Transaction and synthetic identity fraud detection for smaller institutions

OPEN METHODOLOGY · FREE TO RUN, VERIFY, AND ADOPT

THREAT MODEL NOTE · VERSION [VERSION]

Every field is real. That is the *attack*

Identity verification asks whether each element of an identity is genuine. A synthetic identity answers yes to every element, because every element is genuine. The falsity lives in the binding between them, and almost nothing in a standard origination stack examines the binding. This note sets out why, and what follows for detection.

THE VERIFICATION RESULT FOR A FABRICATED PERSON

EVERY ELEMENT CHECKED. EVERY ELEMENT REAL.

ELEMENT PRESENTED	GENUINE?	WHAT THE CHECK ESTABLISHES	WHY IT PASSES	RESULT
Social Security number	Yes, issued	That the number was issued by the Social Security Administration and is structurally valid.	It was issued. It belongs to a real person who is not using it.	PASS
Name	Fabricated	That a name is well formed and free of the patterns associated with obvious fabrication.	A plausible name is not a detectable anomaly. Nothing about it is wrong.	PASS
Date of birth	Fabricated	That the applicant is of age and the date is internally consistent with the rest of the file.	Chosen to be consistent. There is no external record contradicting it.	PASS
Address	Yes, deliverable	That the address exists, is deliverable, and is not a known drop.	It is a real address the operator can receive post at.	PASS
Telephone and email	Yes, live	That the contact routes to somebody who can respond to a one time code.	They route to the operator, who responds immediately.	PASS
Credit file	Yes, real	That a file exists with tradelines, payment history, and no derogatory records.	The file is genuine. It was cultivated deliberately over months or years.	PASS

ELEMENT PRESENTED	GENUINE?	WHAT THE CHECK ESTABLISHES	WHY IT PASSES	RESULT
The binding between them	False	That this number, this name, and this date of birth belong to one actual person.	**Not tested by any of the checks above.**	NOT CHECKED

The verification stack is not failing. It is succeeding at the wrong question. Every row above is a correct answer. Six correct answers about six real elements do not combine into an answer about whether a person exists, and no amount of improvement to the individual checks changes that, because the defect is in what is being asked rather than in how well it is being asked.

FOR
Fraud leads, compliance officers, boards
No technical background assumed

ANSWERS
Why the controls you have do not catch this
And what would

READING TIME
About 25 minutes

SCOPE
Origination and early account life
Not card present or transaction fraud

COMPANIONS
Feature specification, implementation guide
[LINK]

VERSION
[VERSION]
Published [DATE]

LICENSE
CC BY 4.0
Copy it into your own training

ARCHIVED AT
[DOI]
Permanent identifier

1 The claim, stated once

Everything after this section is support for a single proposition, which is worth reading slowly because most of the confusion in this area comes from missing it.

> THE PROPOSITION
>
> **Identity verification is a per element test, and a synthetic identity is an error of composition.** Because the falsity is in the relationship between elements rather than in any element, a control that examines elements one at a time is structurally incapable of detecting it. This is not a gap to be closed by better data or stricter thresholds. It is a mismatch between the question the control asks and the thing being done.

Three consequences follow, and they shape every other document in this series.

First, detection has to move to a different class of evidence. If element validity cannot separate a constructed identity from a real one, then whatever does the separating must look at coherence between elements, at behaviour over time, or at structure across applicants. Those are the three families the feature specification is organised around, and the organisation is not arbitrary.

Second, the loss is usually not recorded as fraud. There is no victim to complain, because the person whose number was used is often unaware and, in the case of a child, will not discover it for years. The account simply stops paying, and it is booked as a credit loss. An institution asking how much synthetic identity fraud it has may be looking at the wrong ledger.

Third, good behaviour is part of the attack. Every other fraud type an institution defends against involves anomalous conduct. Here the conduct during cultivation is exemplary, deliberately, for months or years. A control tuned to find bad behaviour will not find it, because during the period when intervention is still possible there is no bad behaviour to find.

---

2 Why the identity number is the load bearing component

The other elements are chosen. The number has to be taken, and which numbers get taken follows directly from how the scheme works.

A fabricated name and a plausible date of birth cost nothing and can be generated at will. The identity number cannot, because it must survive a check against an issuing authority. That single constraint determines the target population, and it explains a pattern that otherwise looks arbitrary.

The most valuable number is one nobody is watching. An actively used number generates records that would contradict the fabricated person: employment, tax filings, an existing credit file under a different name. A dormant number generates nothing to contradict. So the population selected against is not the wealthy or the careless. It is the quiet.

WHO IS TARGETED, AND WHY IT IS THEM

POPULATION	WHY THE NUMBER IS ATTRACTIVE	WHY NOBODY NOTICES
Children	A number issued at birth with no credit activity attached to it and none expected for fifteen years or more.	Nobody checks a child's credit file. The first check is often a first student loan or a first card application, by which point the scheme has run for a decade.
Older adults with no recent borrowing	A long established number with no current credit seeking activity attached.	A person who has not applied for credit in years has no reason to look at a report, and the material aimed at them tells them to watch for accounts in *their* name, which is not what appears.
Recently deceased	The number remains valid to structural checks after death, and the person will never contest anything.	Death reporting into commercial systems lags, and access to the authoritative death file has been restricted, which narrows who can screen against it.
People with no credit file at all	No existing file means no contradiction. The constructed file becomes the only file.	Millions of adults have no credit record `[VERIFY a current figure before citing]`. A person who has never borrowed has nothing to check and no habit of checking.
Recent immigrants	A recently issued number with a genuinely thin file, which makes the constructed thin file unremarkable.	A thin file is expected here, so the strongest single indicator in the specification is legitimately present and cannot be used on its own.

A HEURISTIC THAT USED TO WORK AND NO LONGER DOES

Before assignment was randomised in mid 2011, a Social Security number encoded information about where and roughly when it was issued. That made it possible to flag an implausible pairing: a number whose issuance vintage did not fit the claimed date of birth. It was a genuine test of the *binding* between two elements rather than of either element alone, and it was cheap.

Randomisation removed that signal for every number issued since. The change was made for good reasons, including reducing the predictability that made numbers easier to guess. Its side effect was to retire one of the few binding checks that existed, and much of the detection burden this note describes is a consequence of not having replaced it.

3 The lifecycle, and where each control can see

Five phases. The institution's leverage is concentrated in phases one and two, and almost every control it owns activates in phase five.

PHASE 1	PHASE 2	PHASE 3	PHASE 4	PHASE 5
Days	**Weeks to months**	**Months to years**	**Months**	**Days**
Assembly	**Seeding**	**Cultivation**	**Amplification**	**Bust out**
A dormant number is paired with a fabricated name, a chosen date of birth, a controlled address, and live contact details.	Applications are made until one is granted. A single approved account creates the file that every later application will be checked against.	Small balances, paid on time. Sometimes accelerated by attaching the identity as an authorised user to a real, aged account.	Limits rise on the strength of the record. More institutions extend credit, each reassured by what the others have already done.	Every line drawn to its limit within a short window. Nothing is repaid. There is no person to pursue.
WHAT CAN SEE IT	WHAT CAN SEE IT	WHAT CAN SEE IT	WHAT CAN SEE IT	WHAT CAN SEE IT
Coherence checks between elements. Almost nothing else.	Application clustering across your own records. Element multiplicity.	File composition: authorised user structure, tradeline velocity, seasoning gaps.	Limit growth against stated capacity. Utilisation trajectory.	Everything, and too late. This is where the loss is booked.

THE TIMING PROBLEM, STATED PLAINLY

Phases 1 to 4 can run for two years or more, during which the account looks like a good customer and is treated as one. Phase 5 takes days. An institution whose detection activates on delinquency has a control that fires after the money is gone, every time, by construction.

This is why the feature specification separates measurements observable at origination from measurements observable after opening. The first group is where an institution can decline. The second is where it can still limit exposure. After that there is only recovery.

One further property of phase 4 deserves attention because it inverts an intuition most credit processes rely on. **Each institution's willingness to lend is treated by the next institution as evidence.** A file with four seasoned tradelines from reputable lenders reads as a good risk, and every one of those lenders reached that conclusion by looking at the tradelines the ones before

them had granted. The scheme borrows the sector's collective judgement, which no single participant is in a position to audit.

4

Control by control, why each one passes it

Not a criticism of these controls. Each does its job. The purpose of the table is to show that their jobs, taken together, do not add up to the question that matters.

CONTROL	WHAT IT TESTS	WHY A SYNTHETIC IDENTITY PASSES	RESIDUAL VALUE
Document verification	That a presented document is authentic and unaltered.	Often no document is presented at all in remote origination. Where one is, the underlying number is genuine, so a document issued on the strength of it is genuine too.	Effective against forgery and alteration, which is a different threat.
Database matching at account opening	That the name, date of birth, address, and identification number match records held about a customer.	The number is real and the other elements are internally consistent. Once a file exists from phase 2, the elements match the file the scheme itself created.	Catches careless fabrication and inconsistent reuse. Substantial and not sufficient.
Knowledge based authentication	That the applicant knows facts drawn from a credit file.	The questions are generated from the constructed file. **The operator built that file and knows the answers better than any real customer would.**	Close to none against this threat. Actively misleading, because passing it reads as reassurance.
Credit file check	Depth of history, payment record, existing obligations.	The file is real and was cultivated to satisfy exactly this check. Phase 3 exists to produce a good result here.	The raw score is inverted as a signal. The file's *composition* is informative; the score derived from it is not.
Device and network signals	Whether the device or connection is associated with prior abuse.	Operators rotate devices and connections, and there is no prior abuse associated with a first application from a new identity.	Real value in detecting phase 2 batches from one operator, if you retain and compare across applications.
Velocity rules	Too many applications in a short window from one applicant.	Applications are spread across time and across institutions. Within any one institution the pattern is unremarkable.	Useful against smash and grab attempts. Blind to a scheme with patience.
Manual underwriting	An experienced person's judgement on a file.	The file is genuinely good. An underwriter asked whether this is a sound credit will correctly answer yes, because on the evidence available it is.	High, but only when the underwriter is given the composition signals. Judgement without the right inputs reaches the wrong conclusion confidently.
Consent based verification against the issuing authority	**That the number, name, and date of birth belong together**, checked against the authority's own records.	It does not pass this. A fabricated name paired with a real number returns a mismatch, because this checks the binding rather than the elements.	**The strongest single control available.** Limits below.

THE HONEST QUALIFICATION, AND IT MATTERS

The last row is a genuine structural counter, and any account of this threat that omits it is out of date. Where an institution can verify the number, name, and date of birth together against the issuing authority's records, the composition error becomes visible at origination. That is precisely the check this note argues is missing, and it exists.

Four limits keep it from closing the problem. It requires the applicant's consent and a permitted purpose. It is a check at origination and says nothing about an account already open, which covers phases 3 to 5. It carries a cost and an integration burden that falls hardest on exactly the smaller institutions that are targeted. And it does not address the case where a real name is paired with a stolen number belonging to somebody else, or where the number's holder is deceased. **Adopt it where you can, and do not treat it as the end of the matter.** Confirm eligibility and current terms with your compliance officer.

5 Six structural properties that make this hard

Each one independently defeats a class of countermeasure, and they compound.

Component validityDefeats verification	Every element is real, so every element check returns a correct affirmative. Improving the checks improves the confidence with which the wrong answer is delivered.
PatienceDefeats velocity and anomaly rules	A scheme willing to wait two years produces no anomaly during the period when it can still be stopped. Detection tuned to unusual behaviour has nothing to detect until phase 5.
Absence of a complainantDefeats the reporting pipeline	No victim knows to complain. The number's holder is unaware, and the institution absorbing the loss usually classifies it as credit rather than fraud. Nothing enters the channel that would ordinarily surface a pattern.
Label absenceDefeats supervised learning	A scheme never drawn down is never labelled. One that is drawn down is labelled a charge off. A model trained on historical labels learns to predict charge off, which is a different and already solved problem.
Cross institution distributionDefeats single firm analytics	The identity is cultivated across several lenders deliberately, so no single institution observes the whole file's construction. Each sees a fragment that looks ordinary, and each is partly reassured by the others.

Inverted signal directionDefeats intuition	Excellent payment behaviour, low utilisation, and a clean record are the attack in progress. Several conventional risk indicators point the wrong way here, which is why an experienced underwriter without the composition signals reaches the wrong answer with confidence.

6 What detection has to do instead

Three shifts, each following directly from a property above. This is the reasoning that produced the feature specification.

Shift 1
From elements to coherence
Stop asking whether each field is valid and start asking whether the fields belong together, and whether the file's composition is consistent with the person it describes. A file that begins abruptly at forty, or whose apparent history was inherited through authorised user tradelines, is coherent nowhere except on paper.

Shift 2
From moments to trajectories
A single observation cannot distinguish cultivation from good citizenship. A trajectory can: limits rising faster than stated capacity supports, utilisation flat for a year and then climbing across every line at once, payment amounts too regular to be a household. This requires retaining history, which is the one preparatory step that cannot be done retrospectively.

Shift 3
From applicants to structure
Operators run batches, so identity elements recur across applicants who have no relationship. One address behind eight unrelated applications, or one number under three names, is visible in an institution's own records without any external data at all, and it is the cheapest signal available.

> **THE CHEAPEST OF THE THREE IS THE THIRD**
>
> Structural clustering needs nothing an institution does not already hold. Five measurements built on application history alone address it directly, and an institution with no bureau feed and no purchased identity data can build a usable manual review queue from them. The implementation guide sets out how, and it takes days rather than months.

7 Why smaller institutions carry more of this than their share

The concentration is not accidental and it is not about competence.

A scheme cultivating an identity across several lenders selects its seeding institution by one criterion: which one will grant the first account. That tends to be an institution with a thinner detection stack, which in practice means a smaller one. The first approval is the most valuable

thing in the whole operation, because it creates the file every subsequent application is measured against.

Smaller institutions are also least able to buy their way out. Commercial detection is priced for large books, sold through procurement cycles they do not run, and integrated into origination systems they do not control. The result is a tier that is targeted precisely because detection there is weakest, and is priced out of the remedy for the same reason it is targeted.

> **WHY OPEN PUBLICATION IS THE RELEVANT RESPONSE**
>
> Detection capability confined to the largest institutions leaves the scheme viable, because the scheme only needs one weak entry point per identity. **The same capability distributed across the smaller institution tier degrades it everywhere at once**, by removing the seeding opportunity rather than by catching the drawdown.
>
> That is the argument for publishing the specification, the implementation guide, and the code openly rather than licensing them. It is not a preference about openness. It follows from the fact that this particular scheme depends on the weakest participant, which means the value of a countermeasure rises with how widely it is held rather than with how exclusively.

8

What this note does not claim

Recorded so that the argument is not read as more than it is.

- **It does not quantify the problem.** No loss figure appears in this document, deliberately. There is no reliable denominator, because there is no complainant and the losses are usually classified as credit rather than fraud. Published estimates vary by an order of magnitude depending on definition, and any figure quoted without its definition should be treated as decoration.
- **It does not describe a detector that exists.** No synthetic identity model has been trained. The feature specification defines measurements; it has not been fitted or evaluated against any label.
- **It does not claim the listed controls are worthless.** Each is effective against the threat it was designed for. The argument is about a specific mismatch, not about the quality of anybody's origination stack.
- **It does not treat the scheme as static.** Operators adapt to whatever is deployed against them. Every measurement in the companion specification has a shelf life, and any institution adopting one needs a re validation date from the day it goes live.

- **It is not legal or regulatory advice.** Eligibility for consent based verification, permitted purpose, and adverse action obligations must be confirmed with your compliance officer and counsel.

> IF YOU READ ONLY ONE SECTION
>
> Read the table on the first page. Six correct answers about six real elements do not combine into an answer about whether a person exists. Everything else in this series follows from that sentence.

Argue with it

This note makes a structural claim, and structural claims are the kind that can be wrong in ways that matter. If a control in Section 4 is described inaccurately, if a limit stated about consent based verification has changed, or if the lifecycle does not match what your institution actually observes, that is worth reporting and it will change the document.

COMPANION DOCUMENTS

- Feature engineering specification [LINK]
- Implementation guide for smaller institutions [LINK]
- Reference implementation [LINK]
- Participant facing module on account takeover and synthetic identity [LINK]

REUSE

Copy this into your own staff training, your board paper, or your regulator submission. Attribution is the only request, so that a reader can check where it came from.

CORRECTIONS

Report a control described inaccurately, a changed regulatory position, or a phase your institution observes differently: [CONTACT]

Ayeni, A. [YEAR]. Threat model note: why field validation fails against synthetic identities, version [VERSION]. [DOI] · Licensed CC BY 4.0.