The Name on the Chip

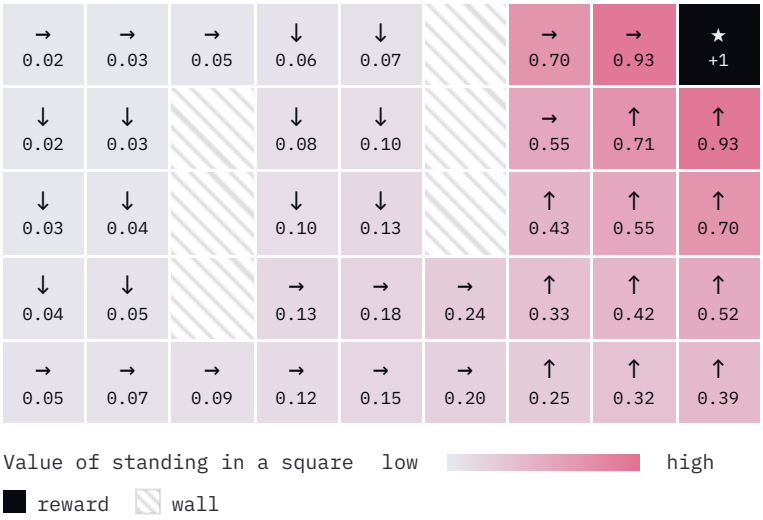
James E. Coleman

STRATEGIC DESIGN · WHITE PAPER SERIES · 2026

ABSTRACT Nvidia's flagship AI processor carries the name of David Blackwell, a mathematician who wrote more than a hundred letters to Black colleges because the universities that trained him would not hire him. This dispatch follows the name: where it came from, what the mathematics does inside the machine, and an evidence-graded account of what has happened to the machines since. Where they went, who rented them, and what a spy story did to a four-trillion-dollar company.

FIG. 1 — BLACKWELL, 1965: REPEAT THE UPDATE AND THE ERROR HAS TO SHRINK

SWEEP 20 · LARGEST CHANGE 0.001



SWEEP

20

One pass of the update over every square. Arrows show the best move so far.

LARGEST CHANGE

0.001

This sweep ÷ last sweep: **0.60**. The theorem caps that ratio at the discount factor, 0.80.

LARGEST CHANGE, SWEEP BY SWEEP





208 billion

transistors in one Blackwell B200, against **80 billion** in the H100 it replaces.[5]



104 letters

sent to historically Black colleges in 1942. **Three offers** came back.[1]



−9.26%

Nvidia's close on **November 5**, not November 9, and the coverage blamed regulation fears.[20]



2.7%

the lead the top US model still held over the top Chinese model in **March 2026**.[15]

§ 01 Origin

A railroad town, a doctorate at 22, *and two doors that stayed shut.*

David Harold Blackwell was born on April 24, 1919, in Centralia, Illinois, where his father, Grover, worked for the Illinois Central Railroad. He entered the University of Illinois in 1935 at sixteen, took a bachelor's degree in 1938 and a master's in 1939, and in 1941, at twenty-two, finished a doctorate under Joseph Doob with a thesis on Markov chains. He was the seventh Foundational Black American to earn a PhD in mathematics in the United States and the first to earn one at Illinois.[1,2]

The credential opened one door and then no more. A Rosenwald Fellowship took him to the Institute for Advanced Study in 1941, and Princeton University objected to his presence on racial grounds. A year later the statistician Jerzy Neyman backed him for a post at Berkeley, and that appointment was blocked by prejudice as well.[1,2] He was twenty-three, with a doctorate from a Big Ten university and a year at the Institute, and no research university in the country would put him on its faculty.

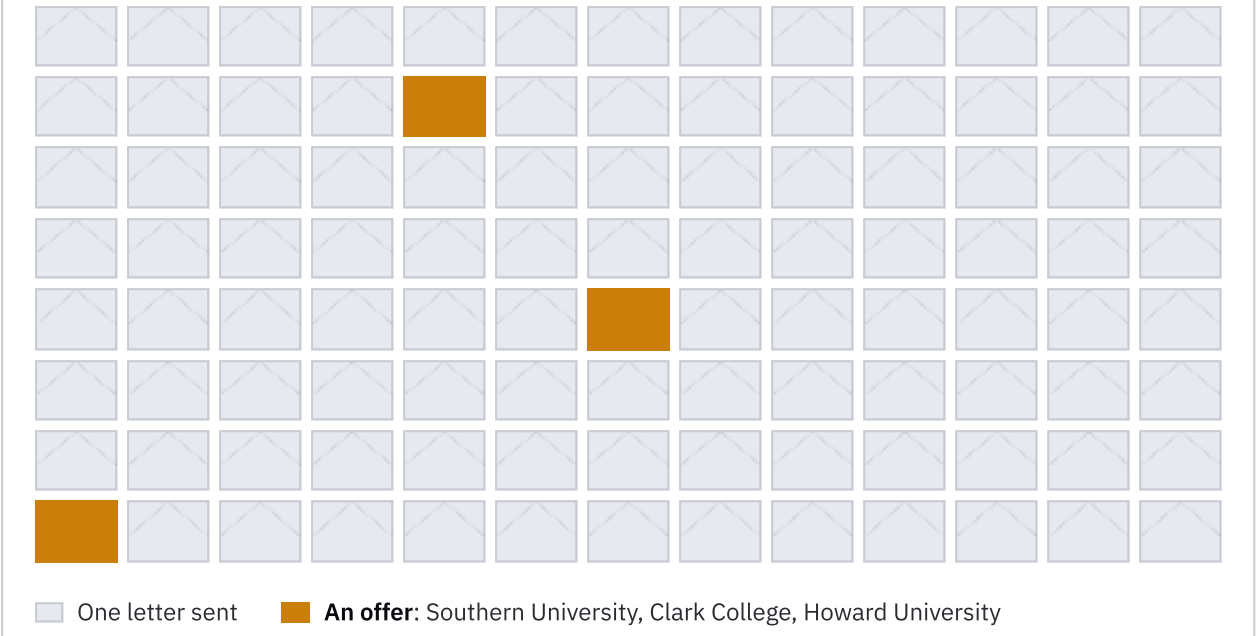
So he did what a person does when the front office will not take the meeting. He wrote letters.

§ 02 The innovation complex

One hundred and four letters. *Three offers.*

Blackwell wrote to the country's historically Black colleges and universities, one at a time, asking for work. One standard account puts the count at 104 letters. The St Andrews biography says 105. Both agree on what came back: three offers.[1,2] He taught at Southern University in Baton Rouge in 1942–43, at Clark College in Atlanta the year after, and in 1944 joined Howard University, where he stayed a decade and was a full professor and head of the mathematics department by 1947.

FIG. 2 — THE JOB SEARCH, 1942 — 3 OFFERS FROM 104 LETTERS



The count is 104 in one biography and 105 in the other; the figure draws 104. Which letters drew the offers is not recorded, so the three filled envelopes are placed for legibility.[1,2]

The Howard decade is the part of the story the chip's marketing leaves out. It was from Howard that he published the 1947 paper behind what is now called the Rao–Blackwell theorem, the 1953 paper on comparing experiments, and, with M. A. Girshick, the 1954 book *Theory of Games and Statistical Decisions*. He spent the summers of 1948 through 1950 at the RAND Corporation.[2,3] Berkeley, which had turned him away in 1942, brought him in as a visitor in 1954 and made him a full professor of statistics in 1955, the first Foundational Black American to hold tenure as a full professor there. He went on to chair the department, supervise more than sixty doctoral students, and in 1965 become the first Foundational Black American elected to the National Academy of Sciences. He died in Berkeley in 2010, at ninety-one. The National Medal of Science came after his death.[1,2]

In the last dispatch in this series I argued that American breakthroughs are born with original issuers: the small, thinly capitalized places where the new is first made, long before the capital shows up. Howard University in 1947 was an original issuer. The institution that carried the risk on David Blackwell, when no research university would, was a Black college. The result of that bet is now printed on the most valuable hardware in the world, and the college's name is not on it.

The pattern has not gone away. HBCUs are about 3% of American colleges. By the United Negro College Fund's 2026 count they award 28% of the STEM bachelor's degrees earned by Black students. Georgetown's Center for Security and Emerging Technology, measuring a different way, puts the HBCU share of Black STEM graduates at 19% in 2021, down from 25% a decade earlier.[8] The two figures use different definitions and different years. I report both rather than choose the flattering one, and either way the conclusion holds: a small set of under-resourced schools is doing a disproportionate share of the country's technical talent formation.

What a Blackwell chip *actually computes.*

Precision first. Nothing in a B200's transistors is "Blackwell's math." Silicon does arithmetic, and it will multiply matrices for anyone. The lineage is in the workloads. The jobs these machines are bought to run are, to a striking degree, industrial-scale versions of problems Blackwell formalized between 1947 and 1965.

Start in 1965. In "Discounted Dynamic Programming," Blackwell studied a decision-maker choosing actions over an unbounded future, with later rewards discounted by a factor below one. He showed that the update which improves an estimate of long-run value is a *contraction*: apply it once, and the distance to the true answer shrinks by at least the discount factor. Apply it again, and it shrinks again. A contraction has exactly one fixed point, so there is one optimal value, repeated updating converges to it at a geometric rate, and there is an optimal rule that never needs to change from one day to the next.[3,4]

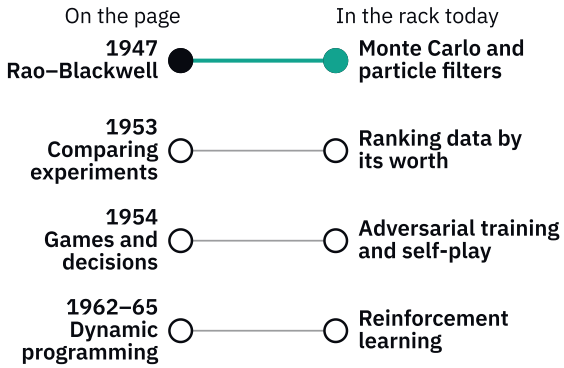
The figure at the top of this paper is that theorem running in your browser. Each sweep is one application of the update. The readout divides the largest change on the board by the one before it, and that ratio never exceeds 0.80, the discount factor, exactly as the paper says it cannot. This guarantee sits underneath value-based reinforcement learning, the family of methods used to train game-playing agents and, increasingly, robots in simulation. I found no source that traces any single modern algorithm straight back to the 1965 paper and I do not claim one. The textbook lineage runs to the discounted framework as a whole.[4]

Then 1947. The Rao–Blackwell theorem says that if you average an estimator over what a sufficient statistic already tells you, the result is never worse and is usually better. In practice it is a recipe for squeezing noise out of a sampling procedure without collecting more samples. Machine-learning researchers still use the name as a verb: Rao-Blackwellized particle filters for robot localization, Rao-Blackwellized gradient estimators in training. On a GPU cluster, lower variance per sample is fewer GPU-hours per result.

And 1954. *Theory of Games and Statistical Decisions* treats statistics as a game against nature and learning as a contest between players. Adversarial training and multi-agent self-play, both staples of the modern stack, descend from that framing.

FIG. 3 — FROM THE PAGE TO THE RACK

STATE 1 OF 4



HOWARD

Rao–Blackwell

Rao–Blackwell. Conditioning an estimator on a sufficient statistic never increases its variance. **Runs today as** variance reduction in Monte Carlo sampling and gradient estimation, and in Rao-Blackwellized particle filters.[3]

HOWARD

Comparison of experiments

Comparison of experiments. A precise test for when one source of information is always at least as useful as another. **Runs today as** the formal basis for ranking sensors, datasets and signals by what they are worth to a decision.[3]

HOWARD AND RAND

Games and statistical decisions

Games and statistical decisions, with M. A. Girshick. Inference framed as a game. **Runs today as** the decision-theoretic frame behind adversarial training and self-play.[3]

BERKELEY

Discounted dynamic programming

Discounted dynamic programming. The update is a contraction, so there is one optimal value and iteration reaches it geometrically. **Runs today as** value iteration and the discounted objective in reinforcement learning.[3,4]

A diagram of the relation, not a chart of data. Each line joins a result Blackwell published to the place it runs in machine learning today.

The right-hand side of each row is an intellectual lineage, not a hardware feature. Three of the four results were written at an HBCU.

When Nvidia announced the platform on March 18, 2024, its release described Blackwell as a mathematician who specialized in game theory and statistics and the first Black scholar inducted into the National Academy of Sciences.[5] The company put the name of the man who proved reinforcement learning's bedrock convergence result on the hardware built to run it. Whether by research or by instinct, the name is exact.

From a diner booth *to a liquid-cooled rack.*

Nvidia is a Main Street origin story too. In 1993 Jensen Huang, then at LSI Logic, and two Sun Microsystems engineers, Chris Malachowsky and Curtis Priem, planned the company in a booth at a Denny's on Berryessa Road in San Jose. It was incorporated on April 5, 1993.[6] The product was a graphics card for video games. Consumer technology, sold in a retail box, to people who wanted a better frame rate.

The arithmetic that draws a game frame turned out to be the arithmetic that trains a neural network, and the company that sold to the arcade ended up supplying the laboratory. Blackwell is where that road arrives. A single B200 carries 208 billion transistors, against 80 billion in the H100, built as two dies on TSMC's 4NP process. A GB200 pairs two of them with one Grace CPU. A GB200 NVL72 puts 72 Blackwell GPUs and 36 CPUs in one liquid-cooled rack, which Nvidia rates at up to thirty times an equal number of H100s for language-model inference. That is the vendor's claim and not an independent benchmark.[5]

FIG. 4 — THIRTY-ONE YEARS, FOUR STEPS — BOOTH TO RACK

1993 · SAN JOSE

3 founders

One booth at a Denny's. A graphics-card company for PC games.[6]

2024 · B200

208 billion

Transistors on one GPU, two dies joined on TSMC 4NP.[5]

GB200

2 + 1

Two Blackwell GPUs and one Grace CPU on a single module.[5]

GB200 NVL72

72 GPUs

Plus 36 CPUs in one liquid-cooled rack, sold as a single machine.[5]

The money followed the machine. For the fiscal year that ended January 25, 2026, Nvidia reported revenue of $215.9 billion, up 65%. In the fourth quarter alone, data-center revenue was $62.3 billion of a $68.1 billion total.[7] During 2026 the company's market value passed $5 trillion.[22] A business planned in a diner booth, selling a product named for a man the research universities would not hire, became one of the most valuable companies ever listed. That is the American innovation complex working the way it actually works: the origin is small, local and overlooked, and the capital arrives late and at the top.

§ 05 The build-up

The family on the payroll. *What the proxy says, and what it does not.*

A claim is circulating that Jensen Huang's children quietly took over Nvidia during the Blackwell build-up. I went to the filing. The record does not support a takeover, and what it does show is still worth a reader's attention.

Nvidia's proxy statement, filed with the Securities and Exchange Commission on May 12, 2026, discloses under related-party transactions that two of Huang's children work at the company. Madison Huang is a senior director of product and technical marketing for physical-AI platforms and received about $1.23 million in total compensation for the fiscal year. Spencer Huang, who joined in July 2022, leads product strategy for robotics software and received about $1.32 million.[9]

The same filing states that neither is an executive officer, that neither reports to their father, and that their pay was set without his involvement.[9] Neither sits on the board. So the accurate sentence is narrower than the rumor. Two of the founder's children joined during the run-up to Blackwell and now hold director-level roles in robotics and physical AI, the lines of business the company promotes as its next market. Talk of succession is press speculation. No filing I reviewed describes a succession plan, for them or for anyone else. A reader is entitled to find that worth watching. A reader is not entitled to call it a takeover.

§ 06 The diversion

Where the machines went. *Case by case, because there is no total.*

The question put to me was how much money was lost when Blackwell-class chips reached China through Malaysia and Singapore. The honest answer is that no government has published that number, and anyone quoting a single figure is adding things that cannot be added. What exists is a set of cases. Each has its own dollar amount, its own standard of evidence, and its own denial.

FIG. 5 — THE CASES, BY THE DOLLAR VALUE NAMED IN EACH NOT ADDITIVE · NO OFFICIAL TOTAL

Shanghai Century Litong

Nvidia purchases it reportedly facilitated, some of them restricted chips



$7.5B

PRESS REPORT The Information, unnamed sources. No charges.[11]

Megaspeed, Singapore

Nvidia chips bought over two years; nearly $2B installed in Malaysia and Indonesia



$4.8B

PRESS REPORT

New York Times, Oct 2025. Under investigation, no charges. Megaspeed and Nvidia deny diversion.[10,11]

US v. Liaw and two others

Servers with H200 and B200 GPUs, allegedly routed through a Southeast Asian front



$2.5B

CHARGED IN COURT Indicted March 2026. Liaw pleaded not guilty; trial set for March 2027.[12]

Singapore fraud case

Servers moved onward through Malaysia



$390M

CHARGED IN COURT Singapore prosecutors, 2025.[12]

US v. Lui / Earthmade Computer

Shipments through Malaysia and Singapore



$300M

CHARGED IN COURT Justice Department, autumn 2026.[12]

US v. Hsu network

H100 and H200 GPUs, October 2024 to May 2025



$160M

CHARGED IN COURT Justice Department, December 2025.[12]

CHARGED IN COURT a government filed the allegation and must prove it

PRESS REPORT journalism from records and unnamed sources, no charge filed

Amounts are the figures named in each case or report. They measure different things (purchases in the top two rows, allegedly diverted goods below) and may overlap, so they are not summed. A charge is an allegation; every defendant is presumed innocent.

Read the bars for what they are. Four are charges filed in court, where a government has put its name to the allegation and a defendant can answer it. In the largest, the March 2026 indictment of Super Micro co-founder Yih-Shyan "Wally" Liaw and two others, Liaw has pleaded not guilty, one co-defendant has pleaded guilty, Super Micro itself is not charged and says it was a victim, and the alleged intermediary has not been charged.[12] The two longest bars are press reports. They measure purchases, not proven diversion, and the companies named deny wrongdoing.

The Megaspeed file is the one that matters for the rest of this story. In October 2025 the New York Times reported that Megaspeed, a Singapore firm, had bought more than $4.8 billion of Nvidia chips in two years and installed nearly $2 billion of them in data centers in Malaysia and Indonesia that allegedly served Chinese customers remotely. A Commerce Department official reportedly found Nvidia equipment still in its boxes at a Malaysian site in 2024. American and Singaporean authorities confirmed they were investigating.[10,11] Megaspeed denies wrongdoing and says inspections by the Commerce Department's export bureau, by Malaysia's trade ministry and by Nvidia found no violations. Nvidia says it visited the sites and found no evidence of diversion.

According to Times reporting relayed by CNBC, the Commerce inquiry ran until August 2025 and was then moved to the Justice Department after outreach from Nvidia's general counsel. A spokesman for the export bureau said on the record that enforcement is conducted by the book. CBS News, citing two law-enforcement sources, reported that the Justice Department's National Security Division is investigating Megaspeed.[10,13] As of this writing I have found no charge against Megaspeed and none against Nvidia.

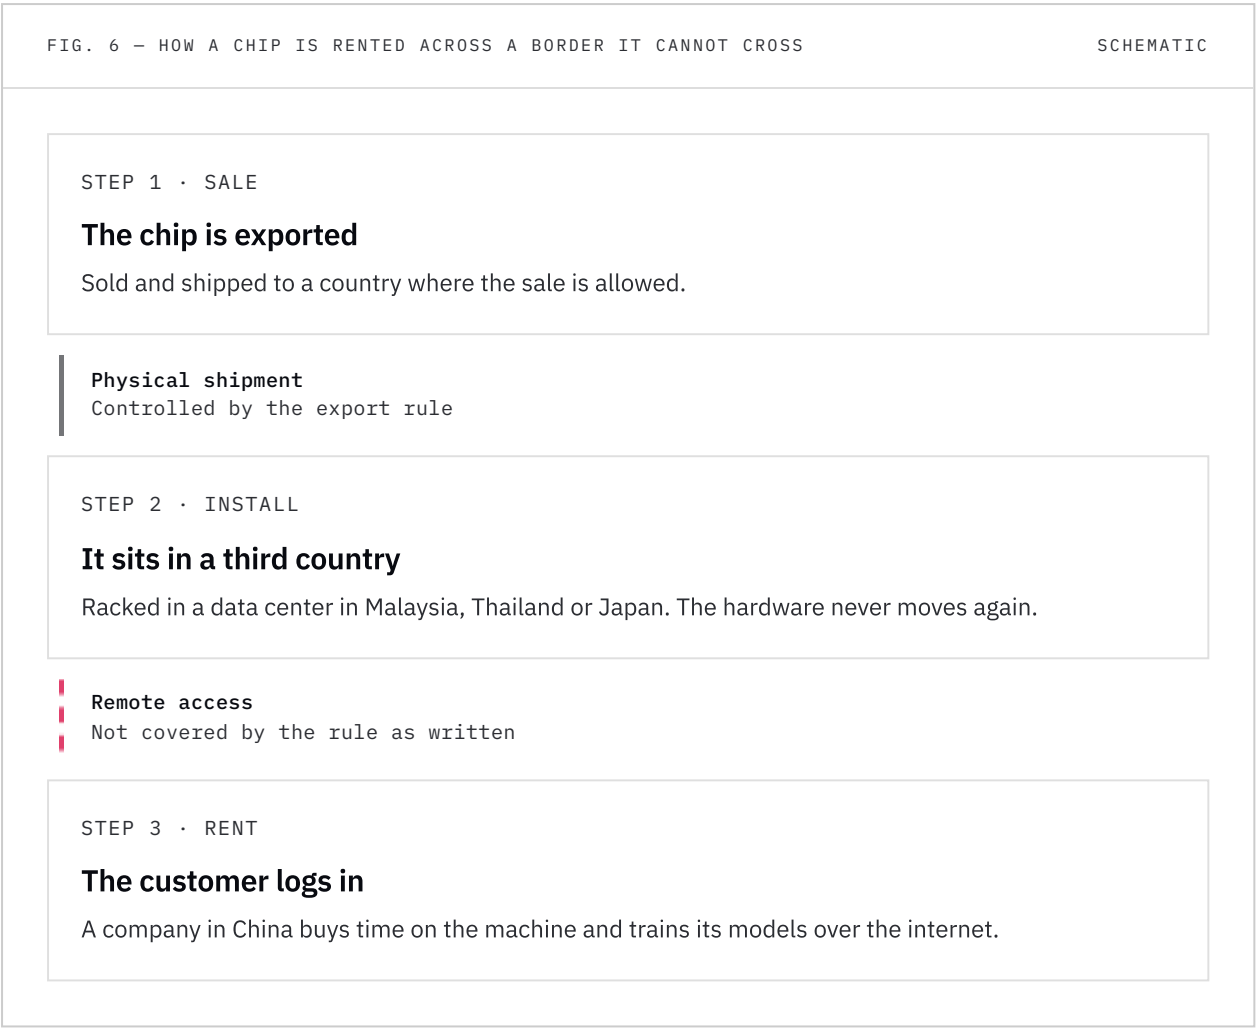
So what was lost? Not Nvidia's revenue. Every chip in every one of these cases was sold and paid for. If the allegations hold, the loss is the policy itself: the lead that the export rule was written to protect, spent a container at a time.

§ 07 Remote access

Renting the chip. *The rule covers the box, not the login.*

American export controls regulate where a chip is shipped. As written, they do not regulate who logs in to it. A Blackwell server lawfully installed in a data center in Southeast Asia can be rented over the internet by a customer in China, and no restricted hardware has crossed a controlled border.

FIG. 6 — HOW A CHIP IS RENTED ACROSS A BORDER IT CANNOT CROSS SCHEMATIC



A schematic of the practice described in the reporting below, not a diagram of any one company's network.[14]

CNBC reported in August and November 2026 that Moonshot AI, ByteDance, Alibaba and Tencent had reached restricted Nvidia compute through data centers in Thailand, Malaysia and Japan; that the White House accused Moonshot of using Blackwell-generation chips in Thailand; that ByteDance worked with a Singapore vendor on clusters in Malaysia; and that Moonshot sought Blackwell access in Europe.[14] These are press reports and one government accusation. I have seen no court finding on any of them.

Congress has noticed the gap. The Remote Access Security Act, which would extend export controls to remote use, passed the House 369 to 22 on January 12, 2026, and is pending in the Senate.[14] Until it or something like it becomes law, renting is the cheapest diversion there is, because under the current rule it may not be diversion at all.

§ 08 The scoreboard

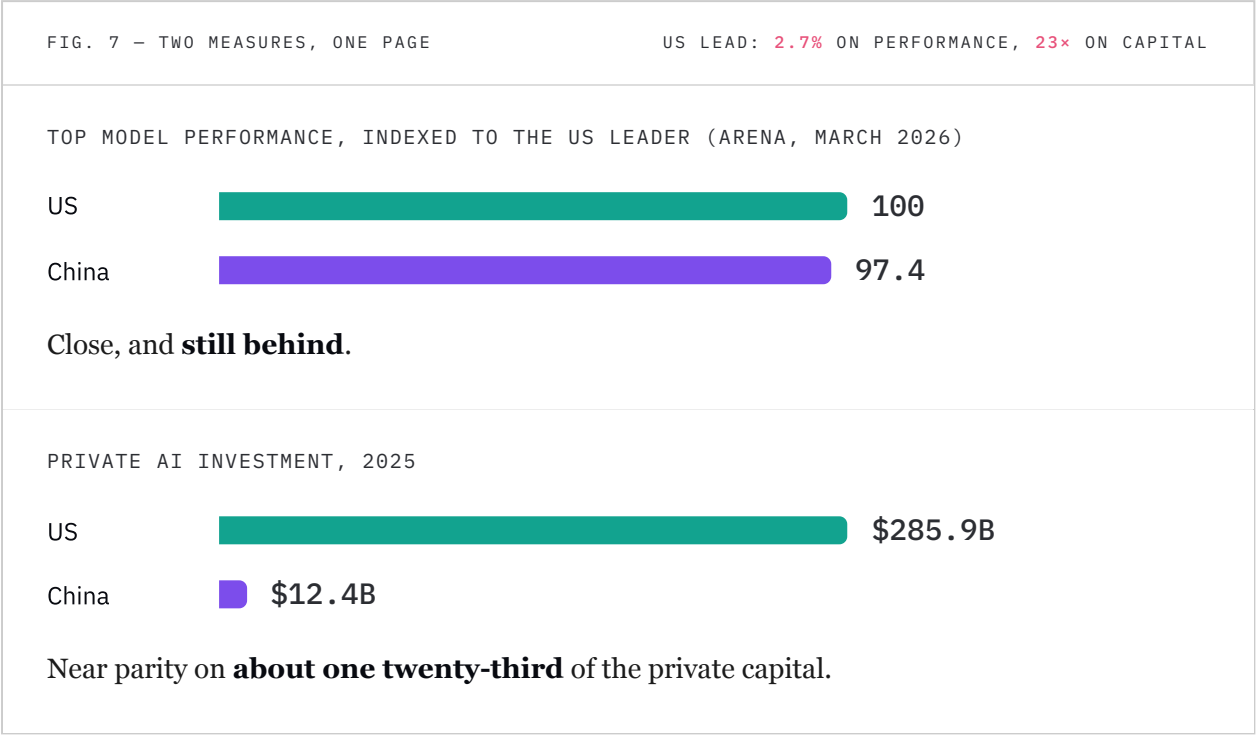
Did it put China ahead? *The measurement says no.*

The framing handed to me is that diverted American chips drove Chinese AI models past American ones. I looked for the evidence and could not support the sentence. Stanford's 2026 AI Index reports

that as of March 2026 the top US model led the top Chinese model by 2.7% on the Arena leaderboard. The gap has nearly closed, and DeepSeek's R1 briefly matched the best American model in February 2025. Nearly closed is not passed.[15]

FIG. 7 — TWO MEASURES, ONE PAGE

US LEAD: 2.7% ON PERFORMANCE, 23× ON CAPITAL



Left panel: the 97.4 is my conversion of Stanford's reported 2.7% lead into an index with the top US model at 100. Right panel: private investment as reported by the AI Index. Two separate scales, shown side by side.[15]

On whether smuggled Blackwells built those models: The Information reported in December 2025 that DeepSeek held roughly 2,000 to 2,300 Blackwell GPUs. Nvidia called the smuggling claims far-fetched. A 2026 report alleged that DeepSeek's V4 model was trained on smuggled Blackwell hardware, and DeepSeek denies it.[11,12]

The defensible sentence is this. Chinese laboratories got close on a fraction of the capital, and some unknown share of their compute is alleged to be American hardware they were not licensed to own or rent. The indefensible sentence is that they are ahead because of it. They are not ahead.

§ 09 The surveillance story

What was said, by whom, *in the order it became public.*

Everything in this section that comes from intelligence is, by the standard of this series, rumor. Not because it is false. Because no reader can inspect it. It was described to reporters by unnamed people briefed on surveillance, and no document has been published. I grade it that way and print every denial beside it.

FIG. 8 — THE PUBLIC RECORD, DATED

EACH ENTRY GRADED BY WHAT A READER CAN CHECK

Oct 9–10, 2025

ON THE RECORD

After the Times report on Megaspeed's purchases, **US and Singaporean authorities confirm they are investigating** the company.[10]

> Megaspeed denies wrongdoing. Nvidia says its own site visits found no diversion.

Late 2025 to spring 2026

RUMOR · INTELLIGENCE-SOURCED

The Times reports that a classified memo, drafted under Secretary of State Marco Rubio before the president's May trip to China, described Alice Huang, a former Megaspeed executive, as an asset of China's Ministry of State Security with a relationship to Jensen Huang, and that the president had been briefed earlier and did not act. Axios later reported that the Commerce Secretary was sent the memo.[11,23]

> President Trump says he had no prior knowledge. Rubio declined to comment. The memo has not been published.

Oct 29, 2026

RUMOR · INTELLIGENCE-SOURCED

The Times reports that US intelligence determined Alice Huang, a Chinese citizen and former television reporter who is no relation to the chief executive, **was passing information to Beijing** and had a close relationship with him. The sourcing is people briefed on surveillance.[11,13]

> Nvidia says the interactions were not unusual, and Huang told Bloomberg they were entirely normal. China rejects the report. I found no public response from Alice Huang in the coverage I reviewed, and she has not been charged with anything.

Nov 4, 2026

ON THE RECORD

Congress gets involved. Rep. Brian Mast, chairman of the House Foreign Affairs Committee, with the chairmen of the Select Committee on the Chinese Communist Party and the Intelligence Committee, writes to Jensen Huang asking for a joint closed-door briefing on Nvidia's protections against foreign espionage and on his interactions with the person the letter calls a publicly reported intelligence asset. The requested date is November 18, with a reply due November 11.[16]

> I read the letter. It rests on press reports, presents no evidence of its own, and contains no dollar figure. Mast later told ABC News he would subpoena Huang if there was no response.[17]

Nov 5, 2026

ON THE RECORD

The letter is made public. **Nvidia closes down 9.26%**, on a day the whole chip sector falls and coverage points to fears of new regulation after the midterm elections.[20]

Nov 8, 2026

ON THE RECORD

Nvidia's board opens an independent review. A committee of independent directors, with outside counsel, will examine reports that the chief executive was targeted by a foreign intelligence effort. The board says it has no reason to believe national security or company secrets were compromised, expresses confidence in Huang, and says he remains chief executive and "fully supports the review."[19] The same day, in a message reported by several outlets, Huang acknowledged having been unfaithful in his marriage and apologized. He named no one.[22]

Nov 9, 2026

ON THE RECORD

China denies it. Foreign ministry spokesperson Guo Jiakun calls the account baseless and a smear.[18] Nvidia falls nearly 9% after the open and closes about 4.27% lower. Huang withdraws from speaking engagements. Senators Elizabeth Warren and Jim Banks call for answers.[20,22]

Nov 10, 2026	PRESS REPORT Fortune reports that the board's committee has retained a former acting US attorney and a New York law firm.[12] > Single-sourced. The board's own statement names no counsel, so I leave the names out.
Ahead	ON THE CALENDAR November 11, the reply the committees asked for. November 17, Nvidia's third-quarter earnings. November 18, the requested briefing in the Rayburn building.[16]

Two notes on the company's posture. First, it is accurate to say Nvidia supports an independent investigation: the board opened one and the chief executive says he backs it. It is also accurate that the statement gives no timeline, names no outside counsel, and makes no commitment to publish what it finds. "Independent" here means independent directors of the same board.[19]

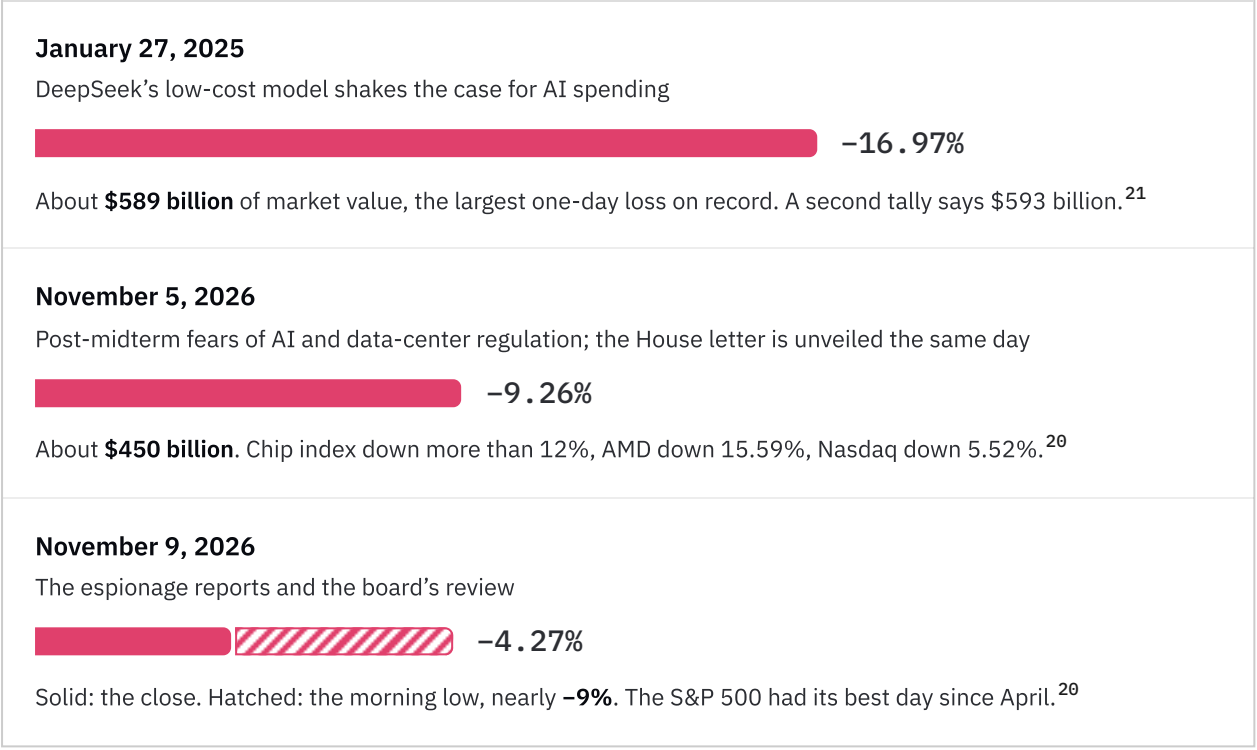
Second, the context a reader should hold alongside it. Nvidia reported a record $2.15 million in federal lobbying for the third quarter of 2026, gave $1 million to the 2025 inaugural fund, and gave $25 million toward the White House ballroom.[22] None of that is evidence of anything in this file. It is the setting in which an export-control inquiry moved from one department to another, and it belongs on the page.

§ 10 The market

What the stock did, *and what it was reacting to.*

The claim is that Nvidia's market value dropped because of spying. Partly. Sort the three largest recent single-day falls by their reported cause and the spy story is the smallest of the three.

FIG. 9 — NVIDIA, THREE BAD DAYS — SINGLE-DAY CLOSE, WITH THE REPORTED CAUSE



Percentages and dollar losses as reported in the cited market coverage. I found no reliable published dollar figure for the November 9 move and do not estimate one.

The number in circulation is that the stock fell 9% on November 9. It did, for part of a morning. Shares were down nearly 9% shortly after the open and recovered to close about 4.27% lower, on a day the broad market rallied. The 9.26% close belongs to November 5, when the entire chip sector fell, AMD dropped further than Nvidia, and the coverage attributed the move to regulation fears after the election. The House letter was unveiled that same day, so the two causes cannot be cleanly pulled apart, and I do not try.

The fair reading: the espionage story cost Nvidia a mid-single-digit share of its value on a day the market rose, which makes it a company-specific penalty and a real one. It is not the $450 billion. It is nowhere near the DeepSeek day. None of this is a view on the stock.

§ 11 The ledger

Confirmed, alleged, rumor. *And what the record does not support.*

The rule of this dispatch is that a claim is graded by what a reader could check for themselves, not by how many times it has been repeated.

CONFIRMED — A document, a filing, a named official or a closing price.

Blackwell wrote to more than a hundred HBCUs, received three offers, and did foundational work at Howard.

SOURCE

Two standard biographies, which differ only on 104 or 105.[1,2]

Nvidia named its 208-billion-transistor platform for him.

SOURCE

Nvidia's own announcement, March 18, 2024.[5]

Two of Jensen Huang's children work at Nvidia, at about $1.23 million and $1.32 million a year. Neither is an executive officer.

SOURCE

SEC proxy statement, May 12, 2026.[9]

The United States and Singapore are investigating Megaspeed.

SOURCE

Officials of both governments, October 2025.[10]

Three House chairmen, led by Rep. Brian Mast, asked Huang for a closed-door briefing on November 18.

SOURCE

The letter itself, dated November 4, 2026.[16]

Nvidia's board opened a review by independent directors, and Huang says he supports it.

SOURCE

Board statement, November 8, 2026.[19]

China denies the espionage account.

SOURCE

Foreign ministry spokesperson, November 9, 2026.[18]

Nvidia closed down 9.26% on November 5 and about 4.27% on November 9.

SOURCE

Market coverage of both sessions.[20]

The House passed a bill to bring remote access under export controls, 369 to 22. The Senate has not.

SOURCE

Congressional record as reported.[14]

ALLEGED Charged in court or reported from records and unnamed sources. Unproven.

About $2.5 billion of GPU servers were diverted through a Southeast Asian front.

ANSWER

US indictment. Liaw pleaded not guilty. Super Micro is not charged and says it was a victim.[12]

Three further schemes moved about $390 million, $300 million and $160 million of hardware.

ANSWER

Charges in Singapore and the United States. Not yet tried, as far as I can establish.[12]

Megaspeed's Malaysian and Indonesian data centers served Chinese customers.

ANSWER

Denied by Megaspeed. Nvidia says it inspected and found no diversion.[10,11]

A Shanghai partner facilitated about $7.5 billion of Nvidia purchases, including restricted chips.

ANSWER

The Information, unnamed sources. I found no on-record response in the coverage I reviewed.[11]

Chinese firms rented restricted Nvidia compute in Thailand, Malaysia and Japan.

ANSWER

CNBC reporting and a White House accusation against one firm. No court finding.[14]

Nvidia's general counsel contacted Commerce before the Megaspeed inquiry moved to the Justice Department.

ANSWER

Unnamed sources. The export bureau says enforcement is by the book.[10]

DeepSeek trained on smuggled Blackwell hardware.

ANSWER

Denied by DeepSeek. Nvidia called the smuggling claims far-fetched.[11,12]

○ RUMOR · INTELLIGENCE-SOURCED

Described by unnamed people briefed on US surveillance. No document published.

A former Megaspeed executive passed information to Beijing and was an asset of China's state security ministry.

ANSWER

Denied by China. No charge. No public evidence.[11,18]

She had a close relationship with Nvidia's chief executive that was part of an intelligence effort.

ANSWER

Nvidia and Huang say their interactions were normal. The board says it has no reason to believe secrets were compromised.[19]

A classified memo saying so reached senior officials in spring 2026, and the president was briefed earlier.

ANSWER

The president denies prior knowledge. Rubio declined to comment.[11, 23]

NOT SUPPORTED In circulation, and contradicted or unbacked by the record I could find.

"Chinese models passed American models because of diverted chips."

RECORD

The top US model led by 2.7% in March 2026.[15]

"The stock fell 9% on November 9 because of spying."

RECORD

Nearly 9% intraday, about 4.27% at the close. The 9.26% close was November 5.[20]

Any single dollar total for chips diverted through Malaysia and Singapore.

RECORD

No government has published one. The cases cannot be summed.

"Huang's children took over Nvidia."

RECORD

Two director-level employees. Not officers, not directors.[9]

> *The mathematics converged in 1965. The facts have not converged yet, and a reporter's job is to say which is which.*
>
> — FIELD DISPATCHES

↓ The series thruline

The gap between where value is created and where capital lands is an infrastructure problem. David Blackwell proved that if you keep applying an honest update, the error shrinks. He proved it after being turned away by the institutions that trained him, from a desk a Black college gave him when it had three of the only offers in the country to make. The chips that carry his name are now the most contested objects in the world economy: bought through intermediaries, rented through logins, argued

over in letters from Congress. The name on the chip is a record of where American invention actually comes from. What happens to the chip is a test of whether the country can keep track of what it makes. This dispatch will be updated as the dates in Fig. 8 arrive.

§ R References — How this was reported

Sources *& method.*

This paper was reported from public documents and published coverage through November 10, 2026. I read the House letter, the Nvidia board statement, the company's SEC filings and its product announcement directly. I did not have direct access to the New York Times, Financial Times, Axios or The Information articles; their reporting is cited through the secondary coverage listed below and is attributed to them by name in the text. Claims sourced to intelligence are labeled as rumor throughout. Figures 1, 2 and 6 are illustrations of cited facts, not new measurements.

1. Wikipedia. "David Blackwell." Biography, appointments and honors; gives 104 letters and Berkeley chairmanship from 1957. en.wikipedia.org.

2. MacTutor History of Mathematics Archive, University of St Andrews. "David Harold Blackwell." mathshistory.st-andrews.ac.uk.

   Gives 105 letters and the chairmanship from 1956. Where the two biographies differ, both figures are reported.

3. Blackwell, D. "Conditional Expectation and Unbiased Sequential Estimation," *Annals of Mathematical Statistics* 18(1), 1947, 105–110. "Equivalent Comparisons of Experiments," *Ann. Math. Statist.* 24(2), 1953, 265–272. With M. A. Girshick, *Theory of Games and Statistical Decisions*, Wiley, 1954. "Discrete Dynamic Programming," *Ann. Math. Statist.* 33(2), 1962, 719–726. "Discounted Dynamic Programming," *Ann. Math. Statist.* 36(1), 1965, 226–235.

4. "Blackwell's contraction mapping theorem." en.wikipedia.org. Bertsekas, D. *Abstract Dynamic Programming*. mit.edu.

   The link from the discounted framework to reinforcement learning is the textbook one. No source consulted traces a specific modern algorithm to the 1965 paper.

5. NVIDIA. "NVIDIA Blackwell Platform Arrives to Power a New Era of Computing," March 18, 2024. nvidianews.nvidia.-com.

   Performance multiples are the vendor's own.

6. Associated Press. "How Nvidia became an AI giant." abcnews.go.com.

7. NVIDIA. "Financial Results for Fourth Quarter and Fiscal 2026," February 25, 2026. sec.gov.

8. United Negro College Fund figures as reported by TheGrio, August 14, 2026. thegrio.com. Georgetown Center for Security and Emerging Technology, "Expanding the Reach." cset.georgetown.edu.

   The 28% and 19% shares use different definitions and years and are not directly comparable.

9. NVIDIA. Definitive proxy statement (DEF 14A), filed May 12, 2026. sec.gov.

10. CNBC. “Singapore, US investigate Nvidia client Megaspeed,” October 10, 2025. cnbc.com. CNBC, Megaspeed explainer, November 9, 2026. cnbc.com.

    Carries the New York Times reporting on Megaspeed’s purchases and on the transfer of the inquiry.

11. TechCrunch. Timeline of the reporting, November 10, 2026. techcrunch.com.

    Secondary source for reporting by the New York Times, the Financial Times and The Information, none of which I read directly.

12. Fortune. “2 scandals, 1 destination,” November 10, 2026. fortune.com.

    Source for the court cases as summarized here. Its report on the board’s outside counsel is single-sourced.

13. CBS News. Report on the Justice Department inquiry. cbsnews.com.

14. CNBC. Remote access to restricted chips, August 19, 2026. cnbc.com. CNBC, Moonshot AI, November 5, 2026. cnbc.-com.

15. Stanford HAI. *AI Index Report 2026*, technical performance chapter. hai.stanford.edu.

16. US House Committee on Foreign Affairs. Letter to Jensen Huang, November 4, 2026. Letter (PDF). Press release.

    Signed by Chairmen Brian Mast, John Moolenaar and Rick Crawford. The letter cites press reporting and makes no independent finding.

17. ABC News. Report on a possible subpoena. abcnews.com.

18. Associated Press via PBS NewsHour. China’s response, November 9, 2026. pbs.org.

19. NVIDIA. “Statement from NVIDIA Board of Directors,” November 8, 2026. nvidianews.nvidia.com. Anadolu Agency report. aa.com.tr.

20. Market coverage. 24/7 Wall St via Yahoo Finance on November 5. finance.yahoo.com. The Motley Fool on November 9. fool.com. Al Jazeera. aljazeera.com.

    One outlet puts the November 5 loss at $447 billion rather than $450 billion.

21. Bloomberg. “Nvidia’s $589 Billion DeepSeek Rout Is Largest in Market History,” January 27, 2025. bloomberg.com. Reuters figure of $593 billion via Business Standard. business-standard.com.

22. CNBC. Lobbying and political giving, November 9, 2026. cnbc.com. The congressional request, November 5. cnbc.com. The board review, November 9. cnbc.com.

23. Axios. Report on the memo's distribution, November 10, 2026. axios.com.
Not read directly; cited as described in other coverage.