transmit security

Fraud in the Dark

Why Only Predictive AI Can See Through the AI Agent Fog

August 2025

Summary

Two major shifts are fundamentally reshaping the fraud landscape.

First, **advanced generative AI is supercharging fraudsters**. Tools once reserved for nation-state actors are now broadly accessible, enabling anyone to launch convincing phishing attacks, generate deepfakes, create fake websites and emails, automate scams across channels, and harvest personal data at scale. The result: **a surge in successful identity and credential theft.**

Second, the **interface between consumers and digital applications is rapidly evolving**. AI agents, such as ChatGPT Agent and others, are becoming the new front end for how users interact with banks, e-commerce platforms, airlines, and more. But these same trusted agents can also be commandeered by fraudsters using stolen credentials. As that happens, **current fraud detection systems go blind**. Traditional controls like device fingerprinting, behavioral biometrics, bot detection, and anomaly rules can't determine who's behind the agent, **a legitimate user or a fraudster.** This isn't a tuning issue; it's a fundamental technological blind spot.

Because **AI agents operate on the consumer side**, banks and merchants cannot fully control or restrict them. As adoption accelerates, fraud losses are set to spike dramatically. This forces financial institutions and online retailers to **rethink fraud prevention in a world they no longer fully control**.

There is one technology designed for this new era: **Predictive AI**.

Unlike legacy systems, Predictive AI doesn't rely on fixed rules or human-like behavior patterns. It continuously learns from massive volumes of real-world data, uncovering intent, detecting subtle fraud signals, and adapting in real time to new attack techniques, even when traffic flows through AI agents.

Transmit Security is the leader in Predictive AI for identity and fraud prevention. Our platform is built from the ground up around Predictive AI, not retrofitted with it. Trained on billions of real-world interactions across hundreds of digital journeys, our models understand what legitimate and fraudulent behavior really looks like, across any channel, device, or agent. We've assembled a team of world-class experts in AI, cybersecurity, and identity to stay ahead of the evolving threat landscape and deliver results where other technologies fail.

A 20-Year Playbook Just Got Obsolete

For the past two decades, fraud prevention has followed a familiar and largely successful playbook. Solutions were designed to detect human behavior anomalies, identify unusual device patterns, spot bot-like activity, and enforce rules based on velocity, location, and transaction logic. These layers, while imperfect, evolved gradually and kept fraud losses within acceptable limits, especially in regulated industries like banking, fintech, and e-commerce.

> "In this new world, stolen credentials plus access to a trusted AI agent equals undetectable fraud."

But this entire model rests on a critical assumption: **the user is human, and their device and behavior are observable.**

That assumption is now collapsing.

With the rise of consumer-ready AI agents and browsers like ChatGPT Agent, Perplexity's Comet, and others, the interaction layer between users and applications is rapidly shifting. Tasks once performed by a person, like logging in, checking balances, buying products, or changing account settings, are now being delegated to autonomous agents. These agents don't type, swipe, hesitate, or make typos. They execute with precision, speed, and consistency. From a fraud detection standpoint, **they would typically look like bots, but they're being legitimately used by real customers.**

This creates a nightmare scenario for traditional fraud controls. Behavioral biometrics can't make sense of agent-driven interactions. Device fingerprinting is bypassed because the AI agent, not the end-user device, initiates the session. Velocity rules are irrelevant because agents behave faster and more systematically than humans. Even bot detection systems struggle, because now **the bots are trusted.**

And fraudsters know it.

In this new world, **stolen credentials plus access to a trusted AI agent equals undetectable fraud.** Fraudsters don't need to build sophisticated automation. They can just instruct a legitimate agent to act on their behalf using compromised data. From the application's perspective, everything looks clean. But behind the agent, it's a criminal pulling the strings.

This isn't a small shift. It's not a bug in the system. It's a paradigm collapse. The tools and models that have defined fraud prevention for the last 20 years are becoming ineffective, not because they were poorly designed, but because **they were never meant for this reality.**

The Answer Is Predictable

To regain visibility and control, enterprises must abandon assumptions hard-coded into legacy fraud stacks. They must embrace technologies that don't rely on human traits or predefined rules but instead **predict intent based on complex, real-time signals, even in agent-driven environments.**

That technology is **Predictive AI.**

And it's not a nice-to-have anymore. It's the only viable path forward.

Predictive AI is a branch of artificial intelligence focused on using historical and real-time data to **predict the likelihood of future outcomes.** In the context of fraud prevention, Predictive AI examines a user session (or agent session) and asks a powerful question:

> ***"Given everything I know about this session and millions like it, how likely is this to be fraud?"***

Unlike rules or heuristics, which rely on fixed assumptions, Predictive AI uses **machine learning algorithms trained on massive datasets.** These algorithms don't just spot known patterns, they learn evolving behaviors, adapt to changing attacker techniques, and discover risk signals humans may never have considered.

One of the most effective models in this space is **CatBoost**, a gradient-boosting machine learning model especially well-suited for fraud detection. It handles **tabular data with high cardinality** (like user IDs, device types, transaction types), manages **missing data gracefully,** and is highly interpretable, meaning it can help explain why a certain transaction was flagged as high risk.

CatBoost is excellent at identifying interactions between features that may not be obvious: for example, how the combination of a rarely-used device, slightly unusual session timing, and a high-value transaction correlates with fraud, even when none of those signals alone would raise a red flag.

Predictive AI models **constantly retrain and improve,** learning from false positives, newly discovered fraud patterns, and shifts in legitimate behavior. They evolve just like attackers do.

This makes Predictive AI not only smarter, but also **more resilient and future-proof** than traditional systems.

No, Your Current Machine Learning Isn't Predictive AI

Every fraud vendor today claims to "use machine learning." And they're not wrong, technically. But the truth is, **most of what's marketed as AI in fraud prevention is little more than outdated ML tooling dressed up as innovation. Many of the fraud prevention systems used by banks and merchants aren't even that, they're just rule-based engines..**

Machine learning (ML) is a broad category. It includes everything from basic decision trees to advanced neural networks. Most of the fraud solutions that claim to use ML rely on models trained infrequently on old data, using narrow signals to flag known attack patterns or fine-tune rules. These systems aren't dynamic, don't adapt in real time, and **were never built to generalize in an environment of unknown threats**, let alone one dominated by AI agents.

Predictive AI is different. It's not a buzzword. It's a highly specialized application of ML that's purpose-built to **anticipate what's going to happen next,** not just detect anomalies in what already occurred. In fraud prevention, that means scoring every session or user action in real time with a probability of fraud, **before it happens,** using vast amounts of structured and unstructured data, from device signals and session context to behavioral cues and historical outcomes.

A true Predictive AI system must:

- Ingest and analyze real-time event streams, not batch logs.
- Evaluate intent, not just match behavior.
- Continuously learn from labeled outcomes (fraud confirmed, user verified, case closed).
- Drive decisions with high precision and recall—not just "maybe flag this."

To highlight the difference, consider behavioral biometrics, often cited as ML-based fraud prevention. These systems monitor how humans type, swipe, or move a mouse to detect impersonation or bot activity. But **behavioral biometrics isn't Predictive AI**. It:

- Focuses on identity verification, not intent scoring.
- Relies on human inputs, which **don't exist** when AI agents operate apps on the user's behalf.
- Looks for known patterns and deviations, not unknown correlations.
- Fails to adjust in real time based on fraud outcomes.

In a world where **AI agents now serve as the interface between users and applications,** behavioral biometrics becomes obsolete. There are no keystrokes to analyze. No mouse movements. Just clean, fast, machine-driven sessions that look identical, whether triggered by a legitimate user or a fraudster with stolen credentials.

This is why **fraud controls rooted in traditional machine learning are collapsing.** The models weren't designed for agent-driven environments. The fraud signals they rely on no longer exist. And most importantly, they can't adapt quickly enough to evolving attacker methods and tools.

Predictive AI can. It takes a data-first approach instead of fraud MO pattern detection approach. It thrives even in signal-poor environments. It doesn't need human behavior, it needs data. It learns from the outcome of every event and uses those learnings to make better predictions the next time. Whether the session was driven by a person or a bot or an AI agent, it doesn't matter. Predictive AI looks beyond surface patterns and toward intent.

IN SHORT:

Machine learning was the last generation's fraud control.

Predictive AI is the next.

Predictive AI Doesn't Age Out

The difference between traditional fraud detection and Predictive AI isn't just in accuracy, it's in resilience. Fraudsters constantly evolve. They test systems, probe weaknesses, and adapt their tactics. Traditional fraud solutions, built around static rules or anomaly detection, inevitably fall behind.

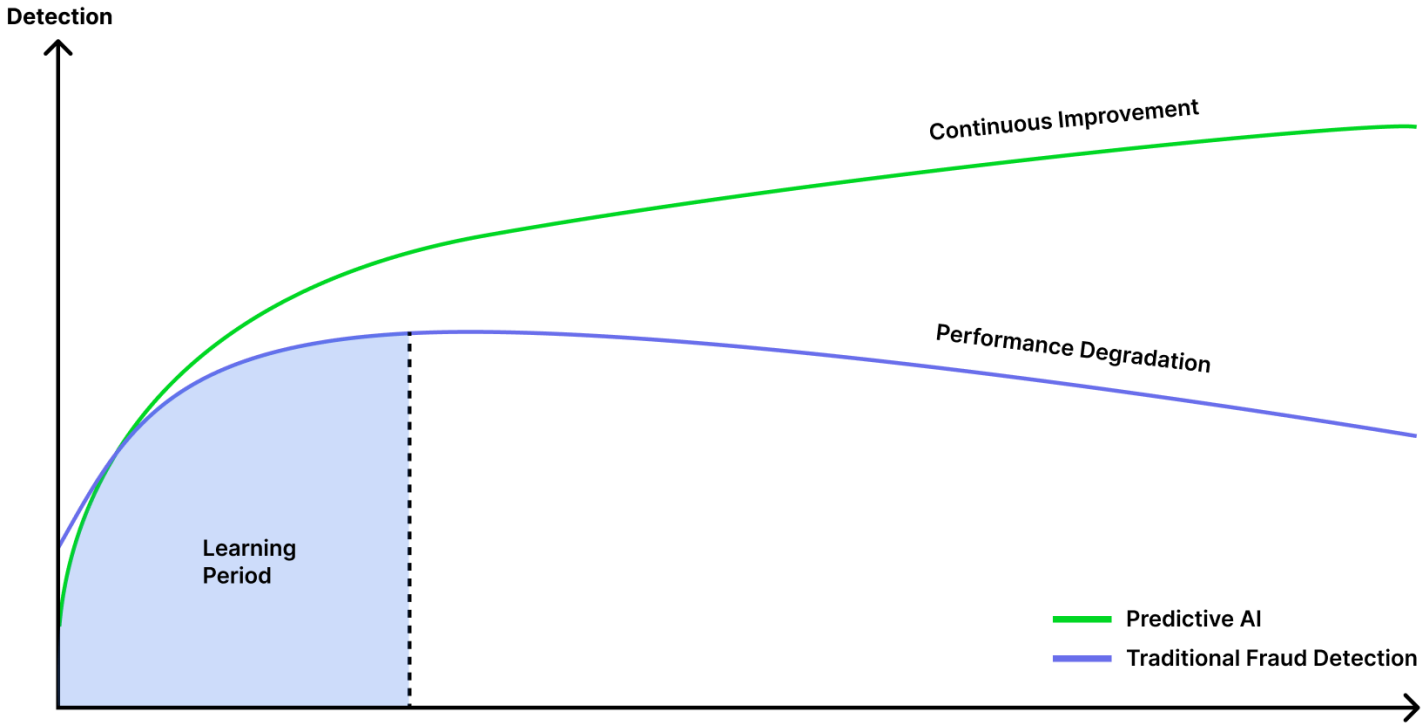

The graph above illustrates this reality. When a traditional system is deployed, it often needs a **learning period** before detection begins. Once tuned, it may achieve impressive results for a time. But as attackers adapt, using malware, scams, or now, AI agents, **detection rates begin to fall.**

Simultaneously, false positives creep upward as systems misclassify legitimate customers who don't fit the outdated pattern.

Traditional fraud detection systems force teams into a reactive cycle: deploy rules, detect a few attacks, then watch as fraudsters adapt and slip through. By the time you've patched the system, the next tactic has already emerged. It's exhausting. It's inefficient. And it's no longer sustainable.

Predictive AI flips that dynamic.

Rather than starting from zero or relying on pre-defined signatures, Predictive AI begins with a **foundation of intent detection,** trained on billions of real-world fraud outcomes. Its performance improves continuously because it doesn't just react to rules, **it learns from every interaction in real time.** As new fraud tactics appear, the model adapts, not weeks later, but **immediately,** fueled by constantly streaming data and outcome feedback loops.

This is visible in the second curve: Predictive AI starts at a strong baseline and steadily climbs, increasing its detection rate over time. Unlike traditional systems, it doesn't degrade. It doesn't need patching. **It evolves because fraud evolves.**

This resilience is critical in a world where fraud is increasingly operated by intelligent agents. When even legitimate users are delegating access to AI, fraud detection needs more than rules. It needs foresight.

Predictive AI is not only more accurate, **it's built to last.**

Predictive AI In Action

While Predictive AI may sound futuristic to some, it's already transforming fraud prevention today. **Transmit Security's Mosaic platform** is deployed across major financial institutions, airlines, telcos, and ecommerce platforms, delivering results that are **impossible to achieve with legacy fraud prevention systems.**

The following real-world examples demonstrate the power and maturity of Transmit Security's Predictive AI:

Global Airline

- **86% detection accuracy** after just one month of deployment, climbing to 96% shortly thereafter.
- Achieved a **100% true positive rate** during a major fraud campaign over the holiday season.
- Detected fraudulent activity invisible to previous solutions, enabling the security team to **eliminate the threat in real time.**

Top Telecommunications Provider

- Reached **82% detection accuracy** in under 30 days before model tuning is even completed.
- Projected **1,000%+ ROI** based on early fraud savings and operational cost reductions.
- Early success created a fast path to **full-funnel expansion across identity and transaction flows.**

Major US Bank

- Achieved a **98% reduction in new account fraud,** largely driven by Predictive AI identifying bot-driven enrollment attacks.
- Realized a **90% decrease in false positives and false negatives,** drastically improving both fraud prevention and customer experience.
- Detected **500% more bot attacks** than the legacy system, dramatically increasing visibility into attack surfaces.
- Reduced operational overhead by **80%,** freeing up fraud analysts and reducing alert fatigue.
- Delivered a **1,300% return on investment (ROI)** - one of the most dramatic improvements ever recorded for a fraud prevention tool at the bank.

Leading Financial Services Provider

- Increased fraud detection rate by **610%,** identifying sophisticated patterns missed by previous-generation tools.
- Cut false positives and false negatives by **98%**, allowing for smarter decisions without introducing customer friction.

All of these institutions had **previously deployed mainstream, rule-based or machine learning fraud tools,** yet Predictive AI outperformed them by a factor of **4x to 20x.**

Whether it's reducing fraud, cutting operational costs, or restoring trust in automated defenses, **Predictive AI isn't just a technology shift, it's a performance revolution.**

Why Transmit Security Leads This Space

Transmit Security isn't new to Predictive AI. We've spent **years building and deploying advanced fraud detection models,** focused specifically on the world's most complex environments: banking, trading, retail, and travel.

Our Predictive AI systems are trained on **billions of events,** across hundreds of customer journeys, with real-world fraud outcomes as ground truth.

We didn't retrofit Predictive AI into our stack, we built our platform around it. And we've invested in **the best minds in AI, cybersecurity, and identity,** ensuring our models stay ahead of both fraudsters and the evolving traffic landscape. We're not building a tool, we're building **the new standard** in fraud defense for the agent-powered internet.

Summary: A New Era, A New Defense

Fraud prevention is at a historic inflection point. AI is no longer just a tool used by defenders, it's now in the hands of attackers and end users alike. Generative AI is enabling phishing campaigns, impersonations, and credential theft at scale. Meanwhile, good AI agents, like ChatGPT Agents, are beginning to operate applications on behalf of legitimate users, changing the very nature of digital traffic.

This dual shift breaks the foundation of legacy fraud prevention:

- **Device identification fails** because devices are new and unlinked.
- **Behavioral biometrics fails** because there's no human behavior to analyze.
- **Bot detection fails** because these are not bots, they are intelligent, authenticated agents.

In this new world, **most fraud becomes invisible,** and organizations face a choice: either challenge customers aggressively and add friction or accept significantly higher fraud losses.

Predictive AI is the only viable path forward. It's not just a better tool, it's a fundamentally different approach, built to see what others miss. It works in the dark. It doesn't need to know who you are, it understands what's likely to happen next. It shifts fraud prevention from reactive to proactive.

> The question is no longer ***if*** Predictive AI is needed, but ***how soon*** organizations are willing to adopt it.

Transmit Security has already proven this at scale: with global banks, telcos, airlines, and financial services firms realizing up to **1,300% ROI, 90–98% reductions in false alerts, and fraud visibility 4x–20x higher** than previous solutions.

The question is no longer *if* Predictive AI is needed, but how soon organizations are willing to adopt it. Because in the era of intelligent agents, fraud doesn't look like fraud anymore. Only Predictive AI can tell the difference.

About Transmit Security

Transmit Security delivers future-proof customer identity experiences in a world where AI is accelerating change across both fraud and user access. We do this by fusing customer identity, fraud prevention, and identity verification into a single, unified system, eliminating silos and enabling rapid adaptation.

At the core is our Predictive AI, continuously learning from real-time signals to detect intent, uncover emerging fraud patterns, and make accurate decisions before damage is done. This fusion-first approach allows leading enterprises to stay ahead of evolving threats while delivering seamless, secure experiences to their customers.

Request a demo

CONTACT US