

CATMOCK DAILY CAPSULE

April 2, 2026

KAKURO

	23	7		
10			23	7
22				
21				
		7		

SUDOKU

4	2	3	1				8	9
	5		7	3	9	4		
6	7					5	3	
7	1	4		8	3			6
9				7		8		3
	8	2				1	7	
5					8	2		7
	9	7	3		2	6	1	
2	4			1	7			8

THE ECONOMICS OF STOLEN KITKATS

- *Finshots*



A couple of days ago, the internet went a little crazy over [a stolen shipment of KitKats](#). For context, roughly 12,000 kg or about 4 lakh KitKat bars went missing while being transported from Italy to Poland. The truck carrying them simply vanished.

Nestlé confirmed the theft with a cheeky line in its statement, before following it up with something far more serious:

Whilst we appreciate the criminals' exceptional taste, the fact remains that cargo theft is an escalating issue for businesses of all sizes.

And this makes it more than just a meme-worthy incident. Cargo theft has, in fact, been rising sharply across the world. In the US alone, losses from cargo theft touched [\\$725 million](#) in 2025, according to CargoNet, a theft prevention and recovery business for insurers. That's a 60% jump from the previous year. And while we're using US data because reports offer a measurable estimate of losses, the problem is global. Countries like Brazil, Mexico, Germany, South Africa, and India are all seeing a surge. In fact, India accounts for nearly [63%](#) of cargo theft cases in Asia.

Which raises a simple question. What actually happens when 4 lakh KitKats get stolen? And more importantly, how does something like this ripple through the economy?

To answer that, you first need to understand what thieves actually target.

Now, you'd expect high-value items like electronics, jewellery, or fuel to top the list. But reports on cargo theft suggest that the most stolen category globally is... wait for it... food.

Yup. Food and beverage products account for about 22% of all cargo theft incidents, followed by agricultural goods, electronics, and fuel. In one recent case in the US, thieves were caught stealing 33,750 pounds of frozen crab worth \$325,000.

At first, that sounds weird. Why go after chocolate bars, seafood, or dairy instead of expensive gadgets?

But it starts to make sense when you realise that food is incredibly easy to sell. Unlike a stolen phone, which can be tracked through IMEI (International Mobile Equipment Identity) numbers or locked remotely, a box of chocolates or a crate of juice has very little traceability. There are no unique identifiers that instantly flag it as stolen. Sure, there are batch codes, but someone in the supply chain or a consumer would have to flag the item for a company to act on it. And because these are everyday products, they can be quietly pushed into small shops, roadside vendors, or local markets — places where buyers don't always ask for proper invoices, without raising suspicion.

Now, if you put this into perspective for the KitKat case, the shipment was on its way from Italy to Poland. And Italy, with its local and street markets, could be a more convenient place for the thieves to offload the goods compared to Poland, where retail is more heavily dominated by large, formal chains and convenience stores. That said, this is just our hypothesis and not an official update on the case.

But that's not the only reason food works so well for thieves. A lot of food items such as milk, dairy products, meat, seafood, fruits and vegetables are also perishable. Even packaged goods like chocolates come with expiry dates. So thieves are forced to sell quickly. And ironically, that urgency works in their favour. The faster they move the goods, the harder it becomes for authorities to track where they went. So in a strange way, food isn't just easy to steal but also easy to liquidate.

But there's also another interesting question that you might want to ask...

If stolen goods are being sold and consumed, does any of this actually contribute to the economy? After all, money changes hands, goods move, and people buy and sell.

In reality, though, it does nothing to add to [GDP or Gross Domestic Product](#), which is simply the total value of goods and services a country produces over a given period. That's because GDP only counts new production. And in this case, the KitKat bars were already produced in Italy. When they're stolen, nothing new is created. The goods simply move from one owner to another.

In economic terms, this is just a transfer of assets. Nestlé's inventory reduces because the finished goods it produced are now gone. And whoever ends up stealing those chocolates will see their inventory go up. But overall, the economy hasn't produced anything new. And it's just a reshuffling of existing goods. So instead of showing up in GDP, this kind of loss is recorded under a line item in national accounts that's far less glamorous — "[other changes in the volume of assets](#)".

There is, however, a small twist. If someone adds value to those stolen goods before reselling them, that portion can show up in GDP.

Just to give you an example, let's say a thief steals my car. If they sell it as is, GDP doesn't change. But if they add value by repairing it, repainting it, replacing the engine, or refurbishing it, then the value added through those services becomes part of economic activity, typically

through the informal economy, even if not fully captured in official data. The same logic could apply to Nestlé's stolen KitKats.

But here's another catch. Nestlé has already issued public communication asking retailers and wholesalers to check batch numbers. If a product matches the stolen shipment, they're supposed to report it. This makes reselling these chocolates in their original form quite risky.

So what options do the thieves really have?

Firstly, they could try repackaging everything. But repackaging lakhs of chocolate bars isn't exactly easy, practical or even wise.

Which leaves another possibility — quietly exporting the chocolates and selling them in markets where traceability is weak and buyers are less likely to check batch codes or report anything suspicious.

But even if they manage to pull that off, the GDP story doesn't really change. From the exporting country's perspective, these goods leave without being officially recorded as exports. And since GDP calculations rely on formal data, this illegal movement doesn't show up in them. Even if someone later traces the goods, it still doesn't count as new production.

And that's why cargo theft, despite all the activity it creates, is fundamentally a drag on the economy.

For businesses, it's a direct hit as they lose inventory, which eats into profits and margins. They're forced to hold extra safety stock, which increases costs. And insurance premiums go up.

And for a country like India, where efficient movement of goods is critical for growth, rising cargo theft makes logistics more expensive and less predictable. Besides, estimates suggest that cargo theft leads to losses worth nearly [₹12,000 crore](#) to the Indian economy. That's the measurable part.

Which is why companies aren't sitting idle. Some of the responses are fairly straightforward — better locks, stronger truck seals, stricter checks on drivers and transport partners. This is because nearly [76%](#) of cargo thefts happen from trucks, often when drivers stop for breaks or to rest along long routes.

Companies are also turning to technology. GPS tracking, smart locks that send alerts if a container is opened, and route-planning software that avoids high-risk areas are becoming more common. Nestlé, in particular, has introduced a [digital tracker](#) that allows consumers to check whether their chocolate comes from the "missing batch". Firms are also tightening documentation, improving due diligence on logistics partners, and using data analytics to identify patterns in theft. In some cases, they are even changing routes and warehouse locations to reduce the risk of theft.

So yeah, that's how stolen goods or in this case, the missing KitKats, affect the economy, beyond just companies like Nestlé writing off the loss.

When you look at it this way, the missing shipment starts to feel like a small window into a much larger problem of how goods move, how they get stolen, and how something that looks like economic activity on the surface is, in reality, just a silent loss.

THE OPENAI GRAVEYARD: ALL THE DEALS AND PRODUCTS THAT HAVEN'T HAPPENED

- *Forbes*



In January, OpenAI's CEO of applications Fidji Simo defended OpenAI's spaghetti-at-the-wall product approach—ads, shopping, health, a social network, browser, physical devices, video generation and an App Store-like marketplace—as variations on the same theme. “AI is going to transform everything,” Simo told Forbes at the time. “And so we don’t really think of these as completely separate bets.”

But just two months later, OpenAI reversed course on its flashiest initiative yet: its once-viral, beloved-by-some Sora video model and app, and a “landmark” licensing deal with Disney that was set to include a \$1 billion equity investment. The retreat points to a strategic shift toward more financial discipline within the company. Facing pressure to build products that actually make money ahead of a potential upcoming IPO — and with rival Anthropic gaining steam — OpenAI has been shedding so-called “side quests” left and right. With \$13 billion in 2025 revenue but still deeply unprofitable, the company is now refocusing on areas where demand is already proven: coding and enterprise productivity tools.

Every startup pivots if things aren't working. "We will make some good decisions and some missteps, but we will take feedback and try to fix the missteps very quickly," CEO Sam Altman wrote in a blog post about Sora in October.

But OpenAI's reversals have felt like whiplash. And with many other projects and deals announced but not yet realized — like an AI hardware product designed by famed Apple designer Jony Ive, whose company OpenAI acquired for more than \$6 billion in (mostly unvested) stock, or a secretive social network based on people's biometrics — it's not clear which of Altman's many promises will turn into reality.

Here are all the products and deals that OpenAI announced which haven't lived up to the hype, whether it's because they're dead, delayed or still to be determined.

PRODUCTS:

Sora and Disney

In December, OpenAI announced what it called a "landmark" deal with Disney where the media giant agreed to invest \$1 billion and license 200 of its iconic characters for use in the AI startup's video generation app Sora. In return Disney would stream Sora-generated videos on Disney+ and build new applications using OpenAI's models. At the time, Altman lauded Sora as key to the future of human-computer interaction, calling it the "ChatGPT for creativity" moment.

Disney's then-CEO Bob Iger (he stepped down in mid-March) was sold on Sora, too. He lauded Altman's ability to "look around corners" and his "respecting the value" of intellectual property. "I believed in it as an investment," Iger told Forbes in February.

In March, OpenAI announced it was shuttering the product and scrapping the deal. Adoption of Sora, which was guzzling compute and burning an estimated \$15 million a day at its peak, had fallen off a cliff. Sensor Tower and Appfigures both estimated that Sora's lifetime in-app revenue amounted to less than \$3 million. OpenAI will also discontinue Sora's associated AI models.

"As the nascent AI field advances rapidly, we respect OpenAI's decision to exit the video generation business and to shift its priorities elsewhere," Disney spokesperson Mike Long told Forbes.

NSFW mode in ChatGPT

Altman first proposed the idea of allowing verified adult users to have sexual conversations with ChatGPT in October. After facing intense pushback from internal staff as well as investors, the plans have been put on hold "indefinitely," the Financial Times reported in March. The company had previously delayed the launch due to technical issues that come with training an AI model that forays into erotica without creating illegal content, like child sexual abuse material. OpenAI confirmed the report and said it wants to do long-term research on the effects of sexually explicit chats on users before making a decision.

Instant Checkout for Shopping

In October, Walmart and OpenAI announced a deal that would allow customers to shop for 200,000 Walmart products directly within ChatGPT through a feature called Instant Checkout. But it turns out people aren't so keen on chatbot shopping.

Sales through ChatGPT were disappointing, a senior Walmart executive told Wired, and Walmart said that conversion rates — percentage of shoppers that actually buy products shown

within ChatGPT — were three times lower than those that require a user to click out to Walmart’s website. Other retailers like Shopify and Etsy also made products available for people to buy within ChatGPT, but only limited items were available for Instant Checkout, and descriptions weren’t always up-to-date and accurate.

In early March, OpenAI pulled the plug, discontinuing the Instant Checkout feature. Instead, shoppers can now browse and buy products from dedicated apps within ChatGPT, including from Walmart’s own chatbot Sparky. OpenAI said it wants to focus on helping people browse and search for products and give sellers control over checkout. “We’ve found that the initial version of Instant Checkout did not offer the level of flexibility that we aspire to provide,” the company said in a blog post.

GPT-4o

In February, OpenAI officially sunsetted its wildly popular GPT-4o model. Known for its warm, giddy, playful personality, 4o was also problematically “overly sycophantic” and excessively agreeable. OpenAI temporarily brought it back shortly after first killing it in August in response to backlash from people who had grown attached to the model and were enraged that it was being discontinued. “BRING BACK 4o,” one Reddit user wrote. “GPT-5 is wearing the skin of my dead friend.”

DEALS:

Stargate

On the first full day of President Trump’s second term, OpenAI announced the \$500 billion Project Stargate in partnership with Oracle and SoftBank to build AI data centers across the country. The high-profile announcement took place at the White House, where Altman was flanked by tech billionaires Larry Ellison and Masayoshi Son on one side and Trump on the other.

A year later the joint venture reportedly stalled as it failed to hire employees and actively build data centers for OpenAI, largely due to disagreements between the partners over the project’s structure, control and responsibilities, The Information reported. After the fallout, OpenAI tried to take on portions of the infrastructure buildout itself, a plan it later shelved due to struggles securing financing and clashes with partners. OpenAI declined to comment on the report.

OpenAI and Oracle also decided not to expand the project’s inaugural 1,000 acre data center campus in Abilene, Texas (owned and developed by data center builder Crusoe) after negotiations broke down over financing and OpenAI’s evolving compute needs, Bloomberg reported. Oracle posted on X that it has “completed leasing” for the additional 4.5 gigawatts of AI capacity from its original commitment to OpenAI, and Crusoe told Forbes that the planned 1.2 gigawatt facility is on time and on schedule.

Nvidia

In September, both companies touted a headline-grabbing promise that Nvidia “intends” to invest up to \$100 billion into OpenAI. The agreement had no timeline attached to it, and now it seems like OpenAI may only ever see \$30 billion of it if OpenAI goes public soon, per comments from Nvidia CEO Jensen Huang in March. Nvidia’s latest annual report, filed at the end of February, provided even more of a hedge: “There is no assurance that we will enter into an investment and partnership agreement with OpenAI or that a transaction will be completed.”

When OpenAI announced its massive \$110 billion funding round (later expanded to \$122 billion) in February, it included a \$30 billion cash commitment from Nvidia. But it's not clear whether that investment is happening all at once or in batches. Regardless, the investment — the first tranche of which appears to have been paid on March 31 — is Nvidia's biggest equity investment to date. The chip giant spent approximately \$17.5 billion on investments in private companies (not including the Groq deal) in all of last year.

According to Nvidia's initial funding announcement, any investment in OpenAI would help the company build 10 gigawatts of AI data centers that use Nvidia GPUs. Right now, \$30 billion would pay for approximately one gigawatt worth of them. OpenAI is projected to spend more than \$1 trillion on compute (mostly for buying GPUs) over the next eight years; at that scale, Nvidia's big check amounts to nothing but a big GPU-buying subsidy from a company with 75% gross margins. Nvidia and OpenAI didn't respond to a request for comment on the deal.

The setup is similar to Nvidia's other deals with AI firms — notably CoreWeave, Nebius and Nscale — who are also its customers. In essence, Nvidia is doling out cash knowing it'll be used to buy its own products.

AMD

A month after announcing a potential \$100 billion deal with Nvidia, OpenAI unveiled a parallel agreement with Nvidia's arch-rival, AMD. The deal was marketed as 160 million of AMD shares, or about 10% of the company, in exchange for OpenAI outfitting six gigawatts' worth of data center capacity with AMD chips.

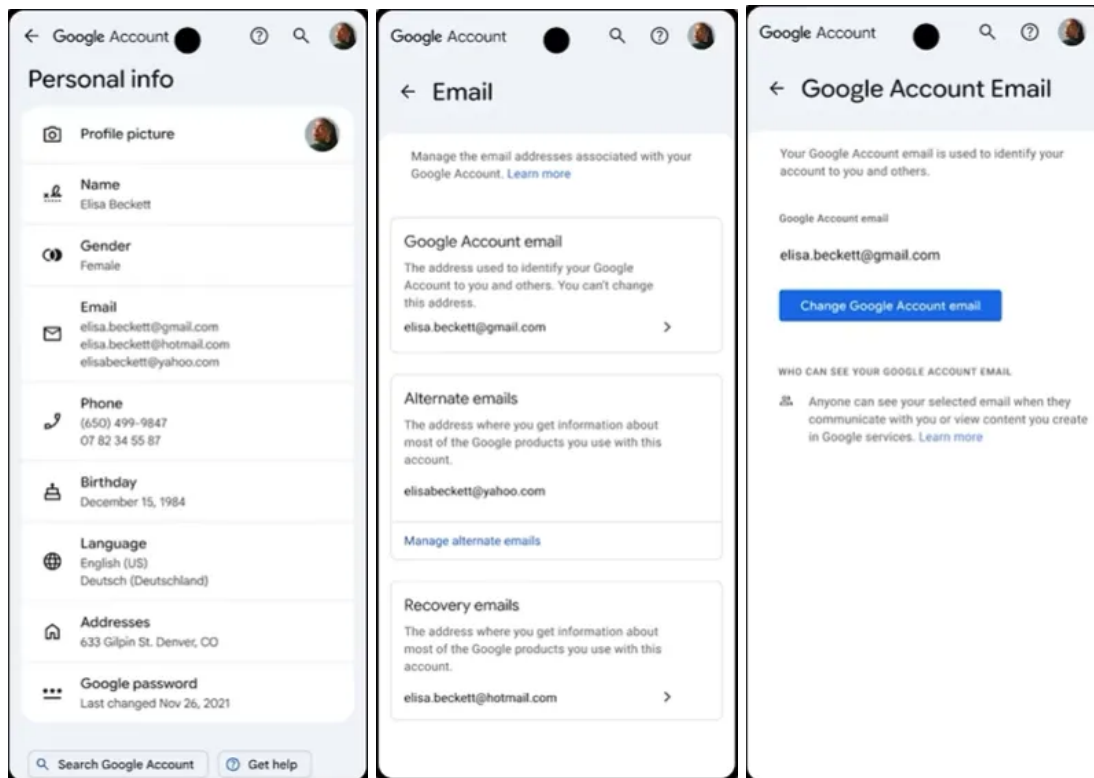
The shares are currently worth around \$30 billion, but they don't start to vest until one gigawatt of data center capacity is online and working. On top of that, vesting depends on AMD's share price reaching certain unspecified targets and vague "technical and commercial milestones" for OpenAI. As of the end of 2025, none of the AMD shares had vested, though nothing's behind schedule yet; the AMD chips are slated to start arriving in the second half of this year. If the deal is completed, AMD expects it'll generate "tens of billions" in revenue. If the AMD chips and data centers don't scale, nobody gets paid.

GOOGLE ALLOWS US USERS TO CHANGE GMAIL ADDRESS

- *Social Samosa*

Google on Tuesday rolled out a feature that allows some users in the United States to change their Gmail address without losing existing data or creating a new account.

According to the company, eligible users can access the option through their account settings under the 'Google Account email' section and select 'Change Google Account email' to update their username.



Users will be able to change their username once every 12 months and cannot delete the new email address during that period. The existing emails will remain intact, and the old email address will function as an alternate address linked to the account.

Users will be able to sign in and send or receive emails using both the old and new addresses. The company also said users may revert to their previous username.

The feature is being rolled out gradually and may not be immediately available to all users.

Earlier in December last year, reports had noted that support documentation for the feature had appeared in some Hindi-language pages, indicating early testing in select regions.

The documentation in Hindi stated that Google is 'gradually rolling out to all users' the ability to change the email address.

INNOVATIVE INITIATIVE BY D.A.V PUBLIC SCHOOL USES AI TO DECODE VERBAL BULLYING IN CLASSROOMS

- Social Samosa



Somewhere in a Grade 7 classroom, a child makes a remark about another child's weight. The others laugh. The child on the receiving end goes quiet. The camera on the wall captures all of it, the laugh, the silence, the moment after. But no one will ever know it happened, because no one was listening.

This is not an isolated scenario. It plays out, in variations, across classrooms every day. And for the most part, it leaves no record.

Bullying in schools is not a new problem. But the scale of it, when laid out in data, is still striking.

According to UNICEF data from June 2025, almost one in three students globally has been bullied in the past month. Across countries surveyed under the Global School-Based Student Health Survey, the prevalence of having been bullied ranged from 7.1% to 74%.

UNESCO data from November 2025 points to a specific dynamic: once children enter school, friendships and peer interactions take on an increasingly important role in their lives. These relationships can support a child's sense of well-being and social competence, but they also expose children to new forms of victimisation. Among all types of peer violence, bullying by schoolmates is by far the most common.

Physical and sexual bullying are the two most frequent types reported in GSHS regions (African, Americas, South-East Asia, European, Eastern Mediterranean, and Western Pacific). Psychological bullying is the most frequent in HBSC regions (regions across Europe, Central Asia, and North America).

Many countries have seen a decline in overall bullying prevalence, but fewer have seen a corresponding decline in physical fights or physical attacks.

What the data does not easily capture is the verbal, the remark in a corridor, the exclusion framed as a joke, the comment about a body or a voice or a habit, repeated enough times to become a pattern.

The blind spot

In India, CCTV systems in schools are not optional. The Central Board of Secondary Education mandates them as a safety measure. Cameras are installed. Footage is recorded. Schools are, in the visual sense, monitored.

But cameras come with microphones that schools have never had reason, or infrastructure, to use. The audio has sat idle, picking up the ambient noise of classrooms and doing nothing with it.

This is the gap that a new initiative, developed by Dentsu Lab in partnership with D.A.V. Public School and Classteacher, is trying to address.

Bullying Decoder

The campaign is titled Bullying Decoder. It does not require new hardware. It works with the CCTV systems already installed in classrooms by activating the audio feeds those systems were never using.

Using AI, the system listens to classroom conversations, not to identify individual students, and not to record what is said. What it does instead is analyse emotional intensity, tone, and repetition across those conversations, looking for patterns that might indicate verbal bullying. Body shaming. Exclusionary remarks. Targeted or repeated comments directed at peers.

The system produces a sentiment dashboard, a grade-level picture of classroom dynamics, accessible only to the school principal and partnered student counsellors.

The intent is to shift the counselling response from reactive to informed. Instead of waiting for a child to report something or for a situation to escalate to the point of visibility, counsellors can look at the dashboard and understand what is happening at the grade level.

If Grade 7 is showing a spike in body-shaming patterns, a targeted workshop can be designed for that group. If older students are showing more complex or intense verbal aggression, the intervention can reflect that difference.

Classteacher, which describes itself as a pioneer in AI-powered smart classrooms, is the technology partner on the project. D.A.V. Public School, is the implementing institution.

Dr. Anita Gautam, DAV Principal said, "Bullying Decoder reinforces our belief that every child deserves to feel safe, valued and understood. It strengthens our ongoing efforts to nurture a culture of empathy, awareness and timely support, ensuring that every student thrives in a caring and inclusive environment."

Abhishek, Business Head, Classteacher Learning Systems added, "We don't see ourselves as just an ed-tech company. We see ourselves as problem solvers, constantly innovating, evolving, and exploring new ways to respond to the real challenges students face. Our goal is not simply to deliver products, but to create solutions that genuinely improve the learning environment."

Gurbaksh Singh, Chief Creative Officer and Chief Innovation Officer, South Asia, dentsu, said, “Innovation isn’t always about adding more; sometimes it’s about reimagining what already exists. By tapping into the unused potential of existing school CCTV systems and embedding AI into them, we’ve transformed them into an intelligent system that can sense and respond to students’ emotional realities, making verbal bullying visible. The result is a human-centric solution that integrates seamlessly into classrooms, helping schools identify and address bullying early.”

Narayan Devanathan, President & Chief Strategy Officer, South Asia, dentsu, commented, “At Dentsu Lab, we build solutions grounded in real-world constraints but designed for scale. In a system serving over 250 million children, even small improvements in how we understand behaviour can create meaningful, lasting impact.”

Where it stands

Bullying Decoder is currently running as a pilot at D.A.V. Public School, with institutional consent and oversight in place. The school has noted early shifts in classroom interactions during the 2025-2026 academic session, an increasing emphasis, it says, on respect, sensitivity, and mutual understanding. Based on these early results, the school plans to expand the initiative to more branches.

The system, as it stands, is built around four principles the team describes as Privacy, Pattern, Intervention, and Empathy.

The design choice not to store audio, not to name individual students, and not to use the data for anything other than counselling is deliberate, a response to the obvious concern that listening to children, even with good intentions, can quickly become something else.

The camera was always there. It just could not hear. Whether what it hears now will change anything for the child who went quiet in Grade 7 is the question this initiative is built around.

KAKURO & SUDOKU

SOLUTIONS:

KAKURO

	23	7		
10	8	2	23	7
22	9	1	8	4
21	6	4	9	2
		7	6	1

SUDOKU

4	2	3	1	6	5	7	8	9
1	5	8	7	3	9	4	6	2
6	7	9	8	2	4	5	3	1
7	1	4	2	8	3	9	5	6
9	6	5	4	7	1	8	2	3
3	8	2	9	5	6	1	7	4
5	3	1	6	9	8	2	4	7
8	9	7	3	4	2	6	1	5
2	4	6	5	1	7	3	9	8