

SBA 7(a) loan vintages: the number a young book shows you is not the number it will cost


A loan book's charge-off rate is not comparable across vintages of different ages. The 2023 vintage shows 3.5% and projects to 12.2%; at 36 months a vintage has realised only 23.5% of the losses it will suffer. Two of the three hypotheses recorded before the analysis were contradicted by it.

DOMAIN	TOOLS	SCALE
Private credit / small-business lending	SQL · DuckDB · Python · matplotlib	1,697,539 disbursed loans · 35 vintages
SOURCES	WINDOW	PUBLISHED
SBA 7(a) FOIA extract · SBA performance report	FY1991–FY2026, as of 30 June 2026	2 September 2026

Corrected for age, the 2023 vintage is on track for 12.2% losses — the worst since 2009

Cumulative charge-off rate of SBA 7(a) loans by fiscal year of approval. The lower block is the loss already recorded; the upper block is what the maturation of completed vintages implies is still to come.

FY2015–FY2025 · 1,697,539 DISBURSED LOANS · SHARE OF LOANS, NOT DOLLARS



Source: SBA 7(a) FOIA loan-level data, as of 2026-06-30 (U.S. Government Works) · Projection: observed rate x the median development factor of the 24 completed vintages. Whisker = interquartile range of those factors.

JA

Loss already recorded, and above it what the maturation of the completed vintages implies is still to come. The two youngest vintages carry no projection at all.

The eight phases, and what each one settled

Every phase closes on an exit gate checked point by point. When one fails the process goes *back* a phase rather than improvising forward. Here nothing had to go back — but two phases had to correct what an earlier one had got wrong, which is the same discipline.

0 · CHOOSE A client that has to price loss it has not yet seen

A fund buying in today, not an originator reviewing its own back book: only the buyer has to commit to a number for loans too young to have defaulted. Source chosen after an integrity test on the raw file, not before.

1 · ASK A SMART question, and three hypotheses written down first

How does the cumulative charge-off rate evolve with months since approval, per vintage — and what terminal loss do the immature ones imply when read at equal age? The expected answers were recorded before any analysis, which is the only way a later contradiction is credible.

2 · PREPARE The field everyone discards turns out to be the mechanism

ROCCC on the source, six biases written down, and the integrity test on all 1.96 M rows. EXEMPT is defined by the publisher as disbursed-but-not-yet-resolved: not missing data, but the censoring indicator handed over explicitly.

3 · PROCESS A reconciliation that has to hold, asserted in code

1,961,455 approvals down to 1,697,539 disbursed loans, and the pipeline stops if that stops balancing. An assertion caught 32 charge-offs dated after the file's own cut-off — the file contradicting itself.

4 · ANALYSE A 138-month horizon derived by rule, and seven checks

The terminal horizon came from a rule written before it was computed, and came out far longer than expected. V0 to V6; V3 confirmed the bias that most threatened the conclusion rather than clearing it.

5 · SHARE Four figures, four forms, a palette that was measured

No shape reused against the earlier two cases. The palette was validated with a colourblind and contrast checker; a first candidate failed and was discarded rather than nudged.

6 · ACT Three recommendations, and a reproducibility defect found

The dollar-weighted figure needed its own development factors, not a division. Rebuilding the database twice and diffing hashes exposed a non-deterministic pipeline.

7 · PORTFOLIO Published against a contract enforced by code

The schema was verified negatively — a required field was removed to confirm the build fails. The contract grew by one optional field, which is the only way a contract with published cases can grow.

The decisions that defined the case

Eight of the twenty-four recorded in **CASO.md**. The discarded alternative is the line that matters: it shows the choice was reasoned rather than reflexive.

EXEMPT loans stay in the denominator

The publisher defines the status as disbursed but not cancelled, paid in full or charged off — alive at the cut-off. It is the censored population, not an unknown.

Discarded: Dropping them as "outcome unknown", which deletes 297,494 loans and almost the entire recent population, and makes new vintages look nearly default-free.

The SBA FOIA extract, not Freddie Mac or Lending Club

Public domain, no registration, loan-level since 1991, published by the guarantor itself.

Discarded: Freddie Mac, whose licence restricts use to internal purposes and collides with publishing; Lending Club, no longer distributed from origin, so only copies of copies remain.

The terminal horizon is derived from a rule, not chosen

Defined as the age at which completed vintages have realised 95% of their charge-offs, and written down before being computed. It came out at 138 months.

Discarded: A round 120, which looks tidier and would have cut the horizon six points of realisation short.

Charge-offs dated after the cut-off keep their loan and lose their date

They count in the terminal rate because they happened; they cannot sit on the curve because their timing is untrustworthy.

Discarded: Clamping the dates to the cut-off, which invents an age nobody observed and piles 32 loans into one month at the far right — exactly where the case makes its claim.

The two youngest vintages are published as a gap, not a number

Their development factors are $17\times$ and $606\times$. Multiplying a near-zero rate by 606 produces a confident-looking figure resting on nothing.

Discarded: Publishing them with a wide band, which would have repeated the case's own headline error in its final chart.

The dollar projection gets its own development factors

The dollar/count ratio falls with age — 0.456 at 138 months, 0.226 at 33 — so the two curves mature at different speeds.

Discarded: Multiplying the dollar rate by the count factors: the obvious shortcut, and an invented number.

The term finding ships with its confound in the card

Term proxies product and collateral: 20-year loans are real-estate secured, short ones are working capital. It says where the loss sits, not that tenor causes it.

Discarded: Presenting the fourteen-fold spread as a lever, which a risk analyst dismantles with the first question.

The mobile layout defect was reported, not patched

The already-published football case truncates its body text at 390 px in exactly the same way, so it is site-wide.

Discarded: A local fix on the new page, which would have made the newest case look fine and left the other two broken.

The moment the case nearly went the other way

A pipeline that promised reproducibility and did not deliver it

The README tells a third party they can rebuild this case from the raw files. In phase 6, rebuilding the database twice and diffing the SHA-256 of every exported file showed **half the rows of the size-cut export changing between runs**.

Approved amounts cluster on round numbers — \$50,000, \$150,000, \$500,000 — so ordering the quartile window by amount alone left thousands of ties, and `ntile` split them by whatever order the rows happened to arrive in. DuckDB does not guarantee that order when it reads four files in parallel.

It was caught by *running the check*, not by reading the code. Nothing in the SQL looks wrong, and every figure in the case was correct — but a reader who rebuilt it would have got different numbers for one cut and had no way to know which of us was right.

SYMPTOM	CAUSE	FIX
Half of <code>curva_por_importe.csv</code> changed between rebuilds	<code>ntile</code> split tied amounts by row arrival order	Tie-break on date, term, status, charge-off date, lender, NAICS
<code>usd_en_riesgo</code> serialised as ...259.01 then ...259.0099998	Summing millions of doubles in parallel is not associative	Round the sums to cents

The tie-break deliberately includes `estado` and `fecha_fallido`, so any tie that survives is between loans with the *same outcome* and moving them between quartiles cannot move the curve. Two full rebuilds are now byte-identical. **The two-run hash diff is cheap and belongs in any case that promises reproducibility** — this one had promised it in writing for two phases before anyone checked.

What the data says


Each finding was written as a sentence carrying a number before any chart was drawn — the last check that it was closed rather than a topic.

1 · At three years a vintage has shown only a quarter of what it will cost

At three years a vintage has shown only a quarter of the losses it will suffer

Share of a vintage's eventual charge-offs that has already occurred, by months since approval. Measured on the 24 vintages old enough to be complete.

FY1991–FY2014 · 138-MONTH TERMINAL HORIZON · MEDIAN ACROSS VINTAGES



Source: SBA 7(a) FOIA loan-level data, as of 2026-06-30 (U.S. Government Works) · The horizon is the age at which the median completed vintage has realised 95% of its charge-offs. Derived, not chosen.

JA

Share of eventual charge-offs already recorded, by months since approval, across the 24 completed vintages.

23.5% at month 36, 56.6% at month 60, 97.0% at month 120. The median charge-off arrives at month 58 — year five, not year three. This contradicts the third hypothesis and is why the horizon had to be 138 months.

2 · Corrected for age, 2023 is the worst vintage since 2009


3.46% observed against a projected 12.23% (band 10.28–13.35). Every completed vintage from 2010 to 2014 landed between 5.99% and 9.20%, so even the bottom of the band clears them all. In dollars: 0.75 cents per dollar approved becomes 4.24.

3 · The correction changes the level, not the league table

Correcting for age changes the level, not the league table: the five worst vintages are the same

Rank of each vintage by cumulative charge-off rate, read four ways: the raw rate, and the rate at a common age of 36, 60 and 84 months. Rank 1 is the worst.

28 VINTAGES WITH ALL FOUR READINGS OBSERVABLE · FY1991–FY2018



Source: SBA 7(a) FOIA loan-level data, as of 2026-06-30 (U.S. Government Works) · Rank correlation between readings is 0.92 to 0.96.
This contradicts the hypothesis recorded before the analysis.

JA

Rank of 28 vintages under four readings. The five worst are the same in all of them.


Rank correlation 0.92–0.96. **This contradicts the second hypothesis**, and it sharpens the case: the error a fund makes is not ranking the wrong vintage worst, it is believing a young vintage's number.

4 · Two vintages are policy, not performance

The 2020 and 2021 vintages default at half the rate of their neighbours — because the government was paying the instalments

Cumulative charge-off rate over the first 45 months, one panel per vintage, identical axes. CARES Act section 1112 covered principal and interest on many 7(a) loans through the pandemic.

FY2018–FY2022 · FIRST 45 MONTHS · SHARE OF LOANS



Source: SBA 7(a) FOIA loan-level data, as of 2026-06-30 (U.S. Government Works) · Read as credit quality, these two vintages would flatter any underwriting model built on them. They are policy, not performance.

JA

Identical axes across five panels: same age, half the default rate, for a reason that is not credit.

At 36 months FY2020 sits at 0.85% and FY2021 at 0.77%, against 1.58–1.98% for their neighbours. CARES Act section 1112 had the government paying instalments. A model calibrated on those two years understates everything else by about 18%.

What each check ruled out

Seven blocks in `verificar.py`. One of them did not clear its suspicion — it confirmed it, and the conclusion was labelled instead of published clean.

CHECK	WHAT IT RULED OUT	RESULT
V0 · Zero-amount charge-offs	That 180 zero-value charge-offs sit in one place and corrupt the dollar metric	Spread across 25 vintages; readable
V1 · External reconciliation	That the population was built wrong	Max deviation 4 loans across nine fiscal years
V1b · Recovery rates	That "gross, not net" is an unbounded caveat	Mature cohorts recover 34–39%
V2 · Count against dollars	That the finding is about loan size, not vintage	Rank correlation 0.94
V3 · CARES Act relief	That FY2020–21 look good on credit quality	NOT ruled out — bias confirmed, vintages labelled
V4 · Vintage or mix	That the pattern is a change of origination mix	Rank correlation 0.97 after standardising
V5 · Choice of common age	That the conclusion depends on which age is chosen	0.92–0.96 across 36, 60 and 84 months
V6 · Recompute by another route	That the curve itself is a bug	140 cells, 0 discrepancies

V1 is worth a note. It was carried as an open limitation for two phases because the SBA archive that would settle it returned 404 — until the page turned out to link the file with a *relative* URL that only resolves against a legacy host. The check was never impossible; the first attempt resolved the link against the obvious host and took the 404 at face value.

Conclusions, limits and what this case demonstrates

Recommendations

- **Price the book off a maturation curve, not off its observed rate.** For a 2023-like book the loss input moves from 0.75 to 4.24 cents per dollar approved — the uncorrected figure understates by 5.7×.
- **Keep 2020 and 2021 out of any calibration sample,** and use 2016–2019 as the modern reference window.
- **Underwrite the short-term, small-ticket segment on its own terms** instead of at a portfolio average — while remembering that term is a proxy for product, not a lever.

What this case cannot tell you

These loans carry a partial government guarantee: the *shape* of the maturation curve transfers to unguaranteed private credit, the *level* does not. Every rate is gross — bounded at roughly a third by SBA's published recovery tables, but not converted, because recovery is measured against a different amount on a different year axis. And the projection assumes the loss *timing* of 1991–2014 still holds; if today's borrowers fail on a different schedule, the method is wrong in a way it cannot detect on its own.

What it demonstrates

The reflex that separates a credit analyst from someone who can group rows in SQL: knowing that a raw default rate is not comparable across cohorts of different ages, and correcting it. It is not a difficult calculation — it is a difficult *habit*, because the uncorrected number is always available and always looks fine.

And the refusal. The method produces a number for the 2025 vintage. Publishing it would have made the case look more complete and would have been the same mistake it spends four charts exposing.

Full phase log, cleaning log and the twenty-four recorded decisions: [CASO.md](#) in the case repository.