twodots INSIGHTS

Research Report • 2026

FRAUD & BAD DEBT IN RENTAL HOUSING

Why Anti-Fraud Tools Haven't Solved the Problem—and What Actually Will

A data-driven research report for rental housing operators, backed by Two Dots backtest data from NMHC top 50 partners

A Research Report by Two Dots

May 2026

Table Of Contents

Executive Summary 3
Key Findings 4
Section 1: The Macro Backdrop 5
Section 2: Fraud Is Not a Criminal Problem 6
Section 3: Fraud Industrialized 7
Section 4: The Current State of Anti-Fraud “Solutions” 8
Section 5: CPNs: The Fastest-Growing Fraud Type 10
Section 6: Overrides Are the Leading Cause of Bad Debt 12
Section 7: The Ticking Time Bomb Hidden in Your Portfolio 14
Section 8: What a Real Solution Looks Like 15

Appendices
Appendix A: Fraud Types 17
Appendix B: Centralized Screening Teams and the Override Problem 19
Appendix C: Backtest Methodology 20

About Two Dots

Two Dots builds AI underwriting technology that automates and solves problems traditional systems can't. Our conversational agent, Evie, uses natural language to guide applicants through approvals, like handling sophisticated fraud detection and complex income verification. Evie instantly resolves unusual situations and missing information that typically require manual intervention and create delays, while making the experience effortless for applicants and delivering real-time approvals that boost your bottom line.

Learn More at twodots.com

EXECUTIVE SUMMARY

The Hidden Drivers of Rental Housing Bad Debt

- Why has bad debt remained elevated compared to long-run averages? The typical answer is fraud got easier and cheaper – that's true, but it's only a small part of the story.

- Most fraudsters actually intend to pay and simply don't meet the traditional 3x ITR and credit score cutoff criteria. They are driven into fraud by an affordability crisis. Rent has gone from 20% to 50% of a low-income household's spending.

- Two Dots backtest data shows that 80% of applicants who commit fraud still pay rent (for now). When macro conditions worsen, that fraud surfaces in your bad debt number.

- Adoption of new 'anti-fraud' products coincided with improving macroeconomic conditions – as a result, naturally occurring bad debt reduction has masked poor product performance.

- An override culture was created by automating the traditional credit score and income-to-rent cutoffs system – while every year a smaller proportion of the renter population qualifies that way.

- Two Dots backtest data from NMHC top 50 partners shows that the #1 cause of bad debt today is leasing criteria non-compliance.

AT A GLANCE

Key Findings

80%

of applicants who commit fraud still pay rent—they're driven by the affordability crisis, not criminal intent.

Two Dots backtest audit data, 18-month resident portfolio analysis, NMHC top 50 partners

260%+

increase in CPN (synthetic identity) fraud in four years, from 1.5% of applications in 2022 to over 4% in 2026.

Two Dots internal application data, 2022–2026

30-40%

of nonpayment comes from leasing criteria non-compliance—overrides, not missed fraud, are the #1 cause of bad debt.

Two Dots backtest data, NMHC top 50 partners

25-35%

of nonpayment comes from missed fraud—the second-largest driver, but not the one operators typically focus on.

Two Dots backtest data, NMHC top 50 partners

>10%

of offer letters processed by Two Dots are fraudulent—the most-recommended workaround for beating standard screening.

Two Dots internal application processing data

85 bps

current average bad debt across rental housing—historically high despite an otherwise strong economy.

Industry average, varies by asset class

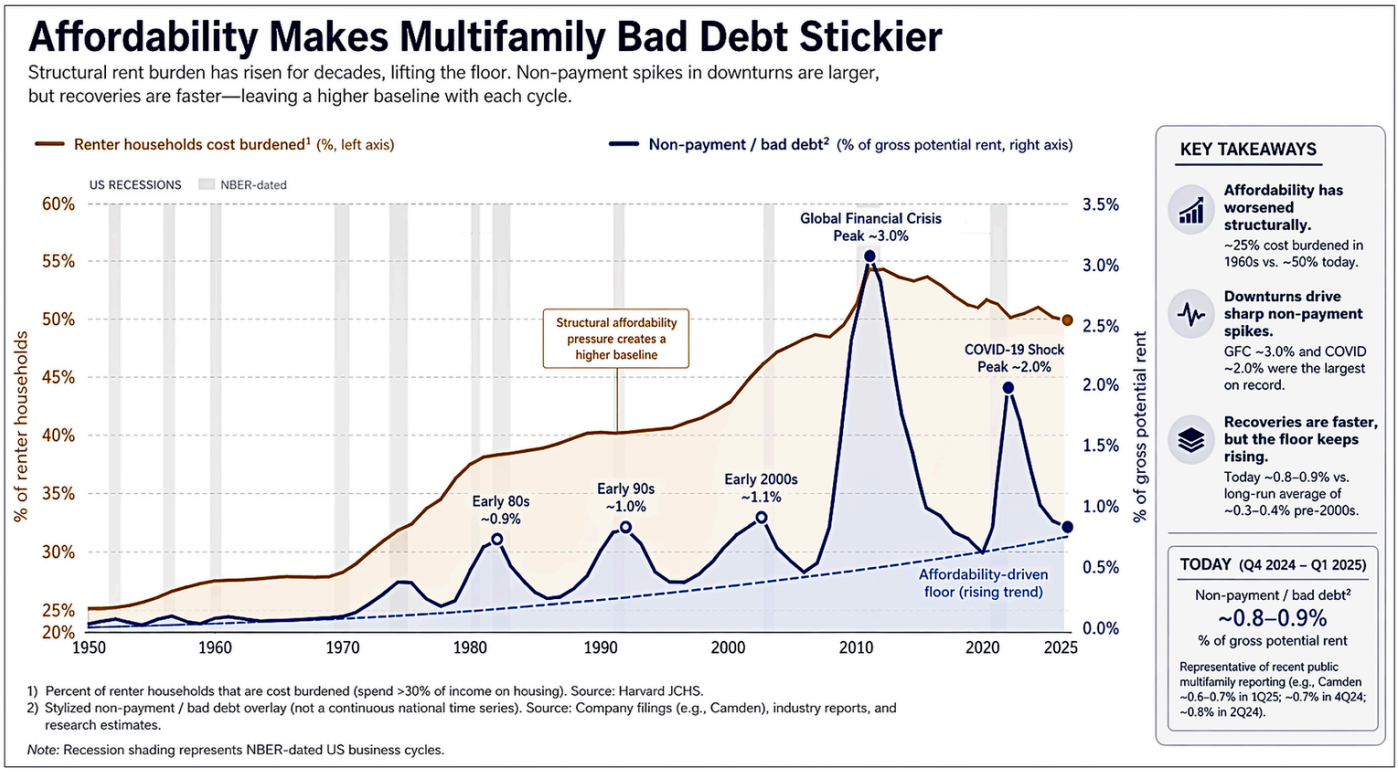


Two Dots Insight: The most useful question is not "Do we have fraud?" The useful question is: **How much of the current resident base got into the portfolio through a path that would fail a serious re-underwrite today?**

SECTION 1

The Macro Backdrop

Non-payment is cyclical – it rises and falls with the macroeconomic environment – amid a long-run rise driven by the affordability crisis. In stress periods like 2008 or 2020, non-payment spikes to 2-3%+ of gross potential rent. Today non-payment is low – inflation has come down and we're at full employment – but bad debt is still elevated vs how low we've gotten in past economic boom times (see chart).



The story of why bad debt has remained stubbornly elevated, largely propagated by screening and 'anti-fraud' vendors, is that fraud became so easy and cheap Americans en masse just decided to start ripping off their landlords.

While it's true that fraud became easier and cheaper – and the reward for committing it in places like Atlanta with 6-9 month eviction timelines encouraged more opportunists – our data proves that true criminals aren't the real bad debt driver.

SECTION 2

Fraud Is Not a Criminal Problem

When you think of who commits fraud, the standard image is a criminal – someone in a dim room, dark clothing, hunched over their computer, unwilling to pay rent.

The thing is, that group of people – looking for financial gain – makes up a relatively small portion of application fraud.

> **Today, the vast majority of multifamily fraudsters intend to pay. They commit fraud not because they want to default, but because they don't meet traditional 3x ITR and credit cutoffs – and they need an apartment**

Median family renters spent 15-19% of their income on rent in the 1960s. Today that number sits between 35-50%, depending on income quartile and region. As rent burden has climbed, the share of applicants who can't meet standard cutoffs honestly has climbed with it.

Many applicants who can't meet income or credit requirements feel they have no option except fraud to find a home.

Two Dots backtest audit data analyzing existing resident portfolios shows that over the last 18 months, only 20% of fraudsters default within the first 60 days. The other 80% pay rent.

We can use this data as a proxy for what proportion of fraudsters are opportunist criminals (~20%) vs those that are driven into it by the affordability crisis (~80%).

80%

of fraudsters pay rent—driven into fraud by affordability, not criminal intent

Two Dots backtest audit data

20%

of fraudsters default within 60 days—the true opportunist criminals

Two Dots backtest audit data

35-50%

of income spent on rent by modern renters, up from 15–19% in the 1960s

U.S. Census Bureau; Joint Center for Housing Studies of Harvard University

SECTION 3

Fraud Industrialized

On the applicant side, affordability drove a massive increase in 'demand' for fraud. While it became easier for people to commit basic types of fraud at home with their computers and the internet, a huge black market industry was created to meet this market need.

Simply google 'apartment approval package' and you will see hundreds of small and medium-sized businesses selling:

- Fake offer letter services
- Synthetic identities with fake SSNs that get past traditional screening
- Fake social security cards to back it up
- Full service 'approval concierge' businesses that apply to many units on your behalf
- Many more purveyors of exotic fraud types

The vast majority of applicant fraud comes from this black market industrial complex, which makes it seem as if you're being targeted by sophisticated criminals when the reality below the surface is often an applicant who wants to pay and can't really afford it.

The multifamily vendor market responded as well to what many operators treated as a five-alarm fire in the Covid era, and suddenly a whole new category of 'anti-fraud' products was born.

SECTION 4

The Current State of Anti-Fraud "Solutions"

This section goes over the major types of products and features that have been marketed as anti-fraud and what types of fraud they actually prevent, if any.

Two Dots Insight: Our backtest data from across the NMHC top 50 shows that many of these products' poor performance at preventing fraud has been masked by the improving macroeconomic backdrop and disappearing local fraud incentives like eviction moratoriums.

See Appendix A for more on specific types of fraud – there are many.

1 Dedicated Document Anti-fraud

During Covid's peak, companies started using a combination of machine learning and offshore review teams to detect fake PDFs. The PDF file type was created in the 1990s to mimic legal documents – so they carry unique metadata which makes it possible to link a document to its original source.

The problem with these companies is fivefold:

1. Needs original PDFs – photos and screenshots destroy metadata, creating friction for applicants and operators.
2. Faking documents is just one type of fraud – e.g. fake SSNs pass through.
3. False positives are common – privacy-related annotations like blacked-out account numbers, or combined documents trigger the fraud detection system.
4. Many documents have no verifiable metadata – offer letters, real or fake, are all made in MS Word or Google Docs.
5. Models have to be trained for each institution (Bank of America, ADP, etc.) – people with smaller banks, or working for small businesses with manual payroll often get unfairly discriminated against.

These tools were marketed as holistic solutions for the multifamily bad debt problem. In reality, they prevented one narrow type of fraud while creating enormous friction for qualified applicants that hurt occupancy.

2 Bank Linking

Bank linking emerged in the mid-to-late 2010s. The pitch: bank linking prevents document fraud by replacing manual paystub uploads with direct deposit data.

The issue is bank linking only addresses one fraud type and is easy to bypass. Fraudsters or legitimate applicants can claim their bank isn't supported. Many qualified applicants refuse to link their account based on privacy concerns. Over 40% of modern renters have complex income types that can be difficult or impossible to verify with bank data such as: offer letters, gig work on prepaid cards, cash income, self-employment, physical check cashing, etc.

This complex income group is also getting larger each year – the share of Americans holding a single standard W-2 job continues to decline.

Bank linking is a genuine workflow efficiency improvement, creating a faster process for your qualified renters that are willing to do it. But it pushes all the fraudsters into the manual review category alongside the rising share of renters with complex income.

3 ID Verification

Companies like Checkpoint ID, which verify identity with pictures of photo IDs, gained widespread adoption during Covid. Some innovative ML models can spot obviously fake IDs, but cell phone images of photo IDs don't carry metadata the way PDFs do – making it extremely difficult to detect fakes consistently. This generates a significant amount of false positives, while good fake IDs still pass the check.

Photo ID verification primarily helps prevent true identity theft, where someone pretends to be an entirely different person. That's a rare fraud type in multifamily since you have to take physical possession of the unit and therefore show your real face at some point. Using a fake SSN[1] with a legitimate ID is far easier than stealing someone's identity.

4 What People Are Using Today

Bank linking and scanning photo IDs were technologies that had use cases across many industries, so Silicon Valley VCs funded tons of companies that sold cheap bank link & ID APIs that any company could easily resell. A wave of 'anti-fraud' startups was born that were simply wrappers around companies like Plaid and Jumio with a PMS integration slapped on.

The property management software companies eventually integrated the same APIs into their own screening workflows as an 'anti-fraud' upsell, which is what the vast majority of operators are using today. The next section discusses the newest fraud type to industrialize – one none of these tools prevent.

[1] Also referred to as a Credit Privacy Number, or "CPN", in following sections.

SECTION 5

CPNs: The Fastest-Growing Fraud Type

260%+ increase in CPN fraud in four years—from roughly 1.5% of rental applications in 2022 to over 4% in 2026

Two Dots internal data shows Credit Privacy Number fraud grew from roughly 1.5% of rental applications in 2022 to over 4% in 2026 – a 260%+ increase in four years. CPNs are now the most common fraud type to get through other systems according to our backtest data, with a higher correlation to default than most other fraud categories.

To understand how the credit system we're all supposed to trust got corrupted, you need to understand how credit files actually get created.

How it works

Credit bureaus do not receive a masterfile of SSNs from the government. They infer identity from your participation in the financial system. You don't get a credit profile along with your birth certificate. Instead, when you open your first bank account at 14 or 16, the bank sends your information to TransUnion or Experian. If it doesn't match an existing file, they create a new one. Real people then accumulate regularly spaced credit checks over time – a pattern fraudsters can imitate.

This used to be rare and hard, but in 2011 the Social Security Administration randomized the number order of all new SSNs because structured numbers[2] were running out. The unintended consequence: any nine-digit sequence can now look like a legitimate SSN.

Example: Let's use a fake person, Jim Johnson, to understand better how this all works.

Say Jim wants to rent an apartment but has bad credit. He applies for a credit-builder card – a product designed for people with thin or no credit – using his real name and a random nine-digit SSN. He gets approved, makes small transactions monthly, and may use services like Boom to backfill fake tradelines. In a few weeks, he has a thin but clean credit file that passes most screening systems. His name is real, his face is real, his documents aren't edited – Jim looks like a low-risk applicant, but the credit identity underneath is fabricated. Even worse, services that do the above for you are widely available for a few hundred bucks.

[2] SSNs were traditionally structured – the first digits corresponded to birth year and geography.

Why the Bureaus Don’t Fix It

The bureaus don't fix this because they face FCRA violations if they merge or delete the wrong files. If TransUnion merges Jim's synthetic profile with a real Jim Johnson's file, and it turns out there really are two Jim Johnsons, the real one now has fraudulent tradelines and gets wrongfully denied. That opens everyone to lawsuits, so the bureaus let duplicate profiles coexist.

Traditional screening systems and some anti-fraud systems can flag some CPNs for your teams to review by detecting patterns like post-2011 issuance and limited tradelines despite older age, but picking between the fraudsters and the false-positive flags is extremely difficult. It also contributes to the current leading cause of bad debt examined in the next section.

See Appendix A for more information on fraud types.

SECTION 6

Overrides Are the Leading Cause of Bad Debt

The single largest driver of bad debt today isn't fraud that slips by completely undetected. It's residents who moved into your buildings despite not meeting the property's leasing criteria – admitted through overrides or coaching that worked around the rules.

30-40%

of nonpayment comes from **non-compliance** (overrides and coaching around criteria)

Two Dots backtest data

VS.

25-35%

of nonpayment comes from **missed fraud**

Two Dots backtest data

Teams aren't acting with bad intentions – there are more qualified applicants who legitimately need overrides than ever before. This is a natural consequence of the affordability crisis combined with 'anti-fraud' features that automate old screening rules – which apply to a much smaller proportion of the modern renters.

How overrides happen

Every anti-fraud tool described in this report produces a flag when it can't make a clean call:

> *This SSN might be fake.*

> *This paystub looks edited.*

> *This applicant has a thin credit history.*

> *This applicant is claiming income that isn't in their bank link.*

Your team has to decide whether it's a false positive, a real fraud, or a complex-but-legitimate situation.

The path of least resistance becomes: approve. Don't investigate the warning too much. Trust their story. Coach the applicant to submit an offer letter. Note that they don't have an SSN and run different checks. Add a co-applicant with stronger credit who never moves in. Push it through as a regional exception.

The real underwriting happens in phone calls, escalations, and judgment calls under time and occupancy pressure.

The implication: most of your bad debt isn't from fraud you missed. It's from applicants that were approved under pressure, with no good way to tell the legitimate, complex case from the risky one.

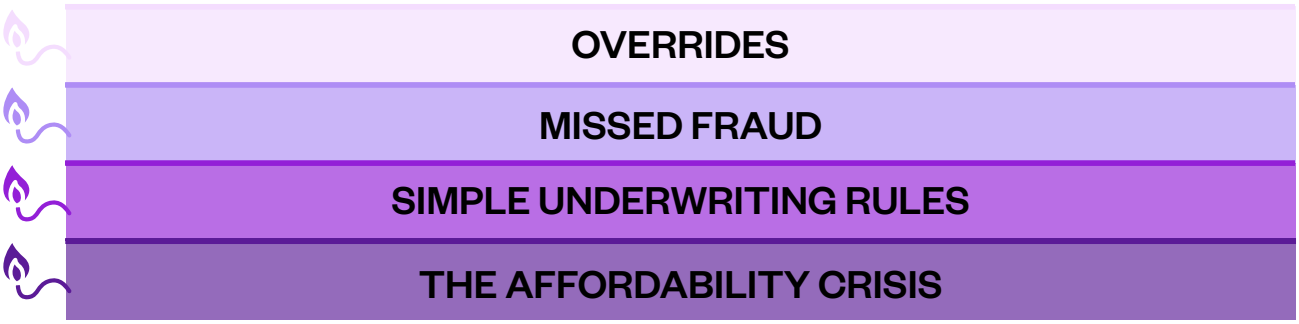
There's a way to do sophisticated underwriting on these rent-burdened and complex applicants. Big banks and subprime lenders do it all the time. That said, very few teams in multifamily are equipped with the same tools, training, and models that a large bank's underwriting department has honed for years. The next section examines what this means for the residents currently in your buildings.

SECTION 7

The Ticking Time Bomb Hidden in Your Portfolio

Most operators look at their current bad debt number and assume the tools they've layered into screening are working. Bad debt is roughly 85bps (varying by asset class) – historically high but slowly trending lower, occupancy is healthy, the dashboards look fine.

What's actually going on is a huge buildup of risk driven by:



Modern anti-fraud tools came of age as the economy was recovering from Covid. From 2021 forward: inflation came down, employment rose, wages climbed. Bad debt fell with the economic tide.

> **Adoption of bank linking, document anti-fraud, and ID verification coincided with that naturally decreasing bad debt – it didn't cause it.**

What happens next?

Rent-burdened fraudsters who are paying today won't have the margin to pay during the next downturn, and many of the override-admitted residents who barely qualified will be in the same position.

The answer isn't more point solutions that create more risk flags or reasons to deny – the structural problem of fraud and overrides is actually driven by there not being enough approvals. The solution is an end-to-end underwriting system.

SECTION 8

What a Real Solution Looks Like

Meet Evie

Evie is a conversational AI underwriting agent trusted by industry leaders – like BH Management, Weidner Apartment Homes, and RPM Living.

Evie uses holistic underwriting to approve more applicants with less risk and optimize property performance:

Prevents overrides and all fraud types
Evie prevents all fraud types and reaches decisions in 100% of applications. If someone submits an offer letter, Evie does online research and contacts employers to authenticate in real-time. No flags and no manual review where overrides flow through.

Approves modern renters, with less risk
Evie underwrites modern applicants with a special scoring model, allowing operators to optimize performance with property-level risk thresholds.

Evie is proactive
Evie reminds applicants to complete the application, guides them through each step, answers questions, gathers additional documentation for complex cases – like thin or no-credit, and asks about guarantors or additional income sources before denial. Where others return false-positives or reasons to deny, Evie finds legitimate reasons to approve.

The first step is a backtest.

We run free backtests on actual portfolio data – every application from the past 12-24 months, screened against Evie, compared against what your existing process produced. The output shows where missed fraud is sitting in your current resident base, where overrides are concentrated, and what holistic underwriting would change.

The diagnosis in this report is general. The backtest is your portfolio's specific version of it.
See Appendix C for backtest methodology.

Book a Demo Today →

See the impact on your specific portfolio.

Two Dots offers a current resident audit that gives a detailed report on how much and what types of fraud are bypassing the current process. No fee or commitment required – just a chance to prove our effectiveness.

Book a Demo Today →

APPENDIX A

Fraud Types

Income document fraud. Altered or fabricated paystubs, bank statements, tax documents, benefits letters. Was the most common fraud type in multifamily until recently being overtaken by synthetic identity fraud. Many applicants who commit this type of fraud do not default – they have passable credit, and are simply rent-burdened.
The correlation between this fraud type and 60-day default is roughly 30-40% during benign economic conditions; that correlation would likely rise during stress periods.

Synthetic identity fraud (CPNs). When an applicant uses a fabricated Social Security Number – a "Credit Privacy Number" or CPN – paired with their real name and identity. The CPN is associated with a clean (often thin) credit profile built specifically for the purpose. Now the most common fraud type in multifamily and the most highly correlated with default. CPNs grew from roughly 1.5% of rental applications in 2022 to over 4% in 2026 – a 260%+ increase. See §5 for the full mechanism.

Traditional identity theft. When an applicant pretends to be a different person entirely, using stolen SSN, photo ID, and other identity documents. Relatively rare in multifamily – under 0.5% of applications – because the fraudster has to take physical possession of the keys and unit, which creates exposure that other fraud types avoid.

Offer letter fraud. Fake offer letters submitted in lieu of paystubs to satisfy income requirements. Often paired with fake employer phone numbers (answered by accomplices) and fake company websites. Over 10% of offer letters Two Dots processes are fraudulent. The most-recommended workaround in online forums for beating standard screening, because most systems cannot meaningfully verify offer letters.

Co-applicant / guarantor fraud. Adding a fake co-applicant – typically with strong income and credit – to an application that would otherwise be denied. The co-applicant has no intention of actually living in or paying for the unit. Often appears suddenly when the primary applicant is about to be denied, after no co-applicant was mentioned in the original application.

Friendly fraud. When a friend or family member voluntarily lends their identity (SSN, photo ID, bank login) to someone with bad credit or a criminal history so that person can pass screening. Sometimes

the lender applies and picks up the keys, after which someone else moves in. No fake documents and no synthetic IDs – just informal trust networks routing around screening criteria.

Inception fraud. When an applicant starts a fake small business and signs up for a real payroll provider like Gusto or QuickBooks. They then transfer money into a business bank account, pay themselves a direct deposit into their personal account, and immediately transfer the money back to the business bank account. This fraud type is relatively rare and is highly correlated with other fraud types like CPNs.

Obfuscation fraud. Deliberately withholding or misstating identity information (typically SSN) to prevent the screening system from accessing real credit, criminal, or rental history. Difficult to handle because legitimate applicants without SSNs (some recent immigrants) and fraudulent applicants hiding bad records produce identical workflow outputs: missing identifiers and incomplete data. New laws in Colorado and California prohibit landlords from asking about immigration status, which limits operators' ability to distinguish the two populations.

APPENDIX B

Centralized Screening Teams and the Override Problem

Some operators have moved screening from site-level decisions to centralized teams or specialists – a model that, on its face, addresses the override problem described in §6. By removing leasing agents from the decision, you remove occupancy pressure from the people making the calls. The reasoning is sound, and centralized teams genuinely improve consistency and training quality.

But centralization shifts the tradeoffs rather than eliminating them.

Speed loss creates its own pressure. Centralized review adds handoffs and delays. Qualified applicants – particularly the rent-burdened, complex-income population the report has discussed throughout – typically have multiple applications out. They accept whichever offer comes first. Your centralized team's slower decision means you lose those applicants to faster competitors. Vacant units pressure operations differently than they pressure leasing agents, but the pressure doesn't disappear.

The detection-vs-adjudication problem persists. Centralized teams are still receiving flags from the same anti-fraud tools described in §4. They still have to determine whether a flag represents a false positive, a legitimate complex applicant, or a fraudster.

Better-trained reviewers make better calls than overworked leasing agents, but they make the same kind of calls – judgment calls without the signal needed to distinguish complex-legit from fraud-mimicking-complex, and detection failures pass through regardless.
Liability stays with the operator. Whether decisions are made by site teams, centralized specialists, or a combination, the operator owns the outcomes. Centralization improves consistency and documentation but doesn't transfer risk.

The result: centralized teams reduce override volume but don't eliminate it, and they don't address detection failures at all. The architecture still produces hidden risk – sometimes more efficiently, sometimes less – but the underlying problem is unchanged.

APPENDIX C

Backtest Methodology

The backtest referenced in §8 evaluates Evie against an operator's actual portfolio data. The output identifies where missed fraud is sitting in the current resident base, where overrides are concentrated by property, correlations to non-payment, and what holistic underwriting would have changed.

Input: The operator provides application data for approved residents at a set of properties.

Process: Each application is re-screened by Evie using the same data the original screening process had access to. Evie renders a decision for every application. Outputs are compared against the operator's actual decisions and the resulting resident outcomes.

Outputs: The backtest report identifies:

- Missed fraud. Applications Evie would have flagged or denied that the existing process approved, segmented by fraud type and current payment status.
- Override concentration. Patterns in override-driven approvals – by property, by leasing agent, by applicant characteristic – and the bad debt outcomes associated with those patterns.
- Bad debt attribution. Current bad debt by approval method: fraud detection failure, override approval.

Privacy and data handling: As a regulated Credit Reporting Agency, Two Dots ensures all handling of application data is FCRA and FHA compliant, and is Soc 2 Type II certified. Data is only used with explicit authorization, and all backtest outputs are returned as a confidential report.

Cost: The backtest is provided at no cost to operators. Two Dots' assumption is that the diagnostic work creates the foundation for a longer engagement; the diagnostic itself is offered as free analysis.

See the impact on your specific portfolio.

Book a Demo Today →

twodots

Henson Orser
Co-Founder & CEO
henson@twodots.com
(415) 699-6416

www.twodots.com

