

TECHNOLOGY & MARKET STRUCTURE

The machines came for the plumbing, *not the judgment*

An audit of where artificial intelligence has actually changed the sports card business, where it has not, and why the fraud data is the only part of this story moving fast enough to reprice assets.

CARDVESTR RESEARCH AUGUST 2026 14 MIN READ

01 / The short answer

Adoption is near universal. Impact is narrow.

Ask ten people in this hobby whether AI is changing the card market and you will get ten confident answers, most of them wrong in the same direction. The optimists describe a future of objective grading and algorithmic price discovery. The pessimists describe a flood of undetectable fakes. Both are arguing about a destination while ignoring the odometer.

Here is what the record actually supports as of August 2026. Artificial intelligence has been absorbed almost completely into the operational layer of the hobby: card identification, cataloging, bulk listing, imaging, marketplace price lookup, and internal fraud triage at the grading companies. Those are real, deployed, at scale, and largely invisible to the person buying a slab.

What AI has not done is touch the two variables that determine investor outcomes. It has not made grading meaningfully more objective from the submitter's point of view. And it has not improved price discovery on the assets that matter, because the constraint there was never computation. It was liquidity.

The one place where AI is demonstrably repricing risk is fraud, and the direction is uncomfortable. Counterfeiters industrialized before the graders finished automating. PSA's own numbers, published for the first time in May 2026, show a fraud problem growing several times faster than the submission volume it hides inside.

EXHIBIT A / AI ADOPTION MAP

CARDVESTR FRAMEWORK · NOT MEASURED DATA



How to read this. The vertical axis is effect on the person who owns the card, not effect on the vendor selling the tool. Note the shape: everything that is fully deployed sits in the middle band, and everything with real repricing power sits either unshipped or in the hands of the adversary. This is a positioning framework, not survey data.

02 / The record

Five years of announcements, one durable change

The hobby's AI story starts in April 2021, when PSA's parent bought Genamint, a machine vision startup, and put its founder in charge of product. The pitch was specific: scan every card, measure it, flag alterations, and fingerprint each copy so the same card could be tracked across resubmissions. PSA said at the time that graders would remain human and the technology would assist them.

Five years later, that is still the architecture. Not because it failed, but because it worked exactly as described and nothing more ambitious replaced it.

APRIL 2021

PSA acquires Genamint. Card fingerprinting, alteration detection, and grader assist enter the workflow. Human graders retained.

2021 TO 2022

Computer-vision-native graders launch. TAG grades on a 1,000-point optical scale with a per-defect digital report. Arena Club launches with scan-based subgrades. Both remain niche by volume.

2023 TO 2024

The distribution layer consolidates instead. Fanatics absorbs PWCC and launches Fanatics Collect. eBay buys Goldin and sells its vault to PSA. AI is a feature inside these platforms, never the thesis.

DECEMBER 2025

Collectors acquires Beckett, putting the two largest grading brands under one roof.

FEBRUARY 2026

PSA raises prices for the second time in roughly six months, citing volume and operating costs. Automation has not lowered the sticker price of a grade.

MARCH 2026

PSA grades more than two million cards in a single month for the first time. Industry volume approaches three million.

MAY 2026

PSA publishes its 2025 Fraud Report, the first time this data has been made public. See Exhibit C.

JUNE 2026

PSA pauses its Value tiers against a reported backlog near ten million cards, pushing the entry price of a grade sharply higher.

JULY 2026

Industry grading volume sets a record near 3.6 million cards. Fanatics Collect demos a consumer value scanner at Fanatics Fest. Hobby Card Index launches with AI scanning across a database of more than seven million cards.

ONGOING

Counterfeit submissions, forged autographs, and AI-fabricated sports content all scale faster than the countermeasures built to catch them.


Automation bought throughput, not objectivity

The grading companies are the most instrumented businesses in the hobby, and their numbers are the cleanest way to test whether AI changed anything. PSA has been imaging every card that enters its facilities, before and after grading, for about four years. It maintains a reference library of more than 270 million card images and runs proprietary models against every submission to flag anomalies. That is a serious industrial vision system by any standard.

The output is throughput. PSA graded 2,175,761 cards in March 2026, the first two-million-card month in hobby history, at an average near 98,900 cards per business day. Its monthly output has roughly doubled in four years. By July, industry-wide volume hit about 3.61 million cards.

EXHIBIT B / GRADING VOLUME, YEAR OVER YEAR

SOURCE: GEMRATE MONTHLY RECAPS, 2026



May dip reflects a pullback from four consecutive record months

What this measures. Growth in cards graded, not growth in grading accuracy. The three months shown are the GemRate recap periods we could verify directly. TCG made up roughly two thirds of PSA's May volume, which matters: the capital and engineering attention follow that mix, not the vintage sports market.

Now look at the same period from the submitter's side. Prices went up twice in about six months. The Value tiers were suspended against a backlog reported near ten million cards. Turnaround remains the hobby's most reliable complaint.

That combination is the finding. If machine vision had genuinely replaced the expensive part of grading, unit economics would improve and the cheapest tier would survive. Instead, the cheapest tier was the first thing to break. The bottleneck was never seeing the card. It was receiving, logging, handling, encapsulating, insuring, and shipping the card, and no vision model touches a single one of those steps.

Industry leadership is candid about this when asked directly. PSA's grading president has described the human read on cardstock and surface, including feel and even smell, as something the models do not replace. CGC's card division points out that algorithms still struggle with subtle surface flaws and foil alignment, which is precisely where 9 becomes 10. Upper Deck's president frames removing subjectivity from grading as the prize, stated in the future tense. Nat Turner, who has been building this capability since 2021, describes the near-term value as helping collectors work out what they just pulled and what is worth submitting.

That is a decision-support tool. It is not a repricing event.

In a field this specialized, a model is only ever as good as what it was fed. Chris Ivy's phrase for the alternative: *“garbage in, garbage out.”*

CHRIS IVY, DIRECTOR OF SPORTS AUCTIONS, HERITAGE AUCTIONS, SPEAKING TO
CLLCT, JANUARY 2026

04 / The pre-grade layer

Why “95% agreement with PSA” is a ceiling, not a score


The loudest AI category in the hobby right now is consumer pre-grading. Point a phone at a card, get a predicted grade and a submit-or-skip recommendation. A dozen or more of these services now compete on subscription pricing, and several advertise something close to 95% correlation with PSA outcomes.

Treat those numbers as marketing until an independent party audits them. Three problems, in order of severity.

The label is not ground truth. These models are trained on PSA grades, so their theoretical ceiling is PSA's own consistency. Vendors in this space cite an industry estimate that PSA graders agree with their own prior grade on resubmission roughly 80 to 85 percent of the time. We have not seen an audited source for that figure, and the vendors quoting it have an obvious interest in it. But note the trap either way: if it is accurate, no model trained on those labels can exceed it, and “95% agreement” would be arithmetically impossible against a noisy target. If it is not accurate, the pre-grade industry is citing a statistic it cannot support.

Correlation is measured on the easy cases. A model that sorts obvious 5s from obvious 10s will post excellent aggregate accuracy while being useless at the only boundary with money on it. On a modern chrome rookie, the entire investment case sits in the gap between a 9 and a 10. Ask a vendor for accuracy conditional on the card being a 9 or a 10, and watch the number move.

Photography is the variable, not condition. Phone-captured images under uncontrolled lighting cannot resolve the microscopic surface and edge detail that decides the top of the scale. The apps know this, which is why they coach you on lighting and angle before scanning.



A model cannot be more consistent than the labels it learned from.
Agreement with PSA is therefore a measure of imitation, not of accuracy.

The practical consequence. Pre-grade tools are legitimately useful for one job: triage. Screening a bulk lot to decide what not to submit has real expected value now that PSA's cheapest tiers are unavailable. Treating a predicted 10 as a reason to pay a 10 price is a category error.

05 / Fraud


The only place AI is moving real money

In May 2026, PSA published a 32-page fraud report covering 2025. It is the most important AI-adjacent document the hobby has produced, and almost none of the commentary about it framed it that way.

PSA estimates it intercepted counterfeit and altered material with a projected market value above 200 million dollars in a single year. The company's grading president has said publicly that this was a deliberately conservative figure and that the true scale could be far larger. Against 19.2 million cards examined, of which about 7.7 million were sports, sports cards still accounted for 43.1 percent of all counterfeit submissions.

+45.3% COUNTERFEIT SUBMISSIONS, YEAR OVER YEAR	+250% COUNTERFEIT RATE AS A SHARE OF TOTAL SUBMISSIONS	70.1% OF COUNTERFEITS WERE MODERN, 1990 TO PRESENT	270M+ CARD IMAGES IN PSA'S REFERENCE LIBRARY
--	--	--	--

Read those first two numbers together. Volume growing 45 percent is a story about a bigger hobby. Counterfeit share of submissions growing 250 percent is a story about a market being actively farmed. Fraud is not riding the wave. It is outrunning it.



Why the Jordan rookie is the real signal. The Wagner rate is a curiosity at thirteen submissions. The 1986 Fleer Jordan was the single most counterfeited card by volume in PSA's 2025 data, at a one-in-four rate. PSA's brand protection lead compares the strategy to counterfeiting twenty dollar bills rather than hundreds: pick the note that gets less scrutiny. Mid-value modern cards with large graded premiums and simple printing are the target profile.

Here is where AI enters the picture on both sides of the table, and where the asymmetry should worry any long-horizon holder.

PSA's stated detection method is a genuine machine learning application: every submission is imaged and compared against a library of authentic, counterfeit, and altered examples, anomalies are flagged for human review, findings are logged and pushed back into training, and known counterfeits are seeded into the workflow as blind tests. The company also reports investing more than 80,000 hours in grader training in 2025 and sharing counterfeit intelligence across grading companies. That is a competent, layered defense.

But look at the single most useful tell PSA describes. Counterfeits typically originate from scans rather than original print files, which leaves print clarity soft under magnification. That tell is an artifact of a specific production limitation. Generative upscaling and print-pattern synthesis attack exactly that limitation. The defense that AI provides here is pattern matching against known fakes. The offense that AI provides is producing fakes that do not match anything known.

OFFENSE / SCALES WITH COMPUTE

Print-file reconstruction from scans
attacks the softness tell directly

Signature and autopen synthesis
forged signatures track the same marquee names

Slab and label replication
real slab, swapped card, valid cert number

Listing and description automation
distribution at zero marginal cost

DEFENSE / SCALES WITH EXAMPLES

Reference library matching
requires a caught example first

Card fingerprinting
strong on resubmission, blind to new fakes

Human tactile and olfactory read
the part PSA says AI cannot replace

Cross-grader intelligence sharing
real, and newly public as of 2026

The defender needs a labeled example of the attack. The attacker needs only a model release.
This is why fraud share grew 250% while submission volume grew 45%.

The unhedgeable part. Fingerprinting protects a card that has already passed through a grader. It does nothing for a card entering the population for the first time, which is exactly where a generative counterfeit enters.

POSITION IMPLICATION

This is the strongest argument we have seen for the Aristocrat approach, and it is not a sentimental one. When authentication risk rises across the population, the premium migrates to assets with long, documented, verifiable histories: cards graded years ago, with imaging on file, sale records at named auction houses, and continuous custody. Fixed supply plus authenticated condition plus documented provenance is not a nostalgia trade. In a market where the counterfeit share of submissions tripled in a year, it is a risk-adjusted one.

06 / Price discovery

The problem AI cannot solve

The most common claim about AI in this hobby is that it will bring institutional price discovery to collectibles. Scan the card, get a value, trade with confidence. eBay's free price guide now identifies a card down to the parallel. Fanatics Collect demonstrated a consumer value scanner at Fanatics Fest in July. Hobby Card Index launched the same month with sold-price aggregation across marketplaces and grading ROI analysis. The tooling is genuinely good and genuinely free.

It does not solve price discovery, because computation was never the constraint.

A card's price is uncertain for reasons no model removes. Comps are thin: a high-grade vintage key may trade a handful of times a year, and the last sale may be a year stale in a market that moved. The asset is not fungible: two copies of the same card in the same grade can differ in eye appeal by a wide margin, and the model sees a label, not the card. And the input data is contaminated. Shill bidding remains a live, widely discussed problem in this hobby, which means a model trained on public sales is trained partly on fabricated demand. Feed it thin, non-fungible, partly fictitious comps and it returns a confident number with no error bar. That is worse than no number, because the false precision travels.

Where these tools do create value is in the boring middle: identifying what a card is, finding what similar copies have done, and estimating whether grading a specific raw card clears its fee. That is real work made faster. Note what it is not. It is not an edge, because everyone has the same tool for free. A capability available to every participant at zero cost is table stakes, and table stakes compress margins rather than create them.

07 / The research surface

Your comps got better, your reading got worse

There is a second-order effect worth naming, because it is degrading the quality of hobby research in real time.

Generative content is flooding the collectibles information layer. In researching this piece we surfaced a large volume of AI-written card content, much of it published by vendors selling the exact tools being reviewed, citing accuracy statistics with no traceable source, and recycling each other's claims until they read as consensus. Some of it is competent. A great deal is confidently wrong about dates, figures, and mechanics.

The sports information layer around the hobby is under the same pressure. A January 2026 study of AI-generated fake content in sports, reported by Reuters, documented fabricated quotes attributed to named athletes and coaches spreading at volume, with researchers warning the material is realistic enough and plentiful enough that ordinary readers cannot filter it. Player narrative drives card demand. Fabricated player narrative drives fabricated card demand.

The practical response is unglamorous. Trace every number you act on to a named primary source: a grading company's own report, a GemRate recap, an auction house's realized results. If a claim cannot be traced past a blog citing another blog, treat it as unverified regardless of how many places you have seen it.

08 / Positioning

Where to use these tools, and where to ignore them

USE CASE	CARDVESTR READ
Cataloging and portfolio tracking	Use it. Scanning a collection into a structured record is the highest-return hour of AI use available to a collector today, and it is largely free. Verify parallels manually on anything you consider a position rather than an item.
Submission triage	Use it, carefully. With PSA's Value tiers paused, deciding what not to submit has real expected value. Use pre-grade tools to eliminate candidates, never to justify paying up.
Predicted grades as a buy signal	Ignore. The 9 to 10 boundary is where the money is and where these models are weakest. Phone photography cannot resolve it.
Algorithmic price targets	Ignore for illiquid assets. On thin comps a model produces false precision. Use it to bound a range, never to set a bid.
Authentication risk	Reprice it upward. This is the material change of the last twelve months. Favor long custody chains, older certifications, named auction provenance, and cards with imaging already on file.
Mid-value modern with large graded premiums	Underwrite the fake risk explicitly. PSA's data shows this is the counterfeiter's preferred profile, not the six-figure grails.
Raw purchases from unfamiliar sellers	Hardest hit category. The economics of counterfeiting improved faster than your ability to inspect. Cert-verified, or pass.

AI has been fully absorbed by this industry and has changed almost nothing about how a card gets valued. It moved cards through buildings faster, made listings cheaper to create, made identification free, and made fraud detection more systematic. Every one of those is an operational improvement that accrues to platforms and graders before it accrues to owners.

The two things that would actually reprice the market, objective grading and honest price discovery, remain unsolved, and the reasons are not computational. Grading is gated by physical logistics and by the fact that its own labels are noisy. Price discovery is gated by liquidity and by contaminated inputs. No model release fixes either.

Meanwhile the adversarial application is compounding. A 250 percent increase in the counterfeit share of submissions in one year, disclosed by the largest grader in the hobby, is the only number in this piece that should change how you allocate. It argues for the same thing the fundamentals already argued for: fixed supply, authenticated condition, documented provenance, and a long paper trail. The machines did not create that thesis. They just made the alternative more expensive.

SOURCES AND METHOD

All figures cited above come from the primary sources listed here. Exhibits A, C and E are CardVestr analytical frameworks and are labeled as such rather than presented as measured data. Individual card prices are deliberately excluded from this piece because third-party price data in this market is unreliable enough that quoting it would undercut the argument.

- PSA 2025 Fraud Report, as reported by Sports Collectors Daily and Sports Collectors Digest, May 2026: [counterfeit rates and detection method](#)
- PSA acquisition of Genamint, April 2021: [company announcement](#)
- GemRate monthly grading recaps, March through July 2026, via [gemrate.com](#) and reporting at Sports Collectors Daily and Sports Illustrated
- Industry executive commentary on AI, cllct panel, January 2026: [how will AI impact the hobby](#)
- Reuters coverage of the Alethea study on AI-generated sports content, January 2026
- Vendor accuracy claims for consumer pre-grading tools are drawn from the vendors' own marketing and are identified as unaudited throughout

CardVestr Research · Market Structure No. 14 · August 2026

This publication is research and commentary, not investment advice. Collectibles are illiquid, unregulated, and can lose value. Verify every figure before acting on it, including ours.