two conclusions follow directly:

- **$\theta = 0$ (customer-paid compute):** the drift of K_3 rescales NVIDIA orders, with static gain g → y from 1.23 at $\varepsilon = 0$ to 2.77 at $\varepsilon = 2$, but it does not move a single root. The stability margin stays at 8.8 whatever happens to token prices.
- **$\theta = 1$ (vendor-financed compute):** the drift enters the stability boundary, $k_7^{cr} = 1.072/(0.816 + 0.546 \cdot 3^{\varepsilon-1})$.

Elasticity ε	Gain multiplier $3^{\varepsilon-1}$	Critical k_7 ($\theta = 1$)	Margin ($\theta = 1$)	Static gain g → y, $\theta = 0$	Static gain g → y, $\theta = 1$
0	0.33	1.07	7.2	1.23	1.27
0.5	0.58	0.95	6.3	1.38	1.45
1	1.00	0.79	5.2	1.62	1.77
1.5	1.73	0.61	4.1	2.04	2.40
2	3.00	0.44	2.9	2.77	3.73

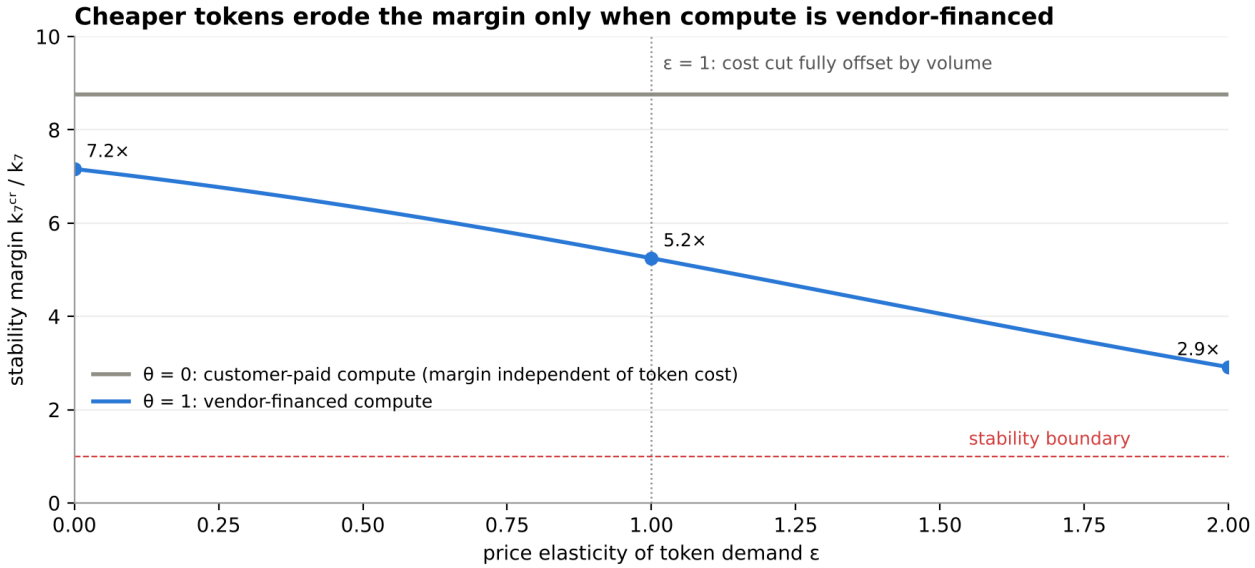


Figure 5. Stability margin after a threefold fall in token cost as a function of the price elasticity of token demand

Three further observations:

1. **Where tokens are generated does not matter for stability; who pays for the GPUs does.** Half of Revolut's tokens come from external vendors and enter the AI-lab branch W_2 as an input signal, i.e. through the numerator of the transfer function. Moving tokens from external vendors to self-hosted GPUs, as Revolut did for coding, redistributes the forward path between W_2 and W_3 and changes the size of NVIDIA's orders, but leaves the characteristic equation unchanged. Only vendor-financed GPU purchases (θ) enter it.
2. **Elastic demand turns efficiency into a destabilising factor.** With vendor-financed compute, a threefold cost fall with $\varepsilon = 2$ cuts the margin from 5.2 to 2.9. The system remains stable at the current financing gain, but more than half of the remaining margin is consumed by the token-price trend alone.
3. **The model becomes time-varying on this horizon.** $T_c \approx 134$ trading days is comparable with the slow roots (time constants 79–221 days), so the frozen-coefficient comparison above is valid for comparing operating points, not for the transient during the drift. The lack of transparency in LLM tokenomics that Pavel Nesterov points out in his post appears here as parametric uncertainty in K_3,eff; this is why the margin is reported over an interval of ε (robust stability): it stays above 1 for every ε up to 2.

7.8. Stability of the bank itself. The model describes the bank's AI compute channel, not its capital, liquidity or credit risk, so it says nothing about Revolut's solvency. Within that scope, three conclusions follow.

- *The bank's channel is stable on its own.* W_3 is a first-order lag with a single real negative pole: it cannot oscillate or overshoot. At the current financing gain all closed-loop roots stay in the left half-plane for any θ.
- *With customer-paid compute the bank is isolated from the market.* For $\theta = 0$ the transfer functions from all three disturbances to the bank's purchases x_3 are identically zero: the 29% fall in NVIDIA's share price does not reach the bank's AI infrastructure. As a private company, Revolut has no feedback through its own share price.
- *With vendor-financed compute (hypothetical scenario) the isolation disappears.* For $\theta = 1$ the static gain from any shock at the market input to x_3 is 0.17. Together, the election shock and the hypothetical risk cut the bank's compute channel by about 3.7% in steady state; 1.6 percentage points of this come from the hypothetical risk alone, i.e. from a risk specific to NVIDIA that has nothing to do with Revolut. The effect arrives slowly (Figure 6): −0.2% after a month, −1.0% after a quarter, −2.2% after half a year and −3.4% after a year, reaching 63% of its final value after about 130 trading days. A shock the market absorbed within a week reaches the bank half a year later, on the horizon of its budget cycle. The bank's sensitivity to its own final demand also rises from 1.00 to 1.27.

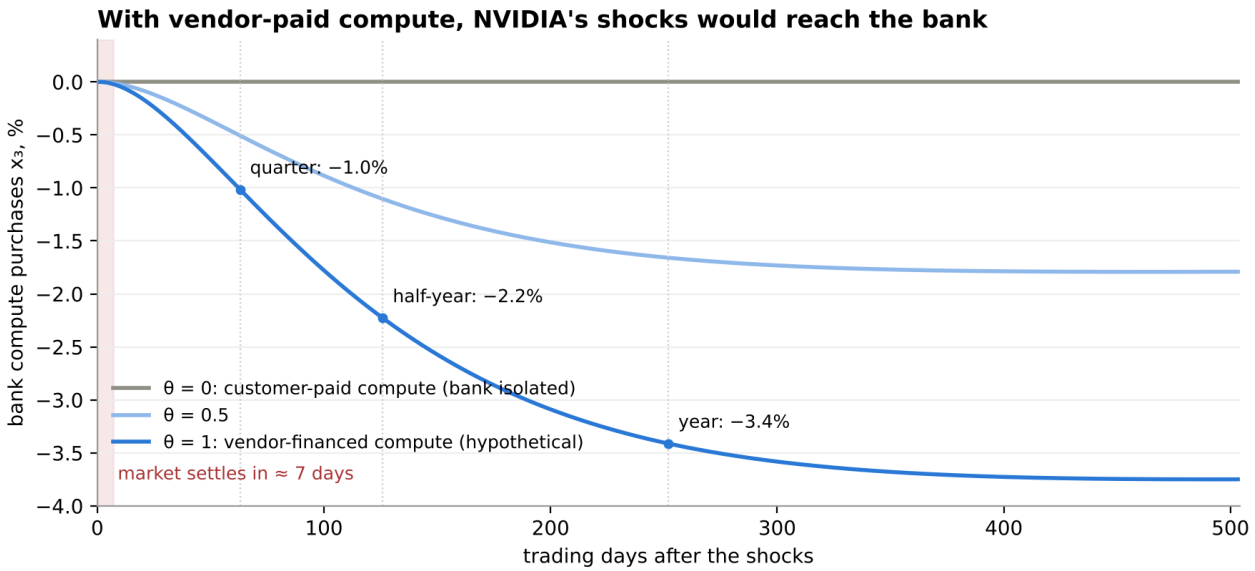


Figure 6. Response of the bank's compute purchases x_3 to the election shock and the hypothetical risk (amplitudes from Section 6) for $\theta = 0$, 0.5 and 1

In the same terms, Revolut's own choices act as hedges. Self-hosting open models on 56 B200 GPUs removes part of its dependence on the neocloud intermediary (Nebius had an election-shock gain of 2.7) and on AI-lab pricing. The half of its tokens still bought from external vendors at opaque prices remains an input uncertainty. In short, **Revolut's AI compute channel is stable on its own; only vendor-financed compute would import NVIDIA's shocks into it, and in our model they take about six months to arrive.**

8. Step responses of stakeholders

Market level. The November 5 session gives each stakeholder's transmission gain of the election shock, K = Δstock/ΔNasdaq (Figure 7). Chipmakers and neoclouds amplified the shock (K = 1.7–3.1); hyperscalers damped it (K = 0.4–0.8) because they already own capacity. Nebius, Revolut's compute provider, had K = 2.7. Revolut itself is private and has no market response.

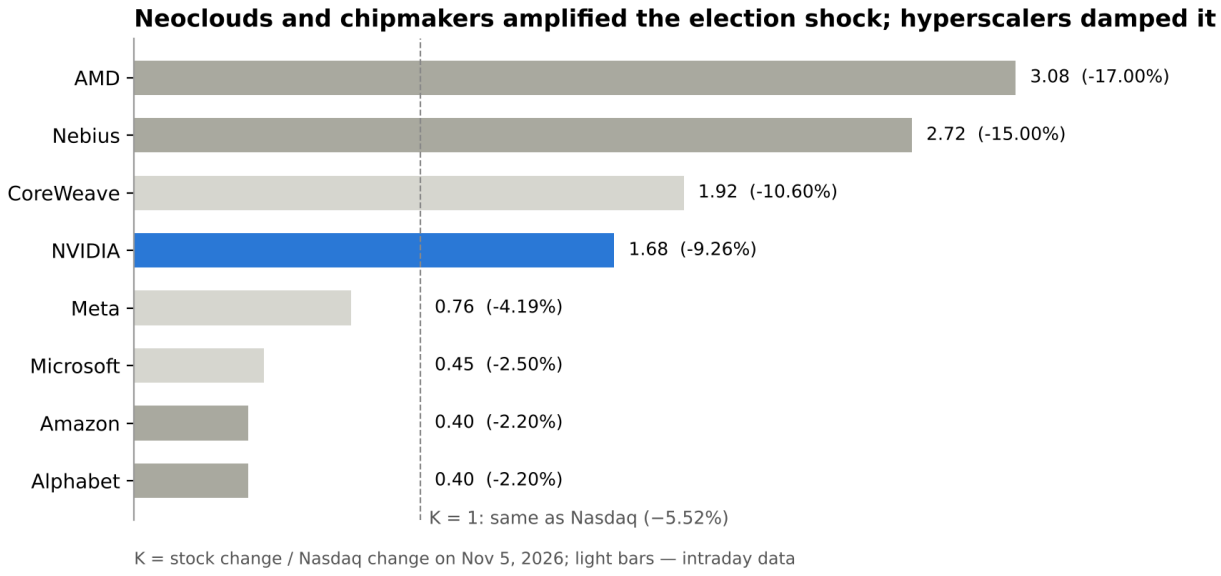


Figure 7. Transmission gain of the election shock for stakeholders, November 5, 2026 (Fox Business, Reuters, ABC, Benzinga)

Fundamental level. Model step responses (Figure 8) show that the neobank channel W_3W_4 is S-shaped: the bank's demand reaches NVIDIA with a delay. The tangent at the inflection point ($t^* = T_3T_4/(T_3 - T_4)\cdot\ln(T_3/T_4)$ = 18.8 d, $h(t^*)$ = 0.275, slope 0.0198 1/d) gives the first-order-plus-dead-time approximation $W_3W_4(s) \approx 1.3\cdot e^{(-5.1s)}/(65.8s + 1)$, with $\tau/T_a = 0.08$.

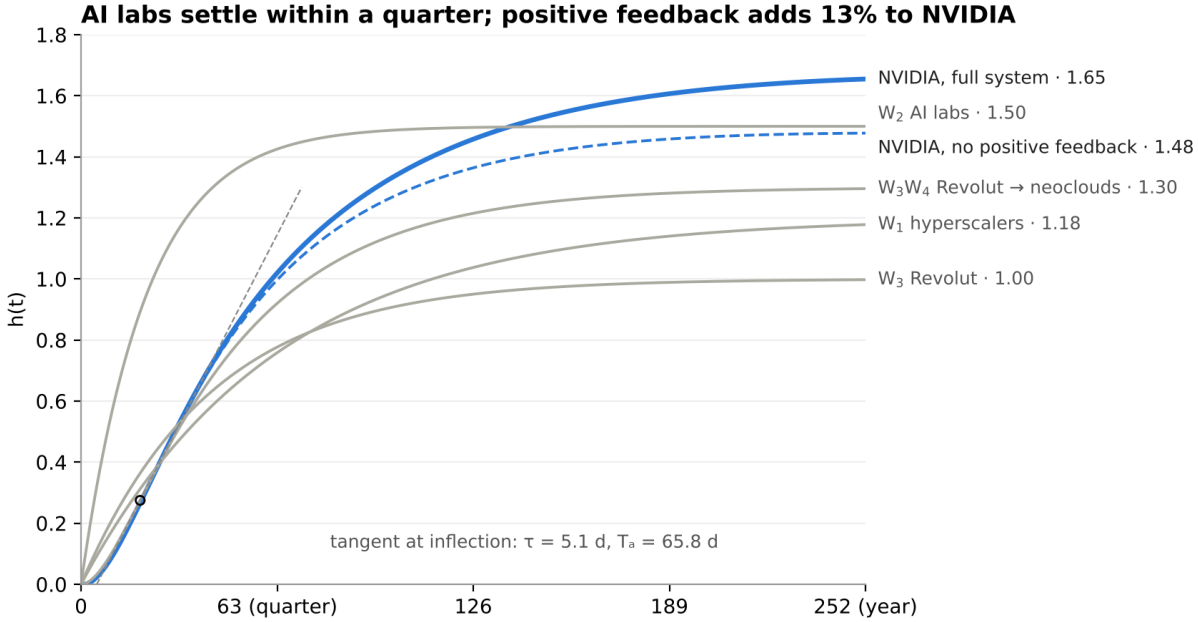


Figure 8. Model step responses of stakeholders and of NVIDIA ($\theta = 0$)

9. Discussion: neobanks as a new element of market control in AI

1. **A neobank is a control element whose sign depends on who pays.** With customer-paid compute ($\theta = 0$) it is a feedforward channel: it transmits final demand to the vendor and cannot change the system's poles. With vendor-financed compute ($\theta \to 1$) it becomes part of the positive loop and shifts the stability boundary. The same company can play either role; what matters is the source of the money behind its GPU hours.
2. **The vendor is hedged against model competition, not against financing risk.** Whether a bank buys a model from a lab (W_2) or trains its own (W_3W_4), the signal reaches the same order junction, which is consistent with NVIDIA's stated aim to support all foundation-model companies rather than pick winners. Revolut's own move illustrates this: it replaced commercial coding agents with open-weight models running on NVIDIA B200 GPUs. Revenue left the AI-lab branch, while NVIDIA still sold the hardware. Competition among model developers redistributes the signal between parallel branches; it does not threaten stability. Stability is threatened by the growing number of branches that are also financed by the vendor.
3. **The hardware intermediary amplifies market shocks.** The neobank's demand passes through a neocloud whose share price moved 2.7 times as much as the Nasdaq on the day of the election shock. Even when the neobank itself is insulated from market shocks, the market valuation of its compute provider is not.
4. **Small demand, large structural effect.** Revolut's PRAGMA models were trained in two days to two weeks on 16–32 GPUs, negligible next to hyperscaler capex. In the model, θ multiplies the whole neobank-and-enterprise channel, so the 40% margin loss describes a scenario in which this sector as a whole is vendor-financed. Revolut matters as a template: for NVIDIA the stake is a way to be first in line when the bank scales its AI infrastructure, and other financial institutions may follow the same pattern.
5. **Efficiency gains do not stabilise the market by themselves.** A threefold fall in token cost multiplies the neobank channel gain by $3^{\varepsilon-1}$. With customer-paid compute this only rescales orders. With vendor-financed compute and elastic demand ($\varepsilon > 1$), cheaper tokens reduce the stability margin further (Section 7.7). Falling token prices and circular financing reinforce each other's destabilising effect.
6. **A monitoring indicator.** The share θ of vendor-financed demand in each stakeholder channel is a measurable quantity: equity stakes, guarantees and compute commitments are partly disclosed. Tracking it,

together with the circular-financing gain k_7 and an estimate of the token-demand elasticity ε, gives a direct estimate of the remaining stability margin $k_7^{cr}(\theta, \varepsilon)/k_7$.

Limitations. The $\theta > 0$ scenarios are hypothetical and do not describe any actual financing arrangement. Only the fast market part (T_6, ξ, T_5 and disturbance amplitudes) is identified from data; the slow-link parameters are expert estimates, and the absolute values of k_7^{cr} should be read as illustrative. The relative effect of θ is more robust: it follows from the structure of the characteristic equation and holds for any parameter set in which the neobank channel carries a non-negligible share of orders. The token-cost analysis rests on two simplifying assumptions: that falling cost per token reflects falling GPU time per token, and that elasticity is constant; both should be checked against disclosed usage volumes.

10. Conclusions

1. NVIDIA's October–November 2026 share-price decline is described by a linear closed-loop model with three step disturbances (background −15.2%, election −9.8%, hypothetical risk −9.3%) and an oscillatory market link with $T = 0.70$ days and $\xi = 0.40$ ($R^2 = 0.96$).
2. A neobank whose compute is paid for by its customers enters the model only as a feedforward path: its pole is fixed and it cannot change stability.
3. In the hypothetical scenario where vendor financing is converted into the neobank's compute purchases, the neobank joins the positive loop. At full coupling the critical financing gain falls from 1.31 to 0.79, the stability margin from 8.8 to 5.2 (−40%), the closed-loop amplification of demand and market shocks rises by 9.4% and the response to demand changes slows by 22%.
4. The coupling changes only the lowest-order coefficients of the characteristic polynomial: it leaves the market's fast reaction intact and acts on the slow fundamental dynamics, over the horizon of capex and financing decisions.
5. Moderate vendor financing improves the degree of stability up to an optimum, but coupling the neobank roughly halves the range in which this holds (k_7^* from 0.50 to 0.27). The share of vendor-financed demand θ is therefore a natural indicator of systemic risk in the AI compute market.
6. Falling token costs, three times cheaper in seven months at Revolut, act as a slow parametric drift of the neobank channel gain. They do not affect stability when compute is customer-paid; when it is vendor-financed, the margin ranges from 7.2 ($\varepsilon = 0$) to 2.9 ($\varepsilon = 2$). Elastic token demand thus amplifies the destabilising effect of circular financing.
7. The bank's own compute channel is stable for any θ. With customer-paid compute it is fully isolated from market shocks; with vendor-financed compute (a hypothetical scenario) it imports them: together, the election shock and the hypothetical risk cut its compute purchases by about 3.7%, with the effect building up over about half a year.

Acknowledgments

This paper was prepared with extensive use of an AI model for source search, calculations, figures, drafting and translation. The author reviewed all results and is responsible for the content.

References

1. Gotrade News. NVIDIA stock (NVDA) daily reports, October 13 – November 9, 2026: Nov 9, Nov 6, Nov 5, Nov 4, Nov 3, Nov 2, Oct 30, Oct 29, Oct 27, Oct 22, Oct 19, Oct 16, Oct 15, Oct 14, Oct 13.
2. StockTitan. Nvidia stock P/E ratio lowest since 2019, November 10, 2026.
3. Fox Business. Stock market report, November 5, 2026; NVDA quote.
4. Trefis. What sent NVIDIA stock lower?, November 6, 2026.
5. 24/7 Wall St. The Nasdaq-100 just had its best day since March without Nvidia's help, November 10, 2026.
6. Advisor Perspectives / Bloomberg. Wall Street peers into abyss of AI trade on midterms risk; AI Chat Daily, AI data center stocks after the midterms.
7. Stakeholder moves on November 5, 2026: Benzinga; GV Wire / Reuters; ABC7; The Financial Express.

8. NVIDIA. Revolut case study; NVIDIA Blog, European financial services organizations embrace open models.
9. American Banker. Revolut valued at $75 billion as Nvidia joins share sale; Finextra, Nvidia invests £2 billion into UK AI ecosystem; Banken.nl, Revolut valuation after Nvidia stake.
10. DEV Community. The second customer; FourWeekMBA, Nvidia revenue breakdown; Daloopa, NVIDIA customer concentration.
11. The Decoder. Nvidia pumps over $40 billion into AI partners; eNCA / AFP, Nvidia's web of AI bets; Gulf News, Nvidia's web of AI bets: smart strategy or bubble risk?.
12. getquin. Nvidia stock forecast after the US election, November 2026.
13. The Register. British fintech Revolut slashes AI coding costs 10x by cutting out Claude and going self-hosted, November 3, 2026.
14. Computing. Revolut turns to self-hosted AI to cut coding costs tenfold over Claude, 2026; Analytics India Magazine, Europe's biggest tech startup cuts AI coding costs by 90% using Chinese model.
15. Pavel Nesterov. Post on Revolut's AI KPI "Blended LLM Inference Cost per 1M Tokens (last 90 days)", LinkedIn, November 2026.
16. F. Desyatirikov. Dynamic Model of the Java Ecosystem Hype Cycle. In: 2024 Conference of Young Researchers in Electrical and Electronic Engineering (ElCon), Saint Petersburg, Russian Federation, January 29–31, 2024. IEEE, 2024. DOI: 10.1109/ElCon61730.2024.10468112.
17. M. Ostroukhov, R. Mikhailov, V. Iashin, A. Sokolov, A. Akshonov, V. Protasov, D. Beloborodov, V. Mullin, R. Y. Enzmann, G. Kolovos, J. Renders, P. Nesterov, A. Repushko. PRAGMA: Revolut Foundation Model. arXiv:2604.08649v2, 2026.
18. A. Tustin. The Mechanism of Economic Systems. London: Heinemann, 1953.
19. A. W. Phillips. Stabilisation Policy in a Closed Economy. The Economic Journal, 64(254): 290–323, 1954.
20. R. Neck. Control theory and economic policy: Balance and perspectives. Annual Reviews in Control, 33(1): 79–88, 2009.
21. T. Andresen. A block diagram approach to macroeconomic dynamics, and why IS/LM is fatally flawed. Department of Engineering Cybernetics, NTNU, Trondheim.
22. S. M. Disney, D. R. Towill. A discrete transfer function model to determine the dynamic stability of a vendor managed inventory supply chain. International Journal of Production Research, 40(1): 179–204, 2002.