

Reclassified to Conform

Chris Martin
Kepler AI
`chris.martin@kepler.ai`

September 22, 2026

Abstract

A company’s financial statements over many years must be assembled from many filings. Each filing may rename, combine, split or move its lines, and restates only the last two or three years to match. Filings provide no consistent, machine-readable mapping between successive statements, and neither a line’s label nor its XBRL tag is fixed. We treat the reconstruction as an optimization problem. A tag that keeps its numbers from one filing to the next is taken to be the same line, and such runs are joined into strands before any edit is weighed. At each boundary between filings, the candidate edits that pass the check their kind requires are listed, each with the strand endpoints it would explain and a price for how much it asks the reader to believe beyond what the filings show. An integer program picks the cheapest set of the listed edits that accounts for every line exactly once, and the chosen edits trace each line through the filings. The result is a statement whose rows follow one line across every year, comparable within each filing’s restated window and marked where a number was restated or derived, in any filing’s layout, with each number cited to the filing that stated it and each declared total re-checked. On 200 companies drawn at random, 99.9% of the totals that can be checked add up. Among the 124 with a usable vendor comparison, the lines the method follows agree with the vendor’s series on 96% of the cells both fill, reach 10% more cells than a tag alone and agree with the vendor on 90% of those. The method also yields a record of what each company changed in its reporting.

*When I use a word, it means just what I choose it
to mean, neither more nor less.*

— Humpty Dumpty, in *Through the Looking-Glass*

1 The problem

Company filings are the case this paper takes up. Anyone analyzing a company wants to see how it has fared over time. A snapshot is not enough; a forecast or a model needs the course of the business across many years, and so many years of reporting.

Reported data is plentiful: public companies file their financial statements with the SEC every quarter and every year, every number tagged and machine-readable. The measures are standardized by GAAP and audited, so each number is reliable for the period it was filed. But which lines a company reports, and how finely it breaks them down, is largely its own decision [7, 27, 31], and those decisions change.

Apple’s cash flow statement is the small case that runs through the paper (Figure 1). Its report for fiscal 2021 shows *proceeds from issuance of common stock* of 880 for 2020. The report for fiscal 2022 has no such line, and *other* for 2020 is 754, the old 880 plus the old (126) [1, 2]. Both numbers were true when filed. This is the easiest case: one line moved, and the arithmetic says where. Usually several lines move at once and more than one explanation fits, and sometimes the company gets it wrong (Section 4, Appendix C). This paper states the question precisely and answers it for every line of every statement.

(a) 10-K for fiscal 2021, cash flow statement

	Years ended		
	September 25, 2021	September 26, 2020	September 28, 2019
Financing activities:			
Proceeds from issuance of common stock	1,105	880	781
Payments for taxes related to net share settlement of equity awards	(6,556)	(3,634)	(2,817)
Other	(129)	(126)	(105)
Cash used in financing activities	(93,353)	(86,820)	(90,976)

(b) 10-K for fiscal 2022, cash flow statement

	Years ended		
	September 24, 2022	September 25, 2021	September 26, 2020
Financing activities:			
Payments for taxes related to net share settlement of equity awards	(6,223)	(6,556)	(3,634)
Other	(160)	976	754
Cash used in financing activities	(110,749)	(93,353)	(86,820)

Figure 1: The financing section of Apple’s cash flow statement in two consecutive annual reports, cut from the filings; the dotted lines mark omitted rows. Each report shows three years; the two share 2021 and 2020, and 2019 appears only in the first [1, 2].

The figure also shows what makes the question answerable. Two filings a year apart differ because the presentation moved and because the business moved, and comparing their newest years cannot tell one from the other. But each filing shows the prior years again: the fiscal 2022 report presents 2020 under its own lines, and the fiscal 2021 report presented the same 2020 under its. With one period under two presentations, the change in presentation can be read apart from the change in the business. That is the evidence the paper works from. It constrains the answer but does not always fix it: several potential sets of edits can fit the same numbers, and Section 4.2 shows a boundary that admits three.

The paper offers a way to discover what changed and a method for rebuilding comparable history where it exists. Its title is the sentence companies print whenever they change their reporting [5]:

Certain amounts in prior periods have been reclassified to conform to the current period presentation.

1.1 The mask

A company has financial properties that outsiders never see directly. A filing shows what the company chooses to show: each line is a total over things the reader cannot see. Call that choice the *mask*. The mask is what changes when a company changes its reporting; what we want is the properties underneath, over time. Names, labels and position on the page say something about how two masks differ, but not much. What separates them is the re-presented years: this year’s filing

applies this year’s mask to last year’s state, and last year’s filing applied last year’s mask to the same state.

The mask can move in many ways. A few recur, and each is drawn in Appendix B. Three change the detail. A line continues under a new name or tag, as when Netflix’s deferred revenue moved to the tag of a new accounting standard [15, 16]. One line becomes several, as when JPMorgan separated asset-management fees from commissions [13, 14]. Several lines become one, as when Apple’s stock proceeds disappeared into *other* [1, 2].

Two change the numbers. Money moves between lines: a line keeps its name, part of what it held now sits elsewhere, and the prior years are restated to match, as when Tesla took leasing out of automotive revenue [22, 25]. Or the company gets it wrong: Palantir attached three years of numbers to the wrong lines [18–20]. And a number can be revised with nothing about the layout changing.

Which lines exist is not enough to tell these apart. A line gone from the newer filing may have ended, gone into a neighbour, been combined with others into a new line, or been rolled into its total, and all four leave the newer filing looking the same. Only the numbers on the year both filings report tell them apart. Section 3 turns this into a precise list of the kinds of edit, and draws the four (Figure 3).

1.2 Analogues in other fields

The problem is not specific to accounting. Wherever something is observed only through a series of reports, each under a classification the reporter is free to change, anyone who wants the thing itself over time has to reconcile the reports first. Every organization that maintains a classification builds the same three things: a list of the kinds of change, a record of which old codes correspond to which new ones, and a warning that combining cannot be undone. The Census Bureau publishes relationship files whenever tracts are redrawn, typed and weighted [26], and the mapping from ICD-9 to ICD-10 medical codes is published in both directions and documented as lossy [4]. Economists keep every published version of a statistical series, so that what was known on a given date can be recovered [6]; that answers what was known and when, and reconciles nothing, since a vintage archive keeps every label as it was.¹ The difference from company filings is who controls the classification. The Census Bureau and the code-set maintainers publish their mappings because they made the changes. In company filings there is no such authority. Thousands of companies each change their own list of lines on their own schedule, and what they share is a vocabulary of tags, not a list of lines. A company knows how its lines changed, since it restates the prior years to conform, but its filings state that only as a sentence in a note, if at all, and no one outside the company has it. So the mapping has to be recovered from the filings themselves.

2 The filings as data

2.1 The structure of a filing

A company files its financial statements in two forms: the document people read, the 10-K or 10-Q, and XBRL, the eXtensible Business Reporting Language, in which every number is a *fact* a machine can read [33]. A fact is a number plus what it measures (a tag, from the shared `us-gaap` dictionary

¹Also: git guesses that a file was renamed from how similar its contents are [10], and Mercurial records the rename [17]; country codes are never reused once retired, after `CS` was assigned twice [12]; trade economists group product codes into families that stay comparable across revisions [21]; and finance has a permanent identifier for a company [3] and nothing like it for a line of its statements.

or the company’s own), the period it covers, its unit, and how it was rounded. A breakdown by product or segment is a fact on the same tag, marked with the product or segment, not a tag of its own. Figure 2 shows two facts as filed. A line of a statement is one tag’s facts, one for each period the statement shows.

(a) Apple, 10-K for fiscal 2021: the number printed as “880”, the context it points to, and its unit

```
<ix:nonFraction unitRef="usd"
    contextRef="i725ce765160e48a2a3b36989bda2ca94_D20190929-20200926"
    decimals="-6"
    name="us-gaap:ProceedsFromIssuanceOfCommonStock"
    scale="6"
    id="id3VybdovL2...3642a79d-e216-4edd-aaef-8b22f84f7992">880</ix:nonFraction>

<xbrli:context id="i725ce765160e48a2a3b36989bda2ca94_D20190929-20200926">
  <xbrli:entity>
    <xbrli:identifier scheme="http://www.sec.gov/CIK">0000320193</xbrli:identifier>
  </xbrli:entity>
  <xbrli:period>
    <xbrli:startDate>2019-09-29</xbrli:startDate>
    <xbrli:endDate>2020-09-26</xbrli:endDate>
  </xbrli:period>
</xbrli:context>

<xbrli:unit id="usd">
  <xbrli:measure>iso4217:USD</xbrli:measure>
</xbrli:unit>
```

(b) Tesla, 10-K for fiscal 2021: automotive leasing revenue, a revenue tag whose context names a product

```
<ix:nonFraction id="F_abccb4fb-2668-4799-8607-ec8ce640f676"
    contextRef="C_bae6b89d-012c-4639-b3ad-49711e152c22"
    name="us-gaap:RevenueFromContractWithCustomerExcludingAssessedTax"
    unitRef="U_USD"
    scale="6"
    decimals="-6"
    format="ixt:numdotdecimal">1,052</ix:nonFraction>

<xbrli:context id="C_bae6b89d-012c-4639-b3ad-49711e152c22">
  <xbrli:entity>
    <xbrli:identifier scheme="http://www.sec.gov/CIK">0001318605</xbrli:identifier>
    <xbrli:segment>
      <xbrldi:explicitMember dimension="srt:ProductOrServiceAxis">
        tsla:AutomotiveLeasingMember
      </xbrldi:explicitMember>
    </xbrli:segment>
  </xbrli:entity>
  <xbrli:period>
    <xbrli:startDate>2020-01-01</xbrli:startDate>
    <xbrli:endDate>2020-12-31</xbrli:endDate>
  </xbrli:period>
</xbrli:context>
```

Figure 2: XBRL facts as filed. A fact is a number with a tag (**name**), a context (whose number, and for which period), a unit, and its rounding (**decimals**; **scale="6"** means the page shows millions). A breakdown is the same kind of fact, with the product or segment named in its context [1, 24].

2.2 The parts of a filing and their reliability

Four parts of a filing matter here: the facts, the layout, the arithmetic, and the prior periods the filing shows again. The facts give the values. The other three tell us different things and can be relied on to different degrees.

The layout gives the order of the page and nothing more. A heading is an abstract element, which the XBRL specification bars from carrying a fact [33], and none does; indentation says where a line sits, not what it adds up to.

The arithmetic says which lines add up to which totals, and with what sign, stated once per tag in a separate file of the filing. A filing must declare it wherever a statement shows lines and their total [31], but a total that does not add up has no effect on whether EDGAR accepts the filing [30, 31], so how often the equations hold is a measurement. Over the 200 companies drawn at random for Section 5 and their 2,826 annual reports, 87.7% of the 104,408 declared totals add up exactly as declared, 89.8% once breakdown rows are folded into their line and the filer’s rounding is allowed, and 98.7% among totals whose every part carries a number; most of the rest have a declared part with no number in that period (Appendix E.1). Breakdowns are the exception: a declared sum relates one tag to another in the same context, so a line’s breakdown is never declared to add up to it [33], and a breakdown can confirm a number but cannot prove one wrong.

The prior periods tell us the most. When a filing shows last year again, it says that these are last year’s numbers under this year’s lines: the two masks over one state that Section 1.1 rests on. The window is short: two balance sheets and three years of income and cash flows [28]. Older years are never redone, so they have no line in the newest filing. The history we want is ten or fifteen years, so no single filing can supply it; each boundary between consecutive filings has to be read on its own, from the two masks that pair shows, and a line’s history is the chain of those readings. Figure 1 in Section 1 shows the window as printed: two consecutive Apple 10-Ks, each with three years, sharing two.

That is all that carries from one filing to the next. No rule requires a filing to map its lines onto the one before. The accounting standard requires only a note saying that prior periods were reclassified to conform, with the nature and size of the change, in words [7]. Nothing in a filing states, in a consistent, machine-readable form, which of its lines continue an earlier filing’s or what became of a line that is gone; the note, where there is one, says so in words. The equations and the re-presented years are the only commitments a company makes across filings, and the rest of the paper builds on those two alone.

3 Formulating the problem

For every measure a line of a company’s statements represents, we want its history: which line it was in each earlier filing, which of its numbers were restated, and where the comparison breaks. The filings do not state it, so it has to be inferred, and usually more than one explanation of a boundary fits the numbers. That is the shape of an optimization problem: an edit involves only a few lines, so the candidates can be listed; a solution is a set of edits that accounts for each line of both filings exactly once, an explanation of the boundary; and among the explanations we want the one that asks the least of the reader beyond what the filings show.

3.1 Input and output

The input is one company’s filings of one statement, in the order published. Each presents the statement as a list of lines; a line has a tag, a label, a place in the layout, and a number for the current year and one or two prior years. Each filing also declares which lines add up to which totals. So every number has two dates, the period it describes and the filing that stated it, and most years are stated by two or three filings that do not always agree. Notes may describe how lines changed, but provide no consistent, machine-readable mapping between filings.

The output is the statement’s history: between each pair of consecutive filings, a set of edits that turns the older list of lines into the newer one, with every line of both accounted for exactly once. A line continues, under its old name or a new one; or it takes part in an edit with other lines; or it ended or started. Nothing is left over and nothing is used twice. The rest follows: rows that

run across years, the filing behind each number, where a comparison breaks, and the table as first or as last reported, which differ only in which filing each year’s number is taken from.

3.2 The kinds of edit

An edit says which lines of the older filing correspond to which of the newer. Each kind is defined by what the numbers must show on the years both filings report, and each moves the mask of Section 1.1 in its own way (Table 1). Three kinds add detail: a **split**, a **carve-out** and a **break out**. Three remove it, each undoing one of the first three: a **merge** undoes a split, an **absorb** a carve-out, a **fold** a break out. A carve-out and an absorb move money between lines, so a continuing line is restated; the other four leave every continuing number as it was. A **rename** moves nothing. **Equal copies** and a **swap** are rarer, and a line nothing explains has **ended** or **started**. Figure 3 draws the four ways a line can be gone from the newer filing that Section 1 listed, each named: a line that ended, an absorb, a fold and a merge.

KIND	WHAT HAPPENED	WHAT THE NUMBERS MUST SHOW	EFFECT ON DETAIL
rename	the same line continues under a new name or tag	the number carries over	none
split	one line ends; new lines are its parts	the parts add up to the old line	more
carve-out	a line that continues hands part of itself to new lines	the old line is restated down by exactly what moved	more; money moves
break out	a total starts showing its parts	the total doesn't change, and the new parts add up to it	more
merge	several lines end; one new line is their total	the parts add up to the new line	less
absorb	lines end inside one or more lines that continue	the receiving lines are restated up by exactly what moved	less; money moves
fold	the parts stop being shown; their total continues	the total doesn't change, and the parts added up to it	less
equal copies	one “basic and diluted” line becomes separate basic and diluted lines, or the reverse	each copy equals the original	more or less
swap	the company attached numbers to the wrong lines	each column of numbers reappears, digit for digit, on another line	none; numbers misplaced
ended, started	a line has no counterpart	nothing; the fallback when nothing else explains it	less; more

Table 1: The kinds of edit. Two more touch numbers only: a **restatement**, a prior year’s number revised, and a **transfer**, part of a split line moving to another section while that section’s total is restated by the same amount. Every kind is drawn, with a made-up and a real example, in Appendix B.

(A) ENDED – THE LINE ENDED

Filing t			Filing t+1		
	Y0	Y1		Y1	Y2
Line A	5	0			
Line B	25	35		35	40
Total	30	35		35	40

← Y1 in both →

ended

nothing is restated: A was 0 in Y1, and nothing in the newer filing accounts for it

(B) ABSORB – IT WENT INTO A NEIGHBOUR

Filing t			Filing t+1		
	Y0	Y1		Y1	Y2
Line A	20	25			
Line B	10	10		35	40
Total	30	35		35	40

← Y1 in both →

went into B

restated

the check, on Y1

$$\begin{array}{r} 25 \\ A \end{array} + \begin{array}{r} 10 \\ \text{old B} \end{array} = \begin{array}{r} 35 \\ \text{new B} \end{array}$$

(C) FOLD – IT WAS ROLLED INTO ITS TOTAL

Filing t			Filing t+1		
	Y0	Y1		Y1	Y2
Line A	20	25			
Line B	10	10			
Total	30	35		35	40

← Y1 in both →

rolled in

unchanged

the check, on Y1

$$\begin{array}{r} 25 \\ A \end{array} + \begin{array}{r} 10 \\ B \end{array} = \begin{array}{r} 35 \\ \text{total, in both} \end{array}$$

(D) MERGE – IT WAS COMBINED WITH OTHERS INTO A NEW LINE

Filing t			Filing t+1		
	Y0	Y1		Y1	Y2
Line A	20	25		35	38
Line B	10	10			
Line E	5	5		5	6
Total	35	40		40	44

← Y1 in both →

combined into C

new line

the check, on Y1

$$\begin{array}{r} 25 \\ A \end{array} + \begin{array}{r} 10 \\ B \end{array} = \begin{array}{r} 35 \\ \text{new line C} \end{array}$$

Figure 3: Four ways a line disappears, each named as a kind of edit from Table 1. In every case line A is gone from the newer filing; only the shared year, Y1, says which happened: nothing restated (ended); a neighbour restated up by exactly the line's amount (absorb); the total unchanged and the parts gone (fold); a new line equal to the old ones' sum (merge). Each side is one filing's statement, older on the left; the shaded column is the year both report; the edit is drawn in blue (dashed where parts are shown or hidden and no number moves); a restated number is circled in orange.

3.3 Multiple explanations

Several explanations can fit the same numbers; Tesla’s 2016 income statement admits at least three (Section 4.2). Whether a line is accounted for depends on every other decision at the boundary, so the explanations cannot be settled line by line, and Section 4 develops the price used to choose among them.

4 Selecting the explanation

4.1 Choosing edits together

Candidate edits compete for the same lines, so selecting them is a weighted set-packing problem. Although this problem is hard in general, the instances here decompose into small connected components and are inexpensive to solve exactly. We therefore use an integer program, avoiding a problem-specific ordering heuristic. Across the random sample, all component solves together took a quarter of a second (Section 4.4.3). Think of a market where edits compete to explain the lines. Each bids a price, no line can be sold twice, and the market clears at the lowest total cost, including the cost of lines left unexplained.

4.2 A simple example

A small case, to show the kinds of edit that can explain one boundary and how they compete for the same lines. Tesla’s 2015 and 2016 annual reports are the case this section and Appendix C return to. The 2015 report shows Automotive revenue of 3,741.0. The 2016 report shows, for the same year, Automotive 3,431.6 under the same tag, a new Automotive leasing line of 309.4, and a new Total automotive revenue of 3,741.0 under a new tag [22, 25]. Three explanations fit, each a set of edits accounting for every line of both reports (Figures 4 to 6).

A, a rename alone. The old line became “Total automotive revenue”; its number carries over exactly. The 2016 Automotive line then keeps the old tag but has to be a new line, and neither its 3,431.6 nor the 309.4 is explained by anything.

EXPLANATION A: THE OLD LINE WAS RENAMED TO THE NEW TOTAL

Tesla, Inc. • Form 10-K, fiscal 2015
0001564590-16-013195
Consolidated statements of operations • \$ millions

	2014	2015
Automotive SalesRevenueGoodsNet	3,007.0	3,741.0

← 2015 in both →

started

started

renamed

Tesla, Inc. • Form 10-K, fiscal 2016
0001564590-17-003118
Consolidated statements of operations • \$ millions

	2015	2016
Automotive unexplained SalesRevenueGoodsNet	3,431.6	5,589.0
Automotive leasing unexplained OperatingLeasesIncomeStatementLeaseRevenue	309.4	761.8
Total automotive revenue tsla:SalesRevenueAutomotive	3,741.0	6,350.8

one number carries over; 3,431.6 and 309.4 come from nowhere

$$\begin{array}{ccc} \underline{3,741.0} & = & \underline{3,741.0} \\ \text{old automotive} & & \text{new total} \end{array}$$

Figure 4: Explanation A, a rename alone.

B, a carve-out. Leasing was taken out of Automotive, which continues restated ($3,741.0 = 3,431.6 + 309.4$), and the total is a new line that adds the two up.

EXPLANATION B: LEASING WAS CARVED OUT OF AUTOMOTIVE




Figure 5: Explanation B, a carve-out.

C, a rename plus a break out. The old line became the total, with its number carried over, and the total now shows its parts: two new lines that add up to it.

EXPLANATION C: RENAMED TO THE TOTAL, WHICH NOW SHOWS ITS PARTS




Figure 6: Explanation C, a rename plus a break out.

B and C both account for every number and differ in what they ask the reader to believe: B, that Automotive was restated and a total appeared; C, that the line changed tag and label and two detail lines appeared. They leave different histories. Under C the total runs 1,921.9, 3,007.0, 3,741.0, 6,350.8 like for like [22, 25], because Tesla's older filings counted leasing inside automotive revenue; under B the Automotive row mixes older years that include leasing with restated years that do not.

4.3 The procedure

The procedure has four steps (Appendix A.1 states them as pseudocode).

1. **Strands.** Join, before anything is weighed, the same tag in consecutive filings with the same numbers on the shared years. Each run is a strand; where a strand ends or starts is an *endpoint*, and endpoints are what is left to explain (Section 4.4.1).
2. **Candidates.** At each boundary, list the candidate edits of every kind: each could explain some of the endpoints and passes the check its kind requires, and is listed with the endpoints it would cover and a price (Section 4.4.2).
3. **Choose the edits.** Find the cheapest set of candidates that covers no endpoint twice, counting an endpoint no candidate covers at the price of leaving it unexplained (Section 4.4.3). Candidates that share no endpoint never interact, so the choice splits into independent problems, one per connected group of candidates, each solved exactly.
4. **Rows.** Read the history off the choice: strands joined by the chosen renames are rows, and every year's numbers come from one filing (Section 4.5).

Only the third step involves a decision. Appendix C.1 walks through the steps on a small example.

4.4 The formulation

Write $F_1, \dots, F_T$ for one company's filings of one statement (Section 3.1), in the order published, with boundary t between F_{t-1} and F_t . Within one filing, $v(k)$ is a tag's number for period k and $d(k)$ its **decimals**, with a subscript to say which tag when it matters; each tag also has a label, the total it is declared to add into, and a weight of +1 or -1.

4.4.1 Strands

A strand is a run of one tag across consecutive filings joined on the strongest evidence the filings offer, the same numbers on the shared years, before any edit is weighed. It is an assumption: a line's definition can change without changing the numbers both filings show, and such a change is invisible to the method. Two numbers *agree* when they differ by no more than half the coarser of their two rounding units, or by no more than ε if that is larger:

$$|v - v'| \leq \max\left(\frac{1}{2} 10^{-\min(d, d')}, \varepsilon\right), \quad \varepsilon = 0.5 \text{ (0.005 per share)}.$$

Here d is the fact's **decimals**, the number of significant digits after the decimal point; Apple's 880 million in Figure 2 carries $d = -6$, rounded to the nearest million, so any number within half a million of it agrees. A number is *material* when $|v| > \varepsilon$.

A tag that appears in F_i and F_{i+1} joins across the boundary when all three hold:

1. the two unit classes do not conflict;
2. no period both filings report disagrees;
3. some material period agrees: a number carried over.

Anything else breaks the run: a restated number, a missing filing, a changed unit, or nothing to compare, when the two filings share no period with a material number: a line that is always zero, a line stated for one date only, a change of fiscal year, or a filing whose comparatives our copy lacks. A break does not mean the two are not the same line; it leaves a pair of endpoints to explain, and a cheap edit may explain them. A heading carries no number and continues with its tag: it is layout, not a line. The rule is conservative because a join is irrevocable, while a missed join is made again whenever the corresponding edit is in the solution.

Each strand q runs from filing $first(q)$ to $last(q)$. The endpoints to explain are

$$N = E \cup S, \quad E = \{ end(q) : last(q) < T \}, \quad S = \{ start(q) : first(q) > 1 \}.$$

4.4.2 Candidate edits

A candidate is one possible edit, written as the set of endpoints it would explain, with a price w_c and a check it must pass to exist at all (Table 1). A candidate is all or nothing: it covers every endpoint in its set or none. It is a set rather than a pair because a merge, a split, a carve-out or an absorb pairs one line with several; matched one to one, the extra lines come out as ended or started and the totals they take part in go unexplained. Each kind has its own check, rather than one shared measure of what fails to add up, because an explanation that simply ends lines is never tested against any arithmetic and under a shared measure would win by default. Every candidate but a rename sits at one boundary t : its ends are strands whose last filing is F_{t-1} and its starts are strands whose first filing is F_t . Each kind is drawn beside its description: a filled dot is where a line's run ends in filing t , an open dot where a run starts in filing $t+1$, and a diamond is one edit, covering every endpoint it touches (Figure 8 in Appendix B shows them all together). Every check below is a sum or an equality on periods both reports show, and a sum is evidence only through the terms that could have made it fail: a term inside the sum's tolerance on the checked period adds nothing, so a part counts only where it carries a material number there. Parts $u_1, \dots, u_m$ *add up* to v when

$$\left| \sum_j u_j - v \right| \leq \max\left(\frac{m}{2} 10^{-d}, \varepsilon\right),$$

with d the coarsest precision among the parts and the total, so that each rounded part may be off by half a unit; every check below that asks for a sum uses this. Each kind is described here by its shape and its check; Appendix A.3 states every rule and bound in full.

Rename. One end and one start, $\{e, s\}$. Any end may pair with any later start, at the next boundary or across a gap, and a tag re-tied to itself after a restatement is a rename too. A pair is refused when the units, the period types (a balance never continues as a flow) or the per-share status conflict, when two different tags carry opposite words (gross and net, current and noncurrent), or, across a gap, when either tag or label appears in a filing the rename skips.

1 end, 1 start (also a line re-tied across a restatement)



Merge and split. Several ends and one start, or one end and several starts, at one boundary: two or three material parts that add up to a material whole on every period all of them report, at least one of those periods material. No part may be left out with the rest still adding up, and parts the filer itself declares under the whole are a break out or a fold, not a merge or split. *Equal copies* have the same shape, with every part equal to the whole. The search is bounded (Appendix A.3).

Carve-out and absorb. A rename $e \rightarrow s$ whose numbers disagree on a shared period, plus one to three lines that start (carve-out) or end (absorb) whose weighted numbers explain the restatement exactly on every period each child reports:

$$\text{carve-out: } v_{\text{old}}(k) = v_{\text{new}}(k) + \sum_j r_j u_j(k),$$


$$\text{absorb: } v_{\text{new}}(k) = v_{\text{old}}(k) + \sum_j r_j u_j(k),$$

where r_j is the ratio of the child's declared weight to the parent's when both add into one total, and otherwise -1 when their balance attributes (debit or credit) are opposite and 1 when not. A declared total above the parent or a child whose restatement equals the movement the edit implies into it is explained by the edit and joins the candidate; a declared total is never itself the parent. One line can also be absorbed into two.


Split with transfer. A split whose parts do not add up to the old line, where the difference, period by period, is exactly the restatement of the section total they add into. The candidate adds that total's rename, $\{e_P\} \cup \{s_j\} \cup \{e_T, s_T\}$, with the total's restatement credited: the same rule a carve-out or absorb applies to every total it moves lines into or out of.

Swap. Two or more lines whose numbers moved onto each other's tags, as a cycle, $\{e_1, s_1, \dots, e_m, s_m\}$: every leg's numbers continue verbatim on the other tag, and every leg's old place is accounted for.


k ends, 1 start




1 end, k starts (equal copies alike)




the line's end and start, the new lines' starts




the line's end and start, the absorbed lines' ends




a split, plus the section total's end and start



1 end and 1 start per line, crossed



Fold and break out. The parts' ends (a fold) or starts (a break out) under a total that runs on across the boundary: at least two parts, each material on a shared period, all declared under the total and adding up to it, with the total's numbers unchanged. If the total's own tag changes at the same boundary the candidate holds that rename too, as explanation C of Tesla does.



Ended and started are not candidates. They are what an endpoint costs when nothing covers it.

4.4.3 The optimization

Given the endpoints N and the candidates C , each $c \in C$ a set of endpoints with a price w_c , choose candidates that use no endpoint twice, at the lowest total price, where an endpoint nothing covers pays its own price p_n for having ended or started:

$$\min_{x \in \{0,1\}^C} \underbrace{\sum_{c \in C} w_c x_c}_{\text{price of the chosen edits}} + \underbrace{\sum_{n \in N} p_n \left(1 - \sum_{c \ni n} x_c\right)}_{\text{price of the strand ends left uncovered}}$$

subject to

$$\underbrace{\sum_{c \ni n} x_c}_{\substack{0: n \text{ uncovered} \\ 1: n \text{ covered}}} \leq 1 \quad \text{for every } n \in N.$$

$x_c = 1$ chooses candidate c . The constraint says no endpoint is covered twice, so the inner sum is 1 if n is covered and 0 if not.

Regrouping the objective by candidate instead of by endpoint (Appendix A.2) shows the same problem as weighted set packing: each candidate saves $s_c = \sum_{n \in c} p_n - w_c$, the price of leaving its endpoints unexplained minus its own price, and the choice is the disjoint set of candidates with the largest total saving. Candidates that save nothing are dropped before solving, since removing one from any answer never makes it worse. The guarantee is narrow: the chosen set is the cheapest among the candidates Section 4.4.2 generates, whose listing is bounded and pruned, under the prices

of Section 4.4.4. An explanation the listing did not produce cannot be chosen, and a different set of prices can choose differently.

The problem splits into independent groups. Two candidates interact only through a shared endpoint, so the connected components of the candidates, two adjacent when they share an endpoint, are independent problems, and each is solved on its own, exactly, as a 0–1 integer program with one constraint per shared endpoint; a group with one candidate takes it without solving anything, and choosing nothing is always allowed, so every group has an answer. The groups are small. Over the 200 companies of Section 5, the 596 statement histories with a boundary, the choice splits into 27,969 groups; 86% hold one candidate and 95% five or fewer, the largest holds 3,351 candidates over 158 endpoints and solves in 41 ms, and all the solves together take a quarter of a second. The answer need not be unique: two sets of candidates can save exactly the same, and the implementation then picks one by a fixed order, which only makes the choice repeatable.

Appendix C.4 lays out Tesla’s 2015 to 2016 boundary as this problem, with its candidates as sets and what each explanation leaves uncovered.

4.4.4 The prices

A price is roughly how much one would have to write down to describe the edit beyond what the filings already say. There are many ways to write such a price, and choosing one is a matter of judgment. The prices decide the answer, since candidates compete for the same endpoints: raising one kind’s price makes it lose more often, lowering it makes it win. Both the form of the prices and their constants took trial and error against cases we knew the answer to, and many others would work. What follows is the structure we settled on.

Prices are counted in digits. $\sigma(v)$ is the number of significant figures of v at its stated precision, trailing zeros removed (1,105 million has four; 880 million has two). Writing a number down costs $c(v) = \alpha \sigma(v) + \beta$, and every price below is counted in that unit, so different kinds of edit compete on equal terms. A number that carries over is free and a changed or new number costs its digits, so an edit cannot win by hiding numbers; evidence that singles out one pairing earns a credit and evidence the filer guaranteed earns none; one event is charged once, however many numbers it touched; and no kind is rewarded for size (the six properties are listed in Appendix A.4).

Every price has the same shape:

$$w_c = \underbrace{A_c}_{\text{the act: how rare this kind should be}} + \underbrace{D_c}_{\text{what differs: names and numbers}} - \underbrace{K_c}_{\text{what matches: the evidence}} .$$

The act is a flat price per kind of edit: a rename has none, a merge, a carve-out or a swap has one, and raising it means that kind needs more evidence before it wins. What differs is charged at its digits, every word of a label and part of a tag that changed and every number that is new or changed. What matches is credited, the same tag or label, a number that agrees, a sum that checks, and a match many lines share is worth little, one that singles out this pairing a lot, and one the filer guaranteed nothing. Charging only for what differs, with no credit, would leave the price saying nothing about evidence: a tag whose numbers moved verbatim to another tag (Palantir, Appendix C.3) would price the same as that tag continuing with a wholesale restatement, and the tie would go to keeping the tag. The credit for numbers that agree is what lets the numbers override the tag. Table 2 says what the three terms are for each kind of edit; Appendix A.4 writes them out in full, and Appendix A.5 gives the constants, which set how much each term weighs. An uncovered endpoint is priced the same way with no act to buy: a start pays for every number it shows, an end only for the numbers that turn up on some other line in the next filing.

KIND	ACT A_c	WHAT DIFFERS D_c	WHAT MATCHES K_c
rename	none	label words and tag parts that changed; numbers on new periods; numbers that changed (a rescale counts once)	same tag, label or note; each number that agrees
merge, split	μ	the whole's name against the parts'; each extra part pays as ended or started; numbers the check never examined	each period where the sum checks; nothing if the filer declared the sum
equal copies	μ , $+\pi$ per extra copy	as merge, split	none
carve-out, absorb	η	the parent's rename, less the restatement the children explain; children pay as started or ended	the parent's rename
split with transfer	μ	a split, plus the total's rename less the restatement the transfer explains	the total's rename
swap	θ , $+\theta_1$ per leg	each leg's numbers as a rename; no name charge, it was one accident	each leg's rename
fold, break out	one line's base back per extra part	parts pay as ended or started; a break out's parts less the one figure per period the total implies, to the digits the check verifies	none: the sum is the filer's declared arc

Table 2: The three terms for each kind of edit. “The parent’s rename” means that line’s rename price, computed the same way.

4.5 The rows

The solution is the cheapest set of edits between each pair of successive filings. The renames among them, including those inside carve-outs and absorbs, join strands end to start, and the joined strands are the rows of the statement. A merge or a split supplies a row with a derived number where no filing states one, the sum of the parts, marked as derived and cited to the lines it adds up (Appendix C.1 and C.5). That number rests on the merge’s check having held on the periods both reports show, two years for an income statement or cash flow statement and one instant for a balance sheet; beyond them it is the check extended, which is why it is marked.

A row spans several filings, which may give different numbers for the same year. So each year of the table is assigned a *ruling* filing, and its column contains only that filing’s numbers. Each row’s number is looked up by the tag the row had in that filing. Every number in a column comes from one filing, so the column should add up as that filing did; Section 5.1 checks that it does. The table as first reported and the table as last reported differ only in which filing rules each year: the filing for that year, or the latest filing that still shows it (Figure 7). If the ruling filing has no such line, the cell is blank. A number is never taken from a later filing.

A row is a lineage, and its numbers need not be comparable throughout. The output tells three things apart. The *lineage* is the strands the chosen edits joined: one line followed through the filings, whatever it was called. *Comparable* is narrower. A carve-out or an absorb changes what a continuing line contains, and the filing that made the change restates only the years it shows again, so a row is like for like within each ruling filing’s window and may change definition at a boundary where such an edit sits. Apple’s “Other” is one lineage with two definitions: 2019, ruled by the fiscal 2021 report, excludes the proceeds from common stock, and 2020 onwards includes them. The table marks every cell a restatement changed, and the edits chosen at each boundary (Section 5.1) say which kind of edit sits at each join of a row, so the boundary where a definition changed can be read off; the table does not restate the earlier years itself, since no filing states those numbers.

Reconstructed is a number made from others: the derived cell of a merge or split, the sum of its parts, marked and cited to the parts.

Automotive

us-gaap:SalesRevenueGoodsNet (the 2015 report's line on this tag is the total's)

AS FIRST REPORTED

the report whose year it is (earliest shown, for years before the first report)

	2013	2014	2015	2016	2017
2015 report	no such line	no such line	no such line		
2016 report		2,874.4	3,431.6	5,589.0	
2017 report			3,431.6	5,589.0	8,534.8
the table				5,589.0	8,534.8

AS LAST REPORTED

the latest report that shows the year

	2013	2014	2015	2016	2017
2015 report	no such line	no such line	no such line		
2016 report		2,874.4	3,431.6	5,589.0	
2017 report			3,431.6	5,589.0	8,534.8
the table		2,874.4	3,431.6	5,589.0	8,534.8

Automotive leasing

us-gaap:OperatingLeasesIncomeStatementLeaseRevenue

AS FIRST REPORTED

the report whose year it is (earliest shown, for years before the first report)

	2013	2014	2015	2016	2017
2015 report	no such line	no such line	no such line		
2016 report		132.6	309.4	761.8	
2017 report			309.4	761.8	1,106.5
the table				761.8	1,106.5

AS LAST REPORTED

the latest report that shows the year

	2013	2014	2015	2016	2017
2015 report	no such line	no such line	no such line		
2016 report		132.6	309.4	761.8	
2017 report			309.4	761.8	1,106.5
the table		132.6	309.4	761.8	1,106.5

Figure 7: Which filing rules each year's cell, for two of Tesla's rows. Each row of a stack is one annual filing and each column a year; a filing shows its own year and the two before. The boxed cell is the ruling filing's, and the table row beneath is what it gives. As first reported takes the filing whose year it is (for 2013 and 2014, before the first filing shown, the earliest that shows them); as last reported takes the latest filing that shows the year. Where the ruling filing has no such line, the cell is blank rather than filled from a later filing.

Appendix C shows all of this on real filings: the candidates, their prices, the choice and the rows for each. Tesla's rows, as first and as last reported, are in Appendix C.4.

5 Results and applications

5.1 Results

The two sets. We use two sets of companies. The *specimen set*: thirty-two companies, three statements each (income statement, balance sheet and cash flow statement), up to fifteen annual reports apiece, 96 statement histories and 800 boundaries between one annual report and the next. The method was developed against them; the specimens of Appendix C are among them, and the constants of Section 4.4.4 were settled by looking at how they came out. The *random sample*: two hundred companies drawn at random from every company our data pipeline has processed with at least eight annual reports: 597 statement histories, one of them a single report with no boundary, and 6,103 boundaries. The sample was held out until the method had been built and first evaluated; after that, its cases informed three revisions to the method, listed in Appendix A.6. The results below are the final method on both sets. The draw includes companies that have since been acquired, gone private or deregistered; a sample of survivors would have been easier to compare with a vendor's series, but it would not represent the filings. Every number below comes from one run of the implementation, with nothing set by hand for any company.

Changes of shape. Table 3 counts, by statement, the boundaries at which the optimization chose at least one change of shape: a merge, split, carve-out, absorb, fold, break-out, swap or transfer. These are chosen edits, not candidates. A change of shape is one the tag cannot follow, and 30 of the 32 specimen companies and 194 of the 200 in the random sample have one somewhere in their history; the cash flow statement carries them most.

	SPECIMEN SET				RANDOM SAMPLE			
	INCOME	BALANCE	CASH FLOW	ALL	INCOME	BALANCE	CASH FLOW	ALL
boundaries	266	267	267	800	1,988	2,066	2,049	6,103
with a change of shape	27	34	69	130	273	298	448	1,019
share	10%	13%	26%	16%	14%	14%	22%	17%

Table 3: Boundaries with a change of shape, by statement. A boundary is one pair of consecutive annual reports of one statement; it counts when the optimization chose at least one edit other than a continuation there.

The edits chosen. Table 4 counts every edit the method chose, from the same runs. Across a break the optimization makes 3,431 continuations on the specimen set and 39,882 on the random sample, a third and a quarter of them changing the tag and the rest keeping it across a restatement or a gap in the numbers; 3,048 and 16,638 numbers are restated. Continuations and restatements dominate; every other kind is rare, and each occurs on both sets. Written out for one company, Apple’s cash flow statement, the boundaries read: *2011 annual report*: renamed the tax paid on net share settlement to a company tag; absorbed *gain or loss on property, plant and equipment* and *other current assets* into *other operating assets*; restated 2 numbers. *2022 annual report*: absorbed *proceeds from issuance of common stock* (880 for 2020) into *other financing* (754); absorbed *purchases and proceeds from sale and maturity of other investments* into *other investing*. The 2022 boundary answers the question Section 1 opened with. Appendix D gives one case per kind, chosen where the names give no help.

Two more cases. The changes of shape are few, so the best way to see what the optimization finds is to follow some of them. Apple’s absorbs (Section 1) and Tesla’s boundary (Section 4.2) have been followed already; Appendix C gives all four in full, as facts, strands, candidates, choice and rows. Palantir filed three years of cash flow numbers on the wrong tags in its fiscal 2021 report and moved them back in 2023: a swap of three legs beats the four same-tag continuations, and the rows are the true series, each named by the tag it ends on (Appendix C.3). JPMorgan’s 2023 report replaces one line of non-interest revenue with two and restates both comparative years into the parts: a split beats a rename with a carve-out, and the old line continues as the sum of its parts, a derived row cited to both (Appendix C.5).

Arithmetic checks. Two checks, both against what the filings declare. First, the rule of Section 3.1 that every line is accounted for exactly once: every line of every filing ends up in exactly one row, note or flag, and nothing is used twice; 0 violations in 96 statements and in the 596 with a boundary. Second, the declared totals: the finished table must still add up, year by year, once the lines are joined across filings. A check is one declared total in one year; it passes, passes within the filer’s rounding, fails, or cannot be run because a line it needs has no number in that column. Table 5 counts them.

	KIND OF EDIT	SPECIMEN	%	RANDOM SAMPLE	%
continuations	tag kept across a restatement or a gap	2,262	25.7	29,747	34.9
	tag changed (rename)	1,169	13.3	10,135	11.9
changes of shape	break out	20	0.2	251	0.3
	fold	35	0.4	299	0.4
	carve-out	36	0.4	273	0.3
	absorb	60	0.7	227	0.3
	merge	28	0.3	139	0.2
	equal copies (split, merge)	11	0.1	144	0.2
	split	9	0.1	102	0.1
	transfer	1	<0.1	25	<0.1
	swap	2	<0.1	4	<0.1
	sever (a strand undone)	15	0.2	338	0.4
ends	ended	1,185	13.5	13,196	15.5
	started	907	10.3	13,765	16.1
amounts	restatement	3,048	34.7	16,638	19.5
all edits		8,788	100	85,283	100

Table 4: Every edit the method chose, all three statements: the specimen set (32 companies, 800 boundaries) and the random sample (200 companies, 6,103 boundaries), each kind with its share of all the edits chosen on that set. A continuation is one line carried across a boundary where the strands did not already carry it.

	ROWS	CHECKS: ONE DECLARED TOTAL IN ONE YEAR				
		PASS	ROUNDING	FAIL	NOT RUN	PASS OR ROUNDING, % OF RUN
specimen set, 96 statements	3,915	5,313	162	15	2,056	99.7
random sample, 597 statements	31,232	47,982	63	67	19,668	99.9

Table 5: Declared totals re-checked in the finished tables. The method also marks 18 and 115 checks not comparable, where the total and a part are stated by different filings.

Of the checks run, 99.7% pass or pass within rounding on the specimen set and 99.9% on the random sample, and no statement type, taken over all companies, is below 99.5%; single company statements go lower, to 93.7% for GE's balance sheet. The specimen set's 15 failures are all in GE and Berkshire, whose statements carry more than one entity in their columns (GE Industrial and GE Capital; Berkshire's segments): the model treats a breakdown row as a line of its own, and a total whose parts sit in another column does not add up on our page. The random sample's 67 are in 33 companies, 33 of them within a tenth of a percent of the total. The checks that cannot be run are mostly years for which no filing declares arithmetic for the line, and otherwise years where a part has no number in that column (Appendix E.3); those years are left unchecked rather than checked against another filing's equations.

Layout invariance. Each history can be shown in any filing’s layout. Rendered in the newest filing’s layout and in the oldest presenter’s, all 96 specimen statements are identical in every cell and every check over a five-year window, and 94 of 96 over fifteen; on the random sample, 564 of 597 over fifteen. Of the 35 that differ, 21 differ only in which checks run; in the other 14, every differing row was read, and the same numbers sit under a differently named row, or a merged line that one layout derives as the sum of its parts appears in the other as the parts alone (Appendix E.3). No cell holds two different numbers for the same line. An earlier run of this check found three wrong renames of a balance to a flow, and the rule of Section 4.4.2 that refuses the pairing came from them (Appendix A.6). The rows come from the strands, so the layout decides how the numbers are ordered, labelled and grouped, and which derived rows exist, not what the numbers are.

Agreement with a vendor. A commercial provider (FMP) publishes standardized series keyed by stock ticker, for companies listed today. Our sets span fifteen years, and 4 of the specimen companies and 75 of the random sample have since deregistered or been delisted, or file without a listing; one more in each set has no series our anchoring could use. That leaves 27 and 124. We anchor each of its fields to one of our lines in the newest year both cover, walk our history backwards under the method and under a floor, the tag alone, the anchored tag followed through every filing whatever its numbers do, and ask, year by year, whether the line we followed carries the number the vendor shows. It is a test of the lineage, not of the rendered cell. Any filing’s version of the number counts, since the vendor holds as-filed numbers and we show the ruling filing’s, and two numbers agree within one part in ten thousand or half a million dollars, whichever is larger (one cent for a per-share figure). A derived cell, the sum a merge or split supplies, is in no filing and is not confirmed by this test. A vendor value of zero is skipped, and a vendor number no filing states is the vendor’s own arithmetic and is left out (1,999 and 7,947 of them), a test made against the vendor’s as-reported feed [8]. Table 6 is the result, with the tag alone as the floor.

The tag alone is a strong floor. On the cells it reaches it agrees with the vendor 98.2% and 96.5% of the time, and the method keeps that (98.4 and 96.6% on the same cells): where a tag runs on with its numbers unchanged the strand rule joins it, so there the two are the same chain. What the optimization adds is the cells the tag cannot reach: 543 on the specimen set, 5.6% more than the tag alone fills, and 5,004 on the random sample, 9.7% more (41 go the other way, so 4,963 net), which takes agreement over all cells from 91.7 to 96.8% and from 83.6 to 91.2%. Every one of those cells lies past a boundary where the tag changed or the line was reshaped, and the edits of Table 4 are what reach them: 3,431 continuations and 202 changes of shape on the specimen set, 39,882 and 1,464 on the random sample. On the added cells the method agrees with the vendor 96.0% and 90.4% of the time, where the tag alone shows nothing; these are the hard cells, and the rate over each method’s own cells (98.2% for both, 96.1 against 96.5%) is lower for the method only because it includes them. The cash flow statement is the hard one for both, on both sets; it is where filers reshape most: 26% of its boundaries on the specimen set and 22% on the random sample carry a change of shape, against 10 to 14% for the other two statements (Table 3).

The vendor’s series are not ground truth, and the vendor has a different goal. It standardizes so that companies can be compared with one another [9], the purpose the SEC gave for standard tags [29]; we standardize one company across time and keep its own lines. Most of the misses are the vendor’s own arithmetic: it nets some lines, gathers others into residual buckets, adjusts per-share figures for later splits, and applies its own definitions of net income and of payables, each of which puts a number in its series that no filing states for the line we anchored. Appendix E.2 sorts every miss; the identity errors the comparison finds or cannot rule out come to 61 cells on the specimen set, 0.6%, and 1,844 on the random sample, 3.1%, most of them one line that the filer kept under

WAY OF JOINING	AGREES	DISAGREES	NO NUMBER	AGREES, %		AGREES, % BY STATEMENT		
				OF ALL	OF SHOWN	INCOME	BALANCE	CASH FLOW
specimen set: 27 companies, 10,404 cells								
the tag alone	9,538	171	695	91.7	98.2	91.7	94.8	88.2
strands and the optimization	10,070	182	152	96.8	98.2	97.2	98.6	94.5
the 543 cells the tag alone cannot fill								
strands and the optimization	521	22			96.0	96.1	96.2	95.7
the 9,709 cells both fill								
the tag alone	9,538	171			98.2	97.5	99.5	97.3
strands and the optimization	9,549	160			98.4	97.6	99.5	97.6
random sample: 124 companies, 59,371 cells								
the tag alone	49,609	1,791	7,971	83.6	96.5	80.3	89.5	78.8
strands and the optimization	54,137	2,226	3,008	91.2	96.1	90.0	94.2	88.3
the 5,004 cells the tag alone cannot fill								
strands and the optimization	4,523	481			90.4	85.3	92.8	93.9
the 51,359 cells both fill								
the tag alone	49,580	1,779			96.5	95.1	97.8	96.0
strands and the optimization	49,614	1,745			96.6	95.2	97.8	96.1

Table 6: Agreement with a data vendor’s series along the reconstructed histories. A cell agrees if the line we followed carries, in some filing, the number the vendor shows for that year; “no number” means our history has no value there. “Of shown” counts only cells where we show a number. The rows for the cells the tag alone cannot fill score the optimization where it adds coverage; the rows for the cells both fill score the two ways of joining on the same cells.

two tags for a while and the optimization left in two rows. These are agreement rates, not accuracy: two series can agree on a wrong number, and a cell that agrees by coincidence is not found by examining disagreements. What the comparison does establish is the coverage gain, and that every cell where the vendor’s standardization departs from the filings counts against us. The method and the tag alone are scored against the same series with the same anchors, so the rows of Table 6 can be compared; the absolute level is set by the vendor’s standardization, not by us.

5.2 Applications

The output is the strands with the edits between them, and several products follow. Each needs both the numbers and the links, with the filing behind each: the filings have every number but no link between one year’s filing and the next (Section 2.2), and a vendor has links, made its own way, but keeps one version of each number, standardized onto its own fields, without citation. What the output is not: it is not standardized across companies, since a row is one company’s line, not “revenue” in general, and putting two companies side by side is the vendor’s job; it does not say whether a change mattered to the business; and its checks reach only as far as the arithmetic the filers declare.

Cited tables. A company’s statements over fifteen years in any filing’s layout, where every cell names the filing and the fact it came from, every break in a line is marked, and a rebuilt total is a

checked formula over the lines it adds up (Section 4.5). What this adds over a vendor’s series is the citation on each number and a mark where a line ends.

Two axes of time. Every number a company reports is reported again, for two or three years, by later filings, and sometimes changed. The strands keep every version, each dated by its filing: a lattice with two axes of time, the period a number is for and the date it was said, of which any table is a slice: the table as it stood on any past date, the history as first reported against as last reported (Section 4.5), or the restatement record of one line. Economists keep such vintages for macroeconomic series (Section 1.2); nothing published keeps them for the lines of a company’s statements.

The record of changes. The method’s by-product is a record of what each company changed in how it reports, boundary by boundary and by kind: on the specimen set, 130 of 800 boundaries carry a change of shape, and Table 4 says which. Companies work this mapping out every time they restate, and the filings carry no consistent, machine-readable form of it (Section 1.2); this is the nearest thing, built from the outside. It says which lines are stable, when a company reorganized its reporting, and how one company’s habits compare with another’s. It does not say why; the reason is in the filing’s own words, and the record says which boundary and which lines to read it for.

An interface for programs. A program asking about a company’s history needs a stable identity for a line across years, a source for every number, and a way to check a table it did not build. The strands are the identity, the cells carry the source, and the declared totals are the check, so the natural interface is a query over strands, cells and edits, with the tables one view of it: *revenue for the last ten years as first reported, with the filing for each; does this table an assistant produced agree with the filings, cell by cell.*

6 Extensions

The method needs no mapping written for a company and no dictionary of related tags. It does not know that `OtherOperatingActivities` is an operating item or that commissions are a kind of non-interest revenue. The evidence it works from is in the filings: the numbers on each line, the prior years each filing shows again, the units, whether a number is stated at a date or over a period, the arithmetic each filing declares, and the debit or credit mark on each line. Names are compared as strings, by the words two labels share and the characters two tags share, with one exception: a short list of word pairs (gross and net, current and noncurrent) that a rename may not cross. A tag the company invented is handled the same way as a standard one. The method should carry over to any series of reports in which each report overlaps the one before.

The method’s input is XBRL facts: a tag, a period, a number and the arithmetic the filing declares. Filings that carry none, annual reports before 2009, most forms other than the 10-K and 10-Q, and reports filed outside the SEC, could be brought in by extracting their numbers and assigning each a tag first. That is the task recent work has set for language models: FinTagging benchmarks the extraction of numbers from the text and tables of a report and their linking to the US GAAP taxonomy [32], and XBRLTagRec assigns tags to numbers by retrieval and re-ranking [11]. Facts produced that way would feed the method unchanged, with two losses: a tag assigned by a model is weaker evidence of identity than one the filer chose, so the prices of Section 4.4.4 would need to weigh it less, and a report with no declared arithmetic gives the checks nothing to foot against, so folds, break-outs and the re-checked totals would need the arithmetic inferred as well.

The record of changes (Section 5.2) also supports questions this paper does not ask. Do reorganizations cluster around new versions of the tag dictionary? Are companies reporting less detail over time? Which lines are stable across a whole industry, and which never are?

7 Conclusion

A financial statement that covers many years has to be built from many filings, and the filings report the same measures differently as the company changes what it shows and how it labels it. Most lines keep their label and tag from one filing to the next and can be followed by name. But neither is fixed: both change while the line continues, and both are reused for lines that are new. A name is only evidence. The filings carry no consistent, machine-readable record of the changes, at most a sentence in a note, so the changes have to be inferred from the filings, and the only certain evidence is the arithmetic each filing declares and the prior years each shows again under its new lines.

We set that inference up as an optimization problem. At each boundary between consecutive filings, the candidate edits of each known kind that pass its check are listed, each with the lines it would explain and a price for how much it asks the reader to believe beyond what the filings show. An integer program picks the cheapest set of the listed edits that accounts for every line exactly once, and the chosen edits trace each line through the filings. The result is a statement whose rows run across every year the filings cover, in any filing's layout, with each number cited to the filing that stated it and each declared total re-checked.

We ran it on 232 companies. Where the line the method follows and a commercial vendor's series both give a number, they agree 96 to 98% of the time. The table reaches 6 to 10% more cells than the tag alone, every one past a change of tag or of shape, and agrees with the vendor on 96% and 90% of those; 99.9% of the declared totals that can be checked add up. A company works out this mapping every time it restates, and its filings carry no machine-readable form of it. The method recovers it from what the filings do state: their arithmetic and the years they show again.

Appendices

A The method in full

A.1 The procedure as pseudocode

Algorithm 1 states the four steps of Section 4.3 as one procedure, from the filings to the edits and the rows.

Algorithm 1 Rebuilding a statement's history from its filings.

Require: filings $F_1, \dots, F_T$ of one company and statement, in the order published

Ensure: the edits between each pair of filings, and the rows that follow from them

- 1: $S \leftarrow$ the strands of $F_1, \dots, F_T$: runs of one tag joined on the same numbers in consecutive filings
 - 2: $N \leftarrow$ every point where a strand ends before F_T or starts after F_1
 - 3: $C \leftarrow \emptyset$
 - 4: **for** each filing boundary t **do**
 - 5: **for** each kind of edit k : rename, merge, split, carve-out, absorb, fold, $\dots$ **do**
 - 6: $C \leftarrow C \cup \text{CANDIDATES}(k, S, t)$, each with its endpoints, its check and its price
 - 7: $C \leftarrow \{c \in C : c \text{ is priced below leaving its endpoints unexplained}\}$
 - 8: $\mathcal{G} \leftarrow$ the connected components of C , two candidates adjacent when they share an endpoint
 - 9: $X \leftarrow \emptyset$
 - 10: **for** each component $G \in \mathcal{G}$ **do**
 - 11: $X \leftarrow X \cup \text{CHOOSE}(G)$, the cheapest set in G using no endpoint twice, solved exactly
 - 12: every endpoint not covered by X has ended or started
 - 13: rows $\leftarrow$ strands linked end to start by the renames in X
 - 14: **return** X , the rows, and every candidate that lost, with its price
-

A.2 The savings form

The optimization of Section 4.4.3 minimizes the price of the chosen edits plus the price of the endpoints left uncovered,

$$\sum_{c \in C} w_c x_c + \sum_{n \in N} p_n \left(1 - \sum_{c \ni n} x_c\right),$$

subject to no endpoint being covered twice. Multiplying out the second term gives $\sum_n p_n - \sum_n p_n \sum_{c \ni n} x_c$. The double sum has one term $p_n x_c$ for every endpoint n and every candidate c that holds it, so it can be added up by candidate instead:

$$\sum_{n \in N} p_n \sum_{c \ni n} x_c = \sum_{\substack{(n, c) \\ n \in c}} p_n x_c = \sum_{c \in C} x_c \sum_{n \in c} p_n.$$

The objective is then

$$\sum_{n \in N} p_n + \sum_{c \in C} x_c \left(w_c - \sum_{n \in c} p_n\right) = \sum_{n \in N} p_n - \sum_{c \in C} s_c x_c, \quad s_c = \sum_{n \in c} p_n - w_c,$$

where s_c is what candidate c saves. The first term does not depend on x , so minimizing the price is maximizing the total saving of a disjoint set of candidates: weighted set packing.

A.3 The candidates in full

The rules of Section 4.4.2, in full: for each kind of edit, the shape of the candidate, the check it must pass to exist, and the bounds on its search.

Rename. One end and one start: $\{e, s\}$.

- Any end may pair with any later start, at the next boundary or across a gap. A tag re-tied to itself, after a restatement or after a break with nothing to compare, is a rename too.
- Refused when the unit classes conflict, when one line is per share and the other is not, when one line is stated at instants and the other over periods (a balance never continues as a flow: the two share no period, so no number could ever test the pairing), or when two different tags carry opposite words (gross and net, current and noncurrent; a tag carrying both words of a pair takes no side).
- A rename across a gap is refused if, in any filing it skips, either tag or either label appears: the line was not missing there, it was somewhere else.

Merge and split. Several ends and one start, or one end and several starts, at one boundary.

- Two or three parts, every one of them material, and the whole material.
- The parts add up to the whole on every period where the whole and every part report a number, and at least one of those periods is material (two, for three parts).
- No part can be left out with the rest still adding up, and the parts are not simply the lines the filer itself declares under the whole: that is a break out or a fold, not a merge or split.
- *Equal copies* have the same shape, with every part equal to the whole instead of adding up to it.
- The search is bounded: pairs and triples only where a boundary has few enough ends (or starts), and only the cheapest few merges and splits per boundary are kept (Appendix A.5).

Carve-out and absorb. A rename $e \rightarrow s$ across boundary t , plus new lines (carve-out) or ended lines (absorb): $\{e, s\} \cup$ the children's starts or ends.

- The rename's numbers must disagree on some shared period, and not all by one ratio (that is a rescale, priced in the rename).
- One, two or three children whose weighted numbers explain the restatement exactly. On every period each child reports,

$$\begin{aligned} \text{carve-out: } v_{\text{old}}(k) &= v_{\text{new}}(k) + \sum_j r_j u_j(k), \\ \text{absorb: } v_{\text{new}}(k) &= v_{\text{old}}(k) + \sum_j r_j u_j(k), \end{aligned}$$

as sums that agree, where $r_j = \omega_j / \omega_{\text{parent}}$ if the child and the parent add into the same total; otherwise $r_j = -1$ when the child's balance attribute (debit or credit) is the opposite of the parent's, and 1 when it is the same or undeclared. So interest expense (-1 into pre-tax income) absorbed by other income ($+1$) enters with -1 , and a liability carved out of a receivable it was netted inside enters with -1 .

- Children are tried one at a time first; a pair only from lines no single child explained, and a triple only from pairs that failed.
- A declared total above the parent or a child, at any height, whose own strand broke at the boundary and whose restatement equals the movement the edit implies into it along the declared arcs on the footing periods, is explained by the edit: its rename joins the candidate and that restatement is credited, as in a split with transfer. Every other continuing line of that total must be unmoved on those periods, and at least one must exist.
- A declared total is never the parent. Its restatement is explained by the movement of its parts, and a total carving a line out of itself would count that movement twice.
- *Absorbed into two*: one ended line c and two restated renames $e_1 \rightarrow s_1, e_2 \rightarrow s_2$ whose changes add up to c on every period it reported: $\{e_c, e_1, s_1, e_2, s_2\}$.

Split with transfer. A split whose parts do not add up to the old line, where the difference, period by period, is exactly the restatement of the section total they add into. The candidate adds that total's rename, $\{e_P\} \cup \{s_j\} \cup \{e_T, s_T\}$, with the total's restatement credited: the same rule a carve-out or absorb applies to every total it moves lines into or out of.

Swap. Two or more lines whose numbers moved onto each other's tags, as a cycle. The candidate is every leg's end and start, $\{e_1, s_1, \dots, e_m, s_m\}$.

- Each leg's numbers continue verbatim on the other tag.
- Each leg's old place is accounted for: its tag is another leg's destination, or it is absent from the new filing, or its run there breaks.

Fold and break out. The parts' ends (a fold) or starts (a break out), under a total that runs on across the boundary.

- At least two parts, each material on a period both reports show, all declared under the total, and the declared sum adds up. Two new lines that are zero in every shared year add up under an old total as well as a real break out would, and are lines that started.
- The total's numbers do not change; a total that restated is an absorb or a split with transfer.
- If the total runs on in one strand, the candidate is the parts alone. If the total's own tag changes at the same boundary, the candidate also holds that rename's end and start, as explanation C of Tesla does.
- A part whose tag has another life across the boundary is admitted only when the candidate's own rename carries that other life, as Tesla's reused automotive tag is. Otherwise a fold and a break out of the same tag would turn a restatement into two edits and cut the row.

Ended and started are not candidates. They are what an endpoint costs when nothing covers it.

A.4 The prices in full

Each price of Section 4.4.4, written out. Constants are in Appendix A.5. Every price respects six things:

1. A number that carries over is free; a changed or new number costs its digits, wherever it appears. An edit cannot win by hiding numbers.
2. Names: the closer two labels, and two tags, the cheaper the pairing.
3. Evidence that singles out one pairing among many (the same tag, the same label, a number that agrees, a sum that checks) earns a credit; evidence shared by many lines earns little; evidence that was guaranteed (a sum the filer declared) earns nothing.
4. One event is charged once: a swap, a carve-out or a change of scale is an act, not one charge per number it touched.
5. No reward for size: every extra line an edit covers pays its own way.
6. A line that ends explains itself, unless its numbers turn up elsewhere; a line that starts brings every number it shows.

Ended and started.

$$p_s = \kappa_S + \sum_k c(v(k)), \quad p_e = \kappa_E + \sum_{k \in \text{echo}(e)} c(v(k)),$$

where $\text{echo}(e)$ is the set of material numbers of the ended line that appear on some other line in the next filing but on no other line in its own.

Rename. The price of what differs, $D = D_{\text{name}} + D_{\text{num}}$, less a credit for what matches, K :

$$w_{e \rightarrow s} = D_{\text{name}}(o, n) + D_{\text{num}}(o, n) - K(o, n).$$

J is the share of label words two labels have in common; sim the share of characters two names have in common.

$$D_{\text{name}}(o, n) = \lambda (1 - J(\lambda_o, \lambda_n)) + \sum_{\text{tag parts that differ}} g(\text{part}),$$

$$\begin{aligned} g(\text{prefix only}) &= \gamma_0, & g(\text{concept}) &= \gamma_C + \delta_C (1 - \text{sim}), \\ g(\text{axis or member}) &= \gamma_A + \delta_A (1 - \text{sim}), & g(\text{breakdown added or removed}) &= \gamma_D; \end{aligned}$$

$$D_{\text{num}}(o, n) = \sum_{k \notin o} c(v_n(k)) + B,$$

$$B = \sum_{k \text{ disagrees}} c(v_n(k) - v_o(k)), \text{ capped at } \rho \text{ if one ratio moved every shared number;}$$

$$\begin{aligned} K(o, n) &= \frac{\kappa(M)}{m_\lambda m'_\lambda} [\text{same label}] + \frac{\kappa(M)}{m_\nu m'_\nu} [\text{same note}] + \kappa(M) [\text{same tag}] \\ &+ \sum_{k \text{ agrees}} \frac{\min(c(v(k)), \kappa(M))}{m_k m'_k}, \end{aligned}$$

where $\kappa(M) = \alpha \log_2(1 + M)$ and m, m' count the lines on the old and new side that share the evidence (a tag is unique on each side). A pair is refused outright, at any price, when the unit classes conflict, when one line is per share and the other not, when one line is stated at instants and the other over periods, or when the two tags carry opposite words.

Merge and split.

$$w = \max\left(\max(\mu + X + \delta_N(1 - \text{sim}), w_{\min}) + U + R - \Phi, w_{\min}\right).$$

- μ is the act; $w_{\min}$ keeps the price positive, before and after the credit.
- X : every part after the first pays its own ended or started price. Equal copies repeat the whole, so each extra copy pays a flat π instead.
- $\delta_N(1 - \text{sim})$: the best name similarity between the whole and any part, or the parts' names run together in any order.
- U : the new side's numbers on periods the check did not examine, at $c(v)$ each. A footing split's later parts already paid in full through X , so only its first part is charged here; an equal split's parts all are.
- Φ : for each period where the whole is material and the sum agrees within a relative tolerance φ , a credit of $\alpha \min(\bar{\sigma}, \sigma(v))$, where $\bar{\sigma} = \log_{10}(1/\varphi)$ is the number of figures such a check can verify. Equal copies earn no such credit.
- R : when the parts all add into one declared total and the whole carries that total's figures (the total under another tag), the sum agrees by construction, and the footed numbers are charged in full at $c(v)$ as if they were new. Parts that add into the whole itself are refused: that is a break out, not a split.

Carve-out and absorb .

$$w = w_{e \rightarrow s} - B^* + \eta + \sum_j p_j + \sum_T (w_{e_T \rightarrow s_T} - B_T^*).$$

- B^* is the part of the parent's restatement on the periods the children explain. It is replaced by one act η ; the children pay their own ended or started prices.
- T ranges over the declared totals above the parent and the children that the edit explains (Section 4.4.2): each contributes its rename and is credited its restatement B_T^* on the footing periods.
- The candidate exists only when $B^* + \sum_T B_T^* > \eta$: restatements cheaper than the act are left to the renames.
- Absorbed into two: $w = (w_1 - B_1^*) + (w_2 - B_2^*) + \eta + p_{ec}$.

Split with transfer. The split's act, extra parts and name term, its unchecked numbers, plus the section total's rename less the restatement the transfer explains:

$$w = \max(\mu + X + \delta_N(1 - \text{sim}), w_{\min}) + U + w_{e_T \rightarrow s_T} - B_T^*.$$

U , the new side's numbers on periods the check did not examine, is charged as for a split. There is no Φ , since the parts do not add up to the old line, and no R .

Swap

$$w = \theta + \theta_1 m + \sum_{i=1}^m (w_i - G_i - L_i)$$

for m legs, where G_i and L_i are leg i 's tag and label charges. The mistagging is one filing error, charged once as $\theta + \theta_1 m$; each leg still pays for its numbers, so a swap is chosen only when the numbers carry.

Fold and break out. The parts pay their own ended or started prices, less one base charge for each part after the first. A break out's parts are also not charged for the figure the total implies, one part per period, to the digits the footing check verifies ($\bar{\sigma}$, as for a split's credit):

$$w_{\text{break out}} = \sum_j p_{s_j} - (k-1) \kappa_S - \alpha \sum_k \min(\bar{\sigma}, \max_j \sigma(u_j(k))), \quad w_{\text{fold}} = \sum_j p_{e_j} - (k-1) \kappa_E,$$

plus the total's rename when its tag changes. The footing earns no evidence credit: it is the filer's declared arithmetic and adds up by construction.

A.5 Constants

The prices in Section 4.4.4 are formulas in the constants below. The formulas follow from the properties the prices must have; the values do not, and were tuned. Table 7 gives the values the implementation uses, set by hand against a bench of known cases and not fitted to the corpus. They are in units of $c(v)$, and only their ratios matter.

SYMBOL	WHAT IT PRICES	VALUE
α	one significant figure	1.2
β	stating a number at all	0.8
κ_S	a line that starts	4
κ_E	a line that ends	2
λ	a label with no word in common	5
γ_o	a new tag prefix only	1
γ_C, δ_C	a new concept: base, fully different name	3, 6
γ_A, δ_A	a new axis or member: base, fully different name	2, 4
γ_D	adding or removing a breakdown	3
ρ	a change of scale	2.5
μ	a merge or split	5
π	each extra equal copy	1
δ_N	a merge or split whose names are fully different	6
w_{min}	the lowest price of a merge or split	0.1
φ	relative tolerance of a sum credit ($\bar{\sigma} \approx 3$)	0.1%
η	a carve-out or absorb ($\mu + \pi$)	6
θ, θ_1	a swap: base, per leg	6, 1
<i>search bounds</i>		
	pairs of parts, when a boundary has at most this many ends (or starts)	60
	triples of parts, likewise	16
	merges and splits kept per boundary	8
	carve-out children: pairs from at most, triples from at most	40, 12

Table 7: One working set of constants for the prices in Section 4.4.4, and the bounds on the candidate search in Section 4.4.2.

A.6 Revisions the random sample prompted

The random sample of Section 5.1 was held out until the method had been built and first evaluated. Its cases then prompted three revisions, all in the rules of Section 4.4.2. DXP’s change from net to gross contract positions (Appendix D) prompted two: the sign a child takes from its balance attribute when it adds into a different total, and the credit a carve-out or absorb earns for the declared totals its movement explains. The layout check of Section 5.1 prompted the third: it found three renames of a balance to a flow, Iron Mountain’s dividends payable to a gain on disposal and, at two companies, the change in restricted cash to the restricted-cash balance in the cash bridge, and a line stated at instants now never continues as one stated over periods. The same case showed that the opposite-words rule had refused Iron Mountain’s true rename, from a tag carrying both “current” and “noncurrent” to one carrying “current”; a tag carrying both words of a pair now takes no side.

B The kinds of edit, illustrated

Every kind of edit in Table 1, drawn twice: a made-up case with round numbers, then a real one the method found, with the accession numbers of both filings and the numbers as filed. What the company printed is set in serif and the edit is drawn in blue (dashed where parts of a total are shown or hidden and no number moves); the shaded column is the year both filings report, where each edit is checked; a restated number is circled in orange. Tags without a prefix are **us-gaap**; some labels are shortened. Figure 8 first puts every kind side by side as the strand endpoints it covers.

Rename, ended, started. A rename says that a line that stops appearing under one tag carries on under another: a different tag, a different product or segment member, or the same tag after the company redid its prior years. Nothing has to add up; a rename is priced on what changed: each part of the tag, each word of the label, and each number that does not carry over. A line that ended has nothing after it, and one that started has nothing before it. Those two have a fixed price, and every other explanation of a line has to be cheaper to be chosen (Figure 9).

Merge and split. In a merge, several lines end and one new line starts whose prior years are their total; a split is the reverse (Figures 10 and 11). The sum is checked on every year both filings report, to the precision the filing states. Every part has to matter: a merge that would still add up without one of its lines is refused, which keeps out coincidences.




Figure 8: Each kind of edit as the strand endpoints it covers. A filled dot is where a line's run ends in filing t ; an open dot is where a run starts in filing $t+1$; the diamond is one edit, and it covers every endpoint it touches. A fold or break out leaves its total's run untouched (dashed). Anything no edit covers has ended or started.




Figure 9: rename, and lines that ended or started. Netflix’s deferred revenue moves to the tag of the new revenue standard, with the same wording and the same December 2017 balance.




Figure 10: merge. JPMorgan Chase shows commercial paper and other borrowed funds as one line from its 2017 10-K.




Figure 11: **split**. JPMorgan Chase separates asset-management fees from commissions in its 2023 10-K; the sum holds on both prior years the newer filing reports.

Equal copies. Some lines repeat rather than add up. When a company with nothing to dilute its shares reports one earnings-per-share line, “basic and diluted”, and later reports two, each new line equals the old one; treating it as a split would double it (Figure 12). The reverse is rarer, one case on the specimen set, at Kraft Heinz, and 28 on the random sample, so that direction is drawn with made-up numbers only.

Carve-out and absorb. “Reclassified to conform to the current period presentation” most often means one of these two. In a **carve-out**, a line continues, a new line starts, and the continuing line’s prior years are restated down by exactly the new line’s amount (Figure 13). An **absorb** is the reverse: lines end, and a continuing line’s prior years are restated up to include them (Figure 14). Signs follow the statement’s own arithmetic, so interest expense moved into “other income” counts negatively, and a line that leaves for the other side of the balance sheet, or for below the sections of a cash flow statement, enters with the sign its balance attribute gives. Several lines absorbed at once are one edit, since none explains the restatement on its own; one line can be absorbed into several, when their restatements add up to it. A total above the parent or a child that moved by exactly what the edit moved into it is explained by the edit, as Tidewater’s operating cash flow is in Figure 13.

Fold and break out. In a fold, the parts of a total stop being shown, but no number moves: the parts were always inside the total, and the total does not change (Figure 15). In a **break out**, a total that continues starts showing its parts (Figure 16). Neither restates anything, so each is checked by the total standing still and by the filing’s own arithmetic adding the parts up to it, in whichever filing shows the parts.

MADE-UP EXAMPLE

Filing t

Filing t+1

	Y0	Y1		Y1	Y2
EPS, basic and diluted x:EpsBasicAndDiluted	1.10	1.20	← Y1 in both → copied	1.20	1.35
				1.20	1.31

the check, on Y1

$$\frac{1.20}{\text{old}} = \frac{1.20}{\text{basic}} = \frac{1.20}{\text{diluted}}$$

TESLA, 2016 TO 2017

Tesla, Inc. · Form 10-K, fiscal 2016

0001564590-17-003118
Consolidated statements of operations · \$ per share

	2015	2016
Net loss per share, basic and diluted EarningsPerShareBasicAndDiluted	(6.93)	(4.68)

Tesla, Inc. · Form 10-K, fiscal 2017

0001564590-18-002966
Consolidated statements of operations · \$ per share

	2016	2017
Basic new line EarningsPerShareBasic	(4.68)	(11.83)
Diluted new line EarningsPerShareDiluted	(4.68)	(11.83)

the check, on 2016

$$\frac{(4.68)}{\text{old}} = \frac{(4.68)}{\text{basic}} = \frac{(4.68)}{\text{diluted}}$$

Figure 12: equal copies, splitting and merging. Tesla reports basic and diluted loss per share separately from its 2017 10-K.

MADE-UP EXAMPLE

Filing t

Filing t+1

	Y0	Y1		Y1	Y2
Product revenue x:ProductRevenue	100	120	← Y1 in both → continues, restated carved out	95	130
				25	35

the check, on Y1

$$\frac{120}{\text{old product revenue}} = \frac{95}{\text{new product revenue}} + \frac{25}{\text{leasing}}$$

TESLA, 2015 TO 2016

Tesla, Inc. · Form 10-K, fiscal 2015

0001564590-16-012195
Consolidated statements of operations · \$ millions

	2014	2015
Services and other tsla:SalesRevenueServicesAndOtherNet	191.3	305.1

Tesla, Inc. · Form 10-K, fiscal 2016

0001564590-17-003118
Consolidated statements of operations · \$ millions

	2015	2016
Services and other tsla:SalesRevenueServicesAndOtherNet	290.6	468.0
Energy generation and storage new line: SalesRevenueEnergyServices	14.5	181.4

the check, on 2015

$$\frac{305.1}{\text{old services and other}} = \frac{290.6}{\text{new services and other}} + \frac{14.5}{\text{energy}}$$

TIDEWATER, 2024 TO 2025: TWO LINES, ONE OUTSIDE THE SECTION

Tidewater Inc. · Form 10-K, fiscal 2024

0001437749-25-000487
Consolidated statements of cash flows · \$ thousands

	2023	2024
Other, net IncreaseDecreaseInOtherOperatingCapitalNet	(1,312)	353
Net cash provided by operating activities NetCashProvidedByUsedInOperatingActivities	104,705	273,849

Tidewater Inc. · Form 10-K, fiscal 2025

0001437749-26-000492
Consolidated statements of cash flows · \$ thousands

	2024	2025
Other, net IncreaseDecreaseInOtherOperatingCapitalNet	6,005	(15,278)
Unrealized foreign exchange loss (gain) new line ForeignCurrencyTransactionGainsAndLosses	(14,275)	23,398
Net cash provided by operating activities NetCashProvidedByUsedInOperatingActivities	282,472	379,108
Effect of exchange rates on cash (below the sections) new line EffectOfExchangeRatesOnCash	(8,623)	16,870

the check, on 2024

$$\frac{353}{\text{old other, net}} = \frac{6,005}{\text{new other, net}} + \frac{(14,275)}{\text{unrealized FX}} + \frac{8,623}{\text{-(8,623) exchange effect}}$$

Figure 13: carve-out. Tesla takes energy generation and storage out of services and other in its 2016 10-K and restates 2015 to match. Tidewater carves two lines out of “Other, net” in its 2025 10-K, one of them below the operating section; operating cash flow moves by exactly what left it.




Figure 14: absorb. Salesforce moves investment income and interest expense into other expense in its fiscal 2020 10-K: one edit, two lines absorbed. Tesla's resale value guarantee line is absorbed into two continuing lines at once.




Figure 15: fold. Netflix stops splitting cost of revenues into subscription and fulfillment costs in its 2012 10-K.




Figure 16: **break out.** Johnson & Johnson splits diluted earnings per share into continuing and discontinued operations after separating Kenvue.

Swap. A swap says the company put a run of filings' numbers on the wrong tags: the numbers on each tag belong to another tag's line (Figure 17; Appendix C.3 follows Palantir's case to its correction). It is considered only when the lines trade places completely, so that every line that lost its numbers got another line's, and chosen only when naming the mistake makes the whole history cheaper. A tag that kept its name but not its meaning is then two lines, one before the swap and one after.

Restatement and transfer. Some edits touch only numbers, on a line whose identity is settled (Figure 18). A **restatement** is a prior year's number that differs from the one first filed. A later filing can also fill in a year an earlier one left blank, or report a whole series at a different scale. A **transfer** goes with a split whose parts do not add up to the old line: the difference moved to another section, whose total was restated by exactly that amount (Figure 19).




Figure 17: **swap**. In the made-up example, lines X and Y keep their tags, but each tag now carries the other line's numbers. In Palantir's 2021 10-K each 2020 financing figure appears under the tag of the line above it.




Figure 18: Edits to numbers only, with made-up numbers.




Figure 19: split with transfer. PetMed Express's amended fiscal 2023 10-K shows sales tax payable as its own line. The parts exceed the old accrued-liabilities line, and total current liabilities for March 2023 is restated by exactly the difference.

C Examples

A made-up example taken through the four steps of Algorithm 1 (Section 4.3), then the four cases of Section 5.1 in full: the facts, the strands, the candidates, the solution, and the rows that follow. The four are from one run of the implementation, with nothing set by hand for any company; prices are in the units of Section 4.4.4, numbers in millions.

C.1 A made-up example

The filings. Three annual reports (Figure 20), each showing the current year and the year before. Marketing and Selling are replaced by Sales and marketing in the second filing, and Other income is retagged as Other income, net in the third.

Filing 1 · fiscal 2021			Filing 2 · fiscal 2022			Filing 3 · fiscal 2023		
	2020	2021		2021	2022		2022	2023
Revenue	100	120	Revenue	120	140	Revenue	140	150
x:Revenue			x:Revenue			x:Revenue		
Cost of sales	(60)	(70)	Cost of sales	(70)	(80)	Cost of sales	(80)	(85)
x:CostOfSales			x:CostOfSales			x:CostOfSales		
Marketing	(10)	(12)	Sales and marketing	(20)	(22)	Sales and marketing	(22)	(25)
x:Marketing			x:SalesMkt			x:SalesMkt		
Selling	(7)	(8)	Other income	4	6	Other income, net	6	3
x:Selling			x:OtherIncome			x:Nonop		
Other income	5	4						
x:OtherIncome								

Figure 20: The worked example: three made-up annual reports, each with its lines, their tags, and two years of numbers.

Step 1, strands. Revenue and cost of sales carry their numbers through all three filings, so each is one strand (Figure 21). Everything else leaves endpoints: Marketing and Selling end after the first filing (e_1, e_2) and Sales and marketing starts in the second (s_1); Other income ends after the second (e_3) and Other income, net starts in the third (s_2).




Figure 21: Step 1: strands. A bar is a strand, with the number that joined it across each boundary; a filled dot is where a strand ends (e) and an open dot where one starts (s).

Step 2, candidates. Four candidates cover those endpoints (Figure 22): c_1 , a merge of Marketing and Selling into Sales and marketing ($12 + 8 = 20$, which adds up); c_2 and c_3 , a rename of either

line into Sales and marketing, each needing a restatement ($12 \rightarrow 20$, $8 \rightarrow 20$); and c_4 , a rename of Other income ($6 = 6$, carried over).




Figure 22: Step 2: candidates. Each candidate joins the endpoints it would explain; a diamond marks an edit that covers more than two.

Step 3, choose the edits. c_1 , c_2 and c_3 all need s_1 , so at most one can be chosen. The merge adds up, pays nothing for numbers and leaves no endpoint uncovered; either rename would pay for a restatement and leave the other line to pay for having ended. c_4 shares no endpoint with anything and is decided on its own. The chosen set, c_1 and c_4 , covers every endpoint once (Figure 23).




Figure 23: Step 3: the choice. c_1 and c_4 are chosen (solid); c_2 and c_3 lose (dashed grey), because each would share s_1 with the merge.

Step 4, rows. Sales and marketing gets a 2020 figure of $10 + 7 = 17$ from the merge (Figure 24), and Other income is one row across its rename, under the tag it ends on.

	2020	2021	2022	2023	
Revenue	100	120	140	150	
Cost of sales	(60)	(70)	(80)	(85)	
Sales and marketing	(17)	(20)	(22)	(25)	2020 = 10 + 7, from the merge
Other income, net	5	4	6	3	one row across the rename, under its last tag

Figure 24: Step 4: the rows that follow from the choice. The figure in blue is not in any filing; the merge supplies it.

C.2 Apple

The facts. Table 8: three lines are gone from the fiscal 2022 report, and the two “Other” lines changed by exactly what they carried ($754 = 880 + (126)$; $(909) = (791) + (210) + 92$). The report does not say so [1, 2].

	FISCAL 2021 REPORT		FISCAL 2022 REPORT	
	2020	2021	2020	2021
Proceeds from issuance of common stock	880	1,105	—	—
Other (financing)	(126)	(129)	754	976
Purchases of non-marketable securities	(210)	(131)	—	—
Proceeds from non-marketable securities	92	387	—	—
Other (investing)	(791)	(608)	(909)	(352)

Table 8: Five lines of Apple’s cash flow statement for the same two years, as two consecutive reports gave them [1, 2]. Three lines are gone; the two “Other” lines changed by exactly their amounts.

Strands. Figure 25. The three gone lines are strands that end (e_1, e_3, e_4); each “Other” is a strand that ends and one that starts, since no shared number agrees ($e_2, s_1; e_5, s_2$); every other line runs through. Seven endpoints.

Candidates. Figure 26. Six:

- c_1, c_3 : each “Other” continues restated and **absorbs** the lines that vanished ($4.81 + 6.00$; $6.01 + 6.00$).
- c_2, c_4 : each “Other” continues restated; the gone lines ended.
- c_5, c_6 : each “Other” renamed to the other.

Solution. Figure 27: c_1 and c_3 . The absorb’s check explains the restatement, so the restated numbers are not charged, only the act: each absorb saves 15.99. The plain continuation pays for the restated numbers in full and saves 13.19. The two renames save 2.49 and 0.09, and each needs an endpoint an absorb holds.




Figure 25: Apple's cash flow lines as strands across the fiscal 2020 to 2022 reports. A filled dot is a strand that ends, an open dot one that starts; the number on a join is the one that carried.




Figure 26: The six candidates at the fiscal 2021 to 2022 boundary, each joining the endpoints it would cover.

Rows. Figure 28, as last reported. 2019 is ruled by the fiscal 2021 report and keeps the three lines; 2020 onwards by the fiscal 2022 report, with them inside “Other”. Every column adds up as its report did.




Figure 27: The solution: the two absorbs solid, the four candidates not chosen dashed. Every endpoint is covered.

from	2019 2021 report	2020 2022 report	2021 2022 report	2022 2022 report	
Proceeds from issuance of common stock	781	-	-	-	absorbed into Other from 2020
Other (financing)	(105)	754	976	(160)	$754 = 880 + (126)$; $976 = 1,105 + (129)$
Payments for dividends	(14,119)	(14,081)	(14,467)	(14,841)	
Repurchases of common stock	(66,897)	(72,358)	(85,971)	(89,402)	
Cash used in financing activities	(90,976)	(86,820)	(93,353)	(110,749)	with four lines not shown
Purchases of non-marketable securities	(1,001)	-	-	-	absorbed into Other from 2020
Proceeds from non-marketable securities	1,634	-	-	-	absorbed into Other from 2020
Other (investing)	(1,078)	(909)	(352)	(1,780)	$(909) = (791) + (210) + 92$
Cash generated by/(used in) investing	45,896	(4,289)	(14,545)	(22,354)	with the other lines not shown

Figure 28: The rows that follow, as last reported (amber: restated across an absorb; the totals include lines not shown).

C.3 Palantir

The facts. Table 9: in the fiscal 2021 report every column is the 2020 report's column moved down one line. The preferred-stock line is gone and its numbers sit on the stock-options tag; the stock-option numbers sit on the debt tag; and so on down to the sale-of-assets tag, whose own numbers are gone. Labels stayed with their tags. The 2022 report keeps it so and drops the stock-options tag; the 2023 report puts the stock-option numbers back on their own tag and drops the other three [18–20].

Strands. Figure 29. At the first boundary no tag's numbers agree, so every tag is a strand that ends and one that starts (e_1 to e_4 , s_1 to s_4), and the preferred-stock line ends (e_5). The new strands run on into the 2022 report, where the stock-options one ends (e_6); at the last boundary the other three end (e_7 , e_8 , e_9) and the stock-options tag starts again (s_5). Fourteen endpoints.

Candidates. Figure 30. At the first boundary:

	2020 REPORT		2021 REPORT		2022 REPORT		2023 REPORT	
	2019	2020	2019	2020	2021	2022	2021	2022
Issuance of preferred stock	7.5	0.0	—	—	—	—	—	—
Exercise of stock options	16.9	298.8	7.5	0.0	—	—	507.5	86.1
Issuance of debt	544.4	199.4	16.9	298.8	507.5	86.1	—	—
Issuance of common stock	100.0	942.5	544.4	199.4	0.0	0.0	—	—
Sale of assets held for sale	0.0	0.3	100.0	942.5	0.0	0.0	—	—

Table 9: Five financing lines of Palantir’s cash flow statement across four consecutive reports [18–20]. In the 2021 report each column is the 2020 report’s moved down one line; the 2022 report keeps it so; the 2023 report moves the stock-option numbers back.




Figure 29: Palantir's financing lines as strands across the fiscal 2020 to 2023 reports.

- W: a swap of three legs, stock options to debt, debt to common stock, common stock to sale of assets. One act of 9.00; each leg at the lowest price, $w_{\min}$, since along it every number carries verbatim.
- P: the preferred-stock line renamed to the stock-options tag, 8.17: the fourth leg, priced as a rename because the old tag ends.
- r_1 to r_4 : each tag continuing under its own name, with every number restated.

At the last boundary, X: the debt tag’s strand renamed to the returning stock-options tag, at $w_{\min}$, since 507.5 and 86.1 carry.

Solution. Figure 31: W, P and X. The swap covers six endpoints at once and saves 126.46; the same-name continuations save 24.75, 36.75, 30.75 and 19.55, and each needs an endpoint the swap holds. X is the correction, and nothing competes with it. Four strands end uncovered: the



Figure 30: The candidates at the first and last boundaries.

old sale-of-assets series at the first boundary, the stock-options tag's run at the second, and the common-stock and sale tags at the last. The swap separates each tag's run before the boundary from its run after.




Figure 31: The solution: the swap, the preferred-stock rename and the correction solid; the same-name continuations dashed; four strands ended.

Rows. Figure 32. Each row is a chain named by the tag it ends on, and each is one series: the stock-option proceeds run 16.9, 298.8, 507.5, 86.1, 218.2 across three tags and end under their own

tag; the debt and common-stock series end where their borrowed tags were dropped. Years on a borrowed tag are flagged. As first reported, the 2021 column shows the numbers on the tags Palantir put them on.

from	2019 2021 report	2020 2022 report	2021 2023 report	2022 2023 report	2023 2023 report	
Exercise of stock options	16.9	298.8	507.5	86.1	218.2	the option series, via the debt tag
Issuance of common stock	544.4	199.4	-	-	-	the debt series, via the stock tag
Sale of assets held for sale	100.0	942.5	-	-	-	the stock series, via the sale tag
Exercise of stock options	7.5	-	-	-	-	the preferred series, then dropped

Figure 32: The rows that follow, as last reported: one series per row, each named by the tag it ends on (amber: years on a borrowed tag, flagged).

C.4 Tesla

The facts. Table 10: the 2016 report restates Automotive (3,741.0 to 3,431.6), adds Automotive leasing (309.4) and a total (3,741.0) under a new tag, and does the same to Services and other, restated with Energy generation and storage carved out ($305.1 = 290.6 + 14.5$) [22, 25]. Section 4.2 set out three explanations of the automotive change.

	2015 REPORT		2016 REPORT	
	2014	2015	2014	2015
Automotive	3,007.0	3,741.0	2,874.4	3,431.6
Automotive leasing	—	—	132.6	309.4
Total automotive revenue	—	—	3,007.0	3,741.0
Services and other	191.3	305.1	187.1	290.6
Energy generation and storage	—	—	4.2	14.5

Table 10: Five revenue lines of Tesla’s income statement for the same two years, as two consecutive reports gave them [22, 25].

Strands. Figure 33, across the 2015 to 2017 reports [23]. Automotive and Services and other each end at the 2015 report and start again at the 2016 ($e_1, s_1; e_2, s_4$); leasing, the total and energy start there (s_2, s_3, s_5); Total revenues runs through, and at the 2016 to 2017 boundary every strand continues. Seven endpoints.

Candidates. Figure 34: explanations A, B and C of Section 4.2 and the plain restated rename for the automotive endpoints e, s_1, s_2, s_3 , and S for the services endpoints. Table 11 lists the automotive ones as sets with the check each passed; all four hold e , so they are one group and at most one can be chosen. Table 12 prices them: a candidate’s w_c plus the p_n of what it leaves uncovered.

Solution. C and S (Figure 35). A is the cheapest single candidate but leaves the most uncovered; C pays the most and leaves nothing. C’s price is A’s rename plus the two broken-out lines priced




Figure 33: Tesla's revenue lines as strands across the 2015 to 2017 reports.




Figure 34: Every candidate at the 2015 to 2016 boundary.

CANDIDATE					EXPLANATION	WHAT MUST HOLD
	E	S ₁	S ₂	S ₃		
rename to itself (restarted)	✓	✓				$3,741.0 \rightarrow 3,431.6$
rename to the total	✓			✓	A	$3,741.0 = 3,741.0$
carve leasing out	✓	✓	✓		B	$3,741.0 = 3,431.6 + 309.4$
rename to the total, break out	✓	✓	✓	✓	C	both of the above
uncovered, so paid for as started		A	A	B		

Table 11: Tesla's 2015 to 2016 boundary as a choice, with some of its candidates. A tick means the candidate covers that endpoint; no two chosen candidates may share one. The last column is the check each candidate passed; the last row shows which explanation leaves each endpoint uncovered.

EXPLANATION	PRICE W_c	COVERS	SAVES S_c	LEAVES UNCOVERED	TOTAL PRICE
A: rename to the total	10.87	e, s_3	41.13	s_1 (31.60), s_2 (28.00)	70.47
B: restated, leasing carved out	32.94	e, s_1 , s_2	47.06	s_3 (31.60)	64.54
rename to itself, restated	14.94	e, s_1	37.06	s_2 , s_3	74.54
C: rename to the total, break out	55.67	e, s_1 , s_2 , s_3	55.93	nothing	55.67

Table 12: Tesla’s 2015 to 2016 boundary, priced. The endpoints cost $p_e = 20.40$, $p_{s_1} = 31.60$, $p_{s_2} = 28.00$, $p_{s_3} = 31.60$ uncovered. An explanation’s total is its candidate’s price plus the price of what it leaves uncovered; C wins by 8.87.

as lines that start, less one flat starting charge and less a credit for the sum the total verifies over three years. A is cheaper line by line; C is cheaper for the boundary as a whole, and under it the total’s history is like for like.

Rows. Figure 35 shows the chosen edits over all three reports; at the 2016 to 2017 boundary every strand continues and nothing is chosen.

Figure 36 is the table that follows, as last reported. Each rename joins two strands into one row: the old Automotive line and the new total are one row, like for like across all five years, and Services and other is one row across its restatement. The three broken-out or carved-out lines begin in 2014. Each column is ruled by the latest report that shows it: 2013 by the 2015 report, 2014 by the 2016 report, 2015 onwards by the 2017 report; where a joined row crosses a restatement, the ruling filing’s numbers are shown and marked.

Figure 37 is the same rows as first reported, each year ruled by its own report. The total’s 2014 and 2015 are the 2015 report’s 3,007.0 and 3,741.0 under the tag it had then; lines the 2015 report did not show are blank there rather than filled from later, including the new Automotive line, whose tag the 2015 report used for what is now the total. Every column of either table adds up as its ruling report did.

THE EDITS CHOSEN AT THE 2015 TO 2016 BOUNDARY

at the 2016 to 2017 boundary every strand continues and nothing is chosen




Figure 35: Tesla's revenue lines across three annual reports as strands, with the edits chosen at the 2015 to 2016 boundary. Chosen candidates are solid, the others dashed; every endpoint is covered. At the next boundary every strand continues.

from	2013 2015 report	2014 2016 report	2015 2017 report	2016 2017 report	2017 2017 report
Total revenues	2,013.5	3,198.4	4,046.0	7,000.1	11,758.8
Total automotive revenue	1,921.9	3,007.0	3,741.0	6,350.8	9,641.3
Automotive	.	2,874.4	3,431.6	5,589.0	8,534.8
Automotive leasing	.	132.6	309.4	761.8	1,106.5
Services and other	91.6	187.1	290.6	468.0	1,001.2
Energy generation and storage	.	4.2	14.5	181.4	1,116.3

one row across the rename, like for like broken out in 2014

broken out in 2014

one row across the restatement

starts in 2014

Figure 36: The rows that follow, as last reported, in millions. Each column's numbers are its ruling report's, the latest that shows the year. A rename joins two strands into one row; a number in amber was restated across the join; a dot is a year before a line was broken out or started.

from	2013 2015 report	2014 2015 report	2015 2015 report	2016 2016 report	2017 2017 report
Total revenues	2,013.5	3,198.4	4,046.0	7,000.1	11,758.8
Total automotive revenue	1,921.9	3,007.0	3,741.0	6,350.8	9,641.3
Automotive	.	.	.	5,589.0	8,534.8
Automotive leasing	.	.	.	761.8	1,106.5
Services and other	91.6	191.3	305.1	468.0	1,001.2
Energy generation and storage	.	.	.	181.4	1,116.3

as the 2015 report gave them, on its tag the 2015 report's Automotive is the total's

no such line in the 2015 report

no such line in the 2015 report

Figure 37: The same rows as first reported: each column ruled by the report whose year it is. A line the ruling report did not show is blank; a number is never taken from a later report.

C.5 JPMorgan

The facts. Table 13: JPMorgan’s 2023 report replaces one line of non-interest revenue with two, and restates both comparative years into the parts: $14,405 + 6,624 = 21,029$ and $14,096 + 6,581 = 20,677$ [13, 14].

	2022 REPORT		2023 REPORT	
	2021	2022	2021	2022
Asset management, administration and commissions	21,029	20,677	—	—
Asset management fees	—	—	14,405	14,096
Commissions and other fees	—	—	6,624	6,581

Table 13: One line of JPMorgan’s income statement becoming two, for the same two years in two consecutive reports [13, 14].

Strands. Figure 38. The old line joins its strand from the 2021 to the 2022 report and ends there (e_1); the two parts start in the 2023 report (s_1 , s_2). Three endpoints.




Figure 38: JPMorgan's asset management line as strands across the 2021 to 2023 reports.

Candidates. Figure 39. Two, for the same three endpoints:

- S: a split of the old line into the two parts, 26.15.
- R: the old line renamed to asset management fees, with commissions carved out of it.

Solution. Figure 40: S, which saves 19.85. Both pass the same sums, but R pays the name term of a rename to a line with a different label, while a split’s name term is judged against both parts together; R saves 3.83.

Rows. Figure 41. The old line continues as the sum of its parts, marked as derived and cited to both; 22,056 for 2023 is in no filing. The parts are blank before 2021.




Figure 39: The two candidates at the 2022 to 2030 boundary, both covering the same three endpoints.




Figure 40: The solution: the split solid, the rename with a carve-out dashed.

from	2020 2022 report	2021 2023 report	2022 2023 report	2023 2023 report	
Asset management, administration and commissions	18,177	21,029	20,677	22,056	from 2021, the sum of the two parts
Asset management fees	-	14,405	14,096	15,220	starts in 2021
Commissions and other fees	-	6,624	6,581	6,836	starts in 2021

Figure 41: The rows that follow (blue: derived as the sum of the parts).

D One hard case per kind

The four cases of Appendix C were chosen because they are instructive. The cases here were chosen from the edits the optimization made on the two sets so that each can be checked: one edit at one boundary, where the names give little or no help, small enough to verify against the two filings named by accession number, and with the boundary inside the fifteen-year window so the rows that follow can be shown. Each case has two figures. The first shows the two reports side by side with the edit drawn over them, in the conventions of Appendix B: what the company printed is in serif, the edit in blue, the shared year shaded, a restated number circled in orange. The second is the excerpt of the finished statement, as last reported: a restated number in orange, a derived number in blue, a dot where a line has no number in that year. Numbers are as filed, at the filer's own scale; a number in parentheses is negative. Ended and started need no example.

D.1 Rename, to a company's own tag: Tupperware Brands

Tupperware's fiscal 2013 report drops *Indefinite-lived intangible assets* and shows *Trademarks and tradenames, net* under a company tag instead (Figure 42). Tag and label both change and share no word; the only evidence is that the December 2012 balance, 138.4, carries, while the neighbouring lines do not move. The rows (Figure 43) are one line across the rename, and across a second rename in a later report.

TUPPERWARE BRANDS, FISCAL 2012 TO 2013

Tupperware Brands Corporation · Form 10-K, fiscal 2012

0001008654-13-000008

Consolidated balance sheets · \$ millions

	Dec 2011	Dec 2012
Indefinite-lived intangible assets	157.1	138.4
<i>IndefiniteLivedIntangibleAssetsExcludingGoodwill</i>		
Other intangible assets, net	7.2	5.0
<i>FiniteLivedIntangibleAssetsNet</i>		

Tupperware Brands Corporation · Form 10-K, fiscal 2013

0001008654-14-000004

Consolidated balance sheets · \$ millions

	Dec 2012	Dec 2013
Trademarks and tradenames, net	138.4	125.7
<i>new tag tup:TrademarksandTradenamesNet</i>		
Other intangible assets, net	5.0	3.2
<i>FiniteLivedIntangibleAssetsNet</i>		

the check, on December 2012

$$\frac{138.4}{\text{old tag}} = \frac{138.4}{\text{new tag}}$$

Figure 42: Tupperware Brands, balance sheet, fiscal 2012 to 2013: indefinite-lived intangible assets renamed to trademarks and tradenames, net, under a company tag. The December 2012 balance carries.

	2011	2012	2013	2014
Trademarks and tradenames, net	157.1	138.4	125.7	104.2
Other intangible assets, net	7.2	5.0	3.2	1.5
Goodwill	241.4	192.9	181.5	164.7

one row across two tags

Figure 43: The rows that follow, \$ millions: one line across the rename.

D.2 Rename, to a standard's tag: NOV

The new revenue standard gave contract positions new tags and new names. NOV's *Costs in excess of billings* becomes *Contract assets* in the fiscal 2018 report, with the December 2017 balance of 495 carried (Figure 44). The liability side is a limit of the method (Figure 45). *Billings in excess of costs* was 279 for December 2017 and *Contract liabilities* is 519, and the report says why: 240 of advance payments and deferred revenue moved out of *Accrued liabilities*, which is restated from 1,478 to 1,238. That is one move between two continuing lines, and no kind of edit in Section 4.4.2 describes it; a carve-out's children are lines that start, not lines that continue. So the solver leaves a line ended, a line started and a bare restatement of accrued liabilities, three edits where the filing describes one. Every cell holds the right number, but the record of changes is incomplete.

NOV Inc. • Form 10-K, fiscal 2017

0001193125-18-048333

Consolidated balance sheets • \$ millions

	Dec 2016	Dec 2017
Costs in excess of billings	665	495
CostsInExcessOfBillingsOnUncompletedContracts...		

NOV Inc. • Form 10-K, fiscal 2018

0001564590-19-002967

Consolidated balance sheets • \$ millions

	Dec 2017	Dec 2018
Contract assets	495	565
new tag ContractWithCustomerAssetNetCurrent		

the check, on December 2017

$$\begin{array}{ccc} 495 & = & 495 \\ \text{old tag} & & \text{new tag} \end{array}$$

Figure 44: NOV, balance sheet, fiscal 2017 to 2018: costs in excess of billings renamed to contract assets under the new standard's tag.

	2016	2017	2018	
Contract assets	665	495	565	one row across the rename
Billings in excess of costs	440	.	.	ended; the 2018 report has no such line
Accrued liabilities	1,568	1,238	1,088	1,478 before: 240 moved to contract liabilities
Contract liabilities	.	519	458	started: 279 + 240

Figure 45: The rows that follow, \$ millions. The asset line is one row. On the liability side the method leaves a line ended, a line started and a bare restatement, where the filing describes one move of 240.

D.3 Absorb: Palantir

Palantir's fiscal 2021 report has no *Change in fair value of warrants* line, and *Other income (expense), net* for 2020 reads 4,111 where the year before it read 3,300 (Figure 46). The difference is the warrant line, 811, to the dollar, and 2019 moves by the same line's 3. Nothing in either name relates the two, and 0.8 million is small in an income statement measured in billions. Every other line is unchanged. The rows (Figure 47) carry the restated 4,111, marked; the warrant line has no row.

Palantir Technologies Inc. • Form 10-K, fiscal 2020

0001193125-21-060650

Consolidated statements of operations • \$ thousands

	2019	2020
Change in fair value of warrants	(3)	811
FairValueAdjustmentOfWarrants		
Other income (expense), net	(2,853)	3,300
OtherNonoperatingIncomeExpense		

Palantir Technologies Inc. • Form 10-K, fiscal 2021

0001193125-22-050913

Consolidated statements of operations • \$ thousands

	2020	2021
Other income (expense), net	4,111	(75,415)
OtherNonoperatingIncomeExpense		

the check, on 2020

$$\begin{array}{ccccc} 4,111 & = & 3,300 & + & 811 \\ \text{new other income} & & \text{old other income} & & \text{warrants} \end{array}$$

Figure 46: Palantir, income statement, fiscal 2020 to 2021: the warrant line absorbed into other income (expense), net. Amounts as they enter income.

	2020	2021	2022	
Interest income	4,680	1,607	20,309	
Other income (expense), net	4,111	(75,415)	(220,135)	2020: 3,300 + 811

Figure 47: The rows that follow, \$ thousands, as last reported.

D.4 Carve-out: Enterprise Financial Services

Enterprise Financial's fiscal 2023 report adds a line, *Deposit costs*, and restates *Other expenses* down by exactly it in both comparative years: 88,955 becomes 57,873 plus 31,082 for 2022, and 62,675 becomes 48,464 plus 14,211 for 2021 (Figure 48). The carved-out line shares no word with the line it came from, and the line it came from is called *Other*. The rows (Figure 49) show other expense restated for the two years the carve-out reached and deposit costs starting with the earliest year the 2023 report shows.

ENTERPRISE FINANCIAL, FISCAL 2022 TO 2023

Enterprise Financial Services Corp · Form 10-K, fiscal 2022

0001025835-23-000022

Consolidated statements of operations · \$ thousands

	2021	2022
Other expenses	62,675	88,955
OtherNoninterestExpense		

← 2022 in both →

continues, restated

carved out

Enterprise Financial Services Corp · Form 10-K, fiscal

0001025835-24-000017

Consolidated statements of operations · \$ thousands

	2022	2023
Other expense	57,873	73,886
OtherNoninterestExpense		
Deposit costs	31,082	72,293
new line efsc:DepositCosts		

the check, on 2022

$$\begin{array}{rcccl} 88,955 & = & 57,873 & + & 31,082 \\ \text{old other expenses} & & \text{new other expense} & & \text{deposit costs} \end{array}$$

Figure 48: Enterprise Financial Services, income statement, fiscal 2022 to 2023: deposit costs carved out of other expenses.

	2021	2022	2023	
Deposit costs	14,211	31,082	72,293	starts; earlier years from the 2023 report
Other expense	48,464	57,873	73,886	62,675 and 88,955 before the carve-out

Figure 49: The rows that follow, \$ thousands, as last reported.

D.5 Merge: Costco

Costco's fiscal 2023 report shows one line, *Impairment of assets and other non-cash operating activities, net*, where the year before showed *Deferred income taxes* and *Other non-cash operating activities, net*. The sum holds on both years the two reports share: $(59) + (85) = (144)$ for 2021 and $37 + (76) = (39)$ for 2022 (Figure 50). The new tag is the company's own and shares one word with one of the old ones. In the rows (Figure 51) the merged line runs back through 2020 as the cited sum of its parts, and the parts are blank after the merge. A merge on a balance sheet can be checked on one instant only, since consecutive reports share one; here two years foot, which is what the derived 2020 figure rests on.

Costco Wholesale Corporation · Form 10-K, fiscal 2022

0000909832-22-000021

Consolidated statements of cash flows · \$ millions

	2021	2022	← 2021 and 2022 in both →
Deferred income taxes IncreaseDecreaseInDeferredIncomeTaxes	(59)	37	merged
Other non-cash operating activities, net OtherNoncashIncomeExpense	(85)	(76)	

Costco Wholesale Corporation · Form 10-K, fiscal 2023

0000909832-23-000042

Consolidated statements of cash flows · \$ millions

	2021	2022
Impairment of assets and other non-cash operating activities, net new line cost: ImpairmentOfAssetsAndOtherNonCash...	(144)	(39)

the check, on 2022 (and on 2021: (59) + (85) = (144))

$$\begin{array}{rcccl} \underline{37} & + & \underline{(76)} & = & \underline{(39)} \\ \text{deferred taxes} & & \text{other non-cash} & & \text{merged line} \end{array}$$

Figure 50: Costco, cash flow statement, fiscal 2022 to 2023: two operating adjustments merged into one line, checked on both shared years.

	2020	2021	2022	2023	
Impairment of assets and other non-cash operating activities, net	(104)	(144)	(39)	(495)	2020 = (104) + (42), from tl
Deferred income taxes	(104)	-	-	-	merged from 2021
Other non-cash operating activities, net	(42)	-	-	-	merged from 2021

Figure 51: The rows that follow, \$ millions.

D.6 Split: Macatawa Bank

Macatawa's fiscal 2022 report replaces *Principal paydowns on securities* with two lines, one for securities available for sale and one for securities held to maturity; for 2021 the two are 31,075 and 9,470, and the old line was 40,545 (Figure 52). The old tag is the company's own and the new ones are standard tags for proceeds from sales, which is not what the lines say. The rows (Figure 53) keep the old line as the cited sum of its parts from 2021 on.

Macatawa Bank Corporation · Form 10-K, fiscal 2021

0001140361-22-005958

Consolidated statements of cash flows · \$ thousands

	2020	2021
Principal paydowns on securities	41,938	40,545
mcbc:PrincipalPaydownsOnSecurities		

← 2021 in both →

split

Macatawa Bank Corporation · Form 10-K, fiscal 2022

0001140361-23-007466

Consolidated statements of cash flows · \$ thousands

	2021	2022
Paydowns on securities available for sale	31,075	14,918
new line ProceedsFromSaleOfAvailableForSaleSecuritiesDebt		
Paydowns on securities held to maturity	9,470	36,053
new line ProceedsFromSaleOfHeldToMaturitySecurities		

the check, on 2021

$$\begin{array}{c} 40,545 \\ \text{old line} \end{array} = \begin{array}{c} 31,075 \\ \text{available for sale} \end{array} + \begin{array}{c} 9,470 \\ \text{held to maturity} \end{array}$$

Figure 52: Macatawa Bank, cash flow statement, fiscal 2021 to 2022: one paydowns line split into two.

	2020	2021	2022	
Principal paydowns on securities	41,938	40,545	50,971	from 2021, the sum of the two parts
... available for sale	-	31,075	14,918	starts in 2021
... held to maturity	-	9,470	36,053	starts in 2021

Figure 53: The rows that follow, \$ thousands.

D.7 Fold: Berkshire Hathaway

Berkshire's fiscal 2019 report shows *Investment gains (losses)* and *Derivative contract gains (losses)* as the two parts of *Investment and derivative gains/losses*. The fiscal 2020 report shows only the total, relabelled, with 2019 still 72,607 (Figure 54). Two lines vanish and the surviving line's label changes in the same report, which reads as a rename and two lines ended until the declared arithmetic is checked. As last reported (Figure 55) the parts have no row and the total is one row, through the fold and through a change of tag in the 2022 report.

Berkshire Hathaway Inc. · Form 10-K, fiscal 2019

0001564590-20-005874

Consolidated statements of earnings · \$ millions

	2018	2019
Investment gains (losses) GainLossOnInvestments	(22,155)	71,123
Derivative contract gains (losses) GainLossOnDerivativeInstrumentsNetPretax	(300)	1,484
Investment and derivative gains/losses NonoperatingIncomeExpense	(22,455)	72,607

Berkshire Hathaway Inc. · Form 10-K, fiscal 2020

0001564590-21-009611

Consolidated statements of earnings · \$ millions

	2019	2020
Investment and derivative contract gains/losses NonoperatingIncomeExpense	72,607	40,746

the check, on 2019

$$\begin{array}{rcl} 71,123 & + & 1,484 \\ \text{investment gains} & & \text{derivative gains} \end{array} = \begin{array}{r} 72,607 \\ \text{total, in both} \end{array}$$

Figure 54: Berkshire Hathaway, income statement, fiscal 2019 to 2020: two gain lines folded into their total.

	2018	2019	2020	2021
Investment and derivative, contract gains/losses	(22,455)	72,607	40,746	78,542

one row; the 2022 report m

Figure 55: The rows that follow, \$ millions, as last reported.

D.8 Break out: JPMorgan Chase

JPMorgan's fiscal 2019 report shows one line of investment securities, 398,239 at December 2019, with the available-for-sale portion in a parenthetical. The fiscal 2020 report shows the same total under a new tag and, beneath it, held-to-maturity and available-for-sale securities as separate lines, 47,540 and 350,699 for December 2019 (Figure 56). The parts are material in the shared year and add up to a total that did not move; that is what a break out is, and what tells it apart from new lines that are zero in the shared years, which are lines that started (Section 4.4.2). The rows (Figure 57) carry the parts from 2019 and the total as one row across its change of tag.

JPMorgan Chase & Co. · Form 10-K, fiscal 2019

0000019617-20-000257

Consolidated balance sheets · \$ millions

	Dec 2018	Dec 2019
Investment securities (includes \$350,699 at fair value)	261,828	398,239
DebtSecuritiesAvailableForSaleAndHeldToMaturity		

JPMorgan Chase & Co. · Form 10-K, fiscal 2020

0000019617-21-000236

Consolidated balance sheets · \$ millions

	Dec 2019	Dec 2020
Held-to-maturity securities	47,540	201,821
new line jpm:DebtSecuritiesHeldToMaturityNetOfAllowance...		
Available-for-sale securities	350,699	388,178
new line AvailableForSaleSecuritiesDebtSecurities		
Investment securities, net of allowance for credit losses	398,239	589,999
new tag jpm:DebtSecuritiesNetCarryingAmount		

the check, on December 2019

$$\begin{array}{rcccl} 47,540 & + & 350,699 & = & 398,239 \\ \text{held to maturity} & & \text{available for sale} & & \text{total, in both} \end{array}$$

Figure 56: JPMorgan Chase, balance sheet, fiscal 2019 to 2020: investment securities broken out into held-to-maturity and available-for-sale, under a total that changes tag and does not move.

	2019	2020	2021	2022	
Held-to-maturity securities	47,540	201,821	363,707	425,305	starts; 2019 from the 2020 report
Available-for-sale securities	350,699	388,178	308,525	205,857	starts; 2019 from the 2020 report
Investment securities, net	398,239	589,999	672,232	631,162	one row, across the rename

Figure 57: The rows that follow, \$ millions.

D.9 Equal copies: KalVista Pharmaceuticals

KalVista reported one line, *Net loss per share, basic and diluted*, until its fiscal 2022 report, which reports basic and diluted on two lines; for fiscal 2021 each reads (2.42), as the one line did (Figure 58). A split would test whether the two add up to the old line, and they do not; each equals it. The rows (Figure 59) are two lines, each running back through the combined line's years.

KalVista Pharmaceuticals, Inc. • Form 10-K, fiscal 2021

0001564590-21-036441

Consolidated statements of operations • \$ per share

	FY2020	FY2021	← FY2021 in both →
Net loss per share, basic and diluted EarningsPerShareBasicAndDiluted	(1.64)	(2.42)	equal copies

KalVista Pharmaceuticals, Inc. • Form 10-K, fiscal 2022

0000950170-22-012521

Consolidated statements of operations • \$ per share

	FY2021	FY2022
Net loss per share, basic new line EarningsPerShareBasic	(2.42)	(3.36)
Net loss per share, diluted new line EarningsPerShareDiluted	(2.42)	(3.36)

the check, on fiscal 2021: each copy equals the old line

$$\begin{array}{ccc} \underline{(2.42)} & = & \underline{(2.42)} = \underline{(2.42)} \\ \text{basic and diluted} & & \text{basic} \quad \text{diluted} \end{array}$$

Figure 58: KalVista Pharmaceuticals, income statement, fiscal 2021 to 2022: one per-share line becomes two equal copies.

	FY2020	FY2021	FY2022	FY2023	
Net loss per share, basic	(1.64)	(2.42)	(3.36)	(3.33)	2020 and 2021 from the combinec
Net loss per share, diluted	(1.64)	(2.42)	(3.36)	(3.33)	the same

Figure 59: The rows that follow, \$ per share.

D.10 Swap: FormFactor

In FormFactor's fiscal 2012 report the deferred-revenue tag carries *Deferred revenues*, 1,064 for 2012, and the other-deferred- liability tag carries *Deferred rent and other liabilities*, 167. In the fiscal 2013 report the two tags have exchanged both labels and numbers (Figure 60). Following the tag gives each line the other's history; following the label is right, but only a swap says why every number under each tag disagrees with the year before. The rows (Figure 61) follow the numbers.

FORMFACTOR, FISCAL 2012 TO 2013

FormFactor, Inc. • Form 10-K, fiscal 2012

0001445305-13-000585

Consolidated statements of cash flows • \$ thousands

	2011	2012	← 2012 in both →
Deferred revenues IncreaseDecreaseInDeferredRevenue	160	1,064	numbers and labels swapped
Deferred rent and other liabilities IncreaseDecreaseInOtherDeferredLiability	43	167	same tag, not the same line

FormFactor, Inc. • Form 10-K, fiscal 2013

0001039399-14-000007

Consolidated statements of cash flows • \$ thousands

	2012	2013
Deferred rent and other liabilities IncreaseDecreaseInDeferredRevenue	167	83
Deferred revenues IncreaseDecreaseInOtherDeferredLiability	1,064	1,014

the same 2012 numbers, on each other's tag

$$\begin{array}{ccc} \underline{1,064} & \cdot & \underline{167} \\ \text{deferred revenues} & & \text{deferred rent} \end{array}$$

Figure 60: FormFactor, cash flow statement, fiscal 2012 to 2013: two tags exchange their labels and their numbers.

	2012	2013	2014	
Deferred revenues	1,064	1,014	(726)	one row, across two tags
Deferred rent and other liabilities	167	83	248	one row, across two tags; later labelled Other liabil

Figure 61: The rows that follow, \$ thousands: each line follows its numbers, not its tag.

D.11 Carve-out across the balance sheet: DXP Enterprises

DXP's fiscal 2014 report changes its contract positions from a net to a gross presentation. *Trade accounts receivable, net* for December 2013 was 192,003; the new report shows receivables of 186,854, a new asset line, *Costs and estimated profits in excess of billings*, 6,487, and a new liability, *Billings in excess of costs*, 1,338, and it restates both section totals up by 1,338 (Figure 62). The edit is a carve-out of both new lines from receivables: 192,003 is 186,854 plus 6,487 less 1,338, the liability entering with -1 because its balance attribute is the opposite of the receivable's (Section 4.4.2), and the three totals above them, current assets, current liabilities and total assets, are explained by what the edit moved into each. No name relates any of the three lines to the others. The rows (Figure 63) keep receivables as one row, restated where the carve-out reached, and show both totals' restatements marked.

DXP ENTERPRISES, FISCAL 2013 TO 2014

DXP Enterprises, Inc. · Form 10-K, fiscal 2013

0001020710-14-000008

Consolidated balance sheets · \$ thousands

	Dec 2012	Dec 2013
Trade accounts receivable, net AccountsReceivableNetCurrent	174,832	192,003
Total current assets AssetsCurrent	295,702	313,149
Total current liabilities LiabilitiesCurrent	128,452	148,717

DXP Enterprises, Inc. · Form 10-K, fiscal 2014

0001020710-15-000006

Consolidated balance sheets · \$ thousands

	Dec 2013	Dec 2014
Trade accounts receivable, net AccountsReceivableNetCurrent	186,854	239,236
Costs and estimated profits in excess of billings new line CostsInExcessOfBillingsOnUncompletedContracts...	6,487	20,083
Total current assets AssetsCurrent	314,487	386,278
Billings in excess of costs (a current liability) new line BillingsInExcessOfCostCurrent	1,338	8,840
Total current liabilities LiabilitiesCurrent	150,055	207,202

the check, on December 2013

$$\begin{array}{ccccccc}
 \underline{192,003} & = & \underline{186,854} & + & \underline{6,487} & + & \underline{(1,338)} \\
 \text{old receivables} & & \text{new receivables} & & \text{costs in excess} & & \text{billings, a liability}
 \end{array}$$

Figure 62: DXP Enterprises, balance sheet, fiscal 2013 to 2014: two lines carved out of receivables, one on each side of the balance sheet, with both section totals explained.

	2012	2013	2014	
Trade accounts receivable, net	174,832	186,854	239,236	192,003 before the carve-out
Costs and estimated profits in excess of billings	-	6,487	20,083	starts in 2013
Total current assets	295,702	314,487	386,278	313,149 + 6,487 - 5,149
Billings in excess of costs (liability)	-	1,338	8,840	starts in 2013
Total current liabilities	128,452	150,055	207,202	148,717 + 1,338

Figure 63: The rows that follow, \$ thousands, as last reported.

E Measurements in detail

E.1 How often the declared arithmetic holds

The arithmetic says which lines add up to which totals, and with what sign, stated once per tag in a separate file of the filing (Figure 64). A filing must declare it wherever a statement shows lines and their total [31], but a total that does not add up has no effect on whether EDGAR accepts the filing [30, 31], so how often the equations hold is a measurement. We made it over the 200 companies drawn at random for Section 5 and their 2,826 annual reports. Over 99% of declared totals state a number. A check is one declared total in one period, asking whether the parts the same filing states add up to it; there are 104,408. Read as declared, with each breakdown row counted as a part, 87.7% hold exactly; folding the breakdown rows into their line and allowing the filer's rounding brings that to 89.8%. Most of the rest have a declared part with no number in that period; among totals whose every part carries a number, 98.7% add up (Table 14). Breakdowns are the exception: a declared sum relates one tag to another in the same context, so a line's breakdown is never declared to add up to it [33], and a breakdown can confirm a number but cannot prove one wrong.

(a) The label the cash flow statement shows for Apple's tag in Figure 2(a), from the label file aapl-20210925_lab.xml

```
<link:label id="lab_us-gaap_ProceedsFromIssuanceOfCommonStock_77a3..._terseLabel_en-US"
  xlink:label="lab_us-gaap_ProceedsFromIssuanceOfCommonStock"
  xlink:role="http://www.xbrl.org/2003/role/terseLabel"
  xlink:type="resource"
  xmlns:xml="http://www.w3.org/XML/1998/namespace"
  xml:lang="en-US">Proceeds from issuance of common stock</link:label>
```

(b) The declared sum it takes part in, from the calculation file aapl-20210925_cal.xml: it adds into net cash from financing with weight +1

```
<link:calculationArc order="1"
  weight="1.0"
  xlink:arcrole="http://www.xbrl.org/2003/arcrole/summation-item"
  xlink:from="loc_us-gaap_NetCashProvidedByUsedInFinancingActivities_5903..."
  xlink:to="loc_us-gaap_ProceedsFromIssuanceOfCommonStock_63d0..."
  xlink:type="arc"/>
```

Figure 64: What a filing says about a tag, as filed in Apple's 10-K for fiscal 2021. The label and the arithmetic are stated once per tag, in separate files of the same filing, not on each fact [1].

DECLARED TOTALS, ONE TOTAL IN ONE PERIOD	CHECKS	SHARE
checks	104,408	
add up exactly, parts as declared	91,551	87.7%
add up exactly, once breakdown rows are folded into their line	93,320	89.4%
add up within the filer's rounding, breakdown rows folded	93,762	89.8%
do not add up: a declared part has no number in that period	9,561	9.2%
do not add up: every part has a number and the sum is off	1,085	1.0%
totals with every part stated	85,816	
of which add up within the filer's rounding	84,731	98.7%

Table 14: How often a declared total adds up inside its own filing, and why the rest do not, over 200 companies drawn at random and their 2,826 annual reports. The last share is of the totals with every part stated.

E.2 The misses against the vendor, sorted

The vendor of Section 5.1 has a different goal from ours. It standardizes so that companies can be compared with one another [9], the purpose the SEC gave for standard tags [29]; we standardize one company across time and keep its own lines. The vendor maps each filer's lines onto its own fields, nets some, gathers what is left into residual buckets, adjusts per-share figures for later splits, and applies its own definitions of net income and of payables. Each of those puts a number in its series that no filing states for the line we anchored, and Table 15 sorts the misses by them first: 258 of the 334 on the specimen set and 2,683 of the 5,234 on the random sample. Of the rest, the vendor's figure is elsewhere on our page, in another row of the same statement, for 30 and 1,380 cells; in 2 and 129 more the anchored row itself shows the vendor's number in the rendered table, a cell the table has that the lineage test did not reach. The two rows' labels sort those: in about half, 17 and 802, they are one line under two tags, a history the optimization left in two rows, which is our error. It happens where a line's definition changed and the filer kept both lines for a while (cash with and without restricted cash, from 2018: the old tag continues beside the new total, so by our own rule they are two lines, which the vendor joins), and where the tag changed across a gap in our filings, so the two lives share no year and the rename has only names to go on. In the other half the row that has the number is a different line that happens to carry it, or one the vendor prefers (total equity with and without noncontrolling interests, cash with and without restricted cash). What is neither the vendor's standardization nor explained is 44 cells on the specimen set and 1,042 on the random sample; with the split lines, the identity errors the comparison finds or cannot rule out come to 61 cells, 0.6%, and 1,844, 3.1%. These are agreement rates, not accuracy: two series can agree on a wrong number, and a cell that agrees by coincidence is not found by examining disagreements. What the comparison does establish is the coverage gain, and that every cell where the vendor's standardization departs from the filings counts against us. The method and the tag alone are scored against the same series with the same anchors, so the rows of Table 6 can be compared; the absolute level is set by the vendor's standardization, not by us.

WHY THE VENDOR'S NUMBER IS NOT THE ONE WE SHOW	SPECIMEN	RANDOM SAMPLE
the vendor nets two lines (stock issued less repurchased, debt issued less repaid)	121	735
the vendor's own definition (net income, payables)	48	532
per-share figures the vendor adjusts for later splits	45	833
the vendor's residual bucket ("other")	44	583
it is on our page, in another row of the same statement	30	1,380
of which: one line under two tags, a history left in two rows	17	802
of which: a different line that carries the same number	11	462
of which: twin rows of one tag (an opening and a closing balance, a breakdown)	2	116
the anchored row shows it; the lineage test did not reach it	2	129
everything else	44	1,042
all misses	334	5,234

Table 15: The method's misses against the vendor, sorted: first by what the vendor did to the field, then by whether the vendor's number is elsewhere on our page, and what the two rows are.

E.3 The layout differences, and the checks not run

Of the two specimen statements that differ between the newest filing's layout and the oldest presenter's over fifteen years (Section 5.1), GE's cash flow statement differs by one row: the older layout also lists under a segment a figure the newest shows once (share repurchases under "GE"), with the same numbers in both rows, and the checks that use that row move with it; JPMorgan's balance sheet differs only in which checks run, as do 20 of the random sample's 33. The other 13 were rendered under both layouts and every differing row read. In 11 the same numbers sit under a differently named row, one of them a total the newest layout shows twice; in 2, both at one company, a merged line that the newest layout derives as the sum of its parts appears in the older layout as the parts alone. The script that reads the renders marks four rows as changed values: the total shown twice, a dividend the newest layout also lists under the operating partnership, and the two derived rows; in each, wherever both layouts show a number for the row, it is the same number. An earlier run of the check found six more, each a row whose lineage joined a line stated at instants to one stated over periods; under one layout such a row showed the balances and under the other the flows. Three were wrong renames and three were one line a filer had tagged both ways (Appendix A.6).

The checks that cannot be run record why. On the random sample, 14,808 of 19,668 are years for which no filing declares arithmetic for the line, spread across the window; 2,496 have a part with no number in that column; 369 are the first column of a history.

References

- [1] Apple Inc. Form 10-K for the fiscal year ended September 25, 2021. <https://www.sec.gov/Archives/edgar/data/320193/000032019321000105/aapl-20210925.htm>, 2021. EDGAR accession 0000320193-21-000105, filed October 29, 2021. Consolidated statements of cash flows: proceeds from issuance of common stock 1,105, 880 and 781 for fiscal 2021, 2020 and 2019; other (financing) (129), (126) and (105).

- [2] Apple Inc. Form 10-K for the fiscal year ended September 24, 2022. <https://www.sec.gov/Archives/edgar/data/320193/000032019322000108/aapl-20220924.htm>, 2022. EDGAR accession 0000320193-22-000108, filed October 28, 2022. Consolidated statements of cash flows: no proceeds-from-common-stock line; other (financing) (160), 976 and 754 for fiscal 2022, 2021 and 2020.
- [3] Center for Research in Security Prices. CRSP US stock database guide: PERMNO. <https://www.crsp.org>. The permanent security identifier. Accessed 2026-09-16.
- [4] Centers for Medicare & Medicaid Services. ICD-10-CM and ICD-10-PCS general equivalence mappings (GEMs): User’s guide. <https://www.cms.gov/medicare/coding-billing/icd-10-codes>. The guide states that the mappings are not one-to-one and are published separately in each direction. Accessed 2026-09-16.
- [5] Farmhouse, Inc. Form 10-K for the fiscal year ended December 31, 2024, note 2, summary of significant accounting policies: Financial statement reclassification. <https://www.sec.gov/Archives/edgar/data/1811999/000109690625000995/R22.htm>, 2025. EDGAR accession 0001096906-25-000995, filed June 23, 2025. “Certain amounts in prior periods have been reclassified to conform to the current period presentation.” One filer’s instance of a sentence that appears, in this or nearly this wording, in filings of every size.
- [6] Federal Reserve Bank of St. Louis. ALFRED: Archival federal reserve economic data. <https://alfred.stlouisfed.org>. Every vintage of each series, as first published. Accessed 2026-09-16.
- [7] Financial Accounting Standards Board. Accounting standards codification 205-10-50-1, presentation of financial statements: Reclassifications. <https://asc.fasb.org>. When prior-period amounts are reclassified to conform to the current presentation, the notes must say so and describe the nature and magnitude of the change. Read 2026-09-16 via the codification and secondary guidance.
- [8] Financial Modeling Prep. As reported financial statements API. <https://intelligence.financialmodelingprep.com/developer/docs/stable/as-reported-financial-statements>, . Developer documentation for the `financial-statement-full-as-reported` endpoint: “The Financial Raw Filings API provides users with original, unadjusted financial statements directly from company filings.” Read 2026-09-18.
- [9] Financial Modeling Prep. Income statement API. <https://intelligence.financialmodelingprep.com/developer/docs/stable/income-statement>, . Developer documentation. “Competitive Analysis: Use the API to compare a company’s financial performance to its competitors. By analyzing income statements across companies, investors can identify which businesses are leading in profitability and efficiency.” Read 2026-09-18.
- [10] Git project. `git-diff(1): -M, -find-renames`. <https://git-scm.com/docs/git-diff>. Rename detection by content similarity. Accessed 2026-09-16.
- [11] Gang Hu, Qun Zhang, Jingyao Luo, Yile Jiang, Jing Chai, and Haiyan Ding. XBRLTagRec: Domain-specific fine-tuning and zero-shot re-ranking with LLMs for extreme financial numeral labeling. arXiv:2603.25263, 2026.

- [12] International Organization for Standardization. ISO 3166-1: Codes for the representation of names of countries and their subdivisions — part 1: Country code. Technical report, ISO, Geneva, 2020. CS denoted Czechoslovakia until 1993 and Serbia and Montenegro from 2003 to 2006; the standard now reserves a deleted code for fifty years.
- [13] JPMorgan Chase & Co. Form 10-K for the fiscal year ended December 31, 2022. <https://www.sec.gov/Archives/edgar/data/19617/000001961723000231/jpm-20221231.htm>, 2023. EDGAR accession 0000019617-23-000231, filed February 21, 2023. Consolidated statements of income: asset management, administration and commissions 21,029 and 20,677 for 2022 and 2021.
- [14] JPMorgan Chase & Co. Form 10-K for the fiscal year ended December 31, 2023. <https://www.sec.gov/Archives/edgar/data/19617/000001961724000225/jpm-20231231.htm>, 2024. EDGAR accession 0000019617-24-000225, filed February 16, 2024. Consolidated statements of income: asset management fees 14,405 and 14,096, and commissions and other fees 6,624 and 6,581, for 2022 and 2021, in place of the single line of the year before.
- [15] Netflix, Inc. Form 10-K for the fiscal year ended December 31, 2017. <https://www.sec.gov/Archives/edgar/data/1065280/000106528018000069/q4nflx201710k.htm>, 2018. EDGAR accession 0001065280-18-000069, filed January 29, 2018; amended without change to the statements as 0001628280-18-000941, February 5, 2018. Consolidated balance sheets: deferred revenue 618.6 at December 31, 2017, tagged `us-gaap:DeferredRevenueCurrent`.
- [16] Netflix, Inc. Form 10-K for the fiscal year ended December 31, 2018. https://www.sec.gov/Archives/edgar/data/1065280/000106528019000043/form10k_q418.htm, 2019. EDGAR accession 0001065280-19-000043, filed January 29, 2019. Consolidated balance sheets: deferred revenue 618.6 at December 31, 2017, now tagged `us-gaap:ContractWithCustomerLiabilityCurrent`.
- [17] Bryan O’Sullivan. *Mercurial: The Definitive Guide*. O’Reilly Media, 2009. Chapter 5, on copies and renames: `hg rename` records the rename in the repository.
- [18] Palantir Technologies Inc. Form 10-K for the fiscal year ended December 31, 2020. <https://www.sec.gov/Archives/edgar/data/1321655/000119312521060650/d65934d10k.htm>, 2021. EDGAR accession 0001193125-21-060650, filed February 26, 2021. Consolidated statements of cash flows: proceeds from the exercise of stock options 298.8 and 16.9 for 2020 and 2019, tagged `us-gaap:ProceedsFromStockOptionsExercised`.
- [19] Palantir Technologies Inc. Form 10-K for the fiscal year ended December 31, 2021. <https://www.sec.gov/Archives/edgar/data/1321655/000119312522050913/d273589d10k.htm>, 2022. EDGAR accession 0001193125-22-050913, filed February 24, 2022. Consolidated statements of cash flows: the same stock-option proceeds, 507.5, 298.8 and 16.9 for 2021, 2020 and 2019, tagged `us-gaap:ProceedsFromDebtNetOfIssuanceCosts`; each of the neighbouring lines’ numbers is likewise on the tag of the line below it.
- [20] Palantir Technologies Inc. Form 10-K for the fiscal year ended December 31, 2023. <https://www.sec.gov/Archives/edgar/data/1321655/000132165524000022/pltr-20231231.htm>, 2024. EDGAR accession 0001321655-24-000022, filed February 20, 2024. Consolidated statements of cash flows: stock-option proceeds 218.2, 86.1 and 507.5 for 2023, 2022 and 2021, back on `us-gaap:ProceedsFromStockOptionsExercised`.

- [21] Justin R. Pierce and Peter K. Schott. A concordance between ten-digit U.S. harmonized system codes and SIC/NAICS product classes and industries. *Journal of Economic and Social Measurement*, 37(1-2):61-96, 2012.
- [22] Tesla, Inc. Form 10-K for the fiscal year ended December 31, 2016. https://www.sec.gov/Archives/edgar/data/1318605/000156459017003118/tsla-10k_20161231.htm, 2017. EDGAR accession 0001564590-17-003118, filed March 1, 2017. Consolidated statements of operations: automotive revenue 3,431.6 and automotive leasing 309.4 for 2015, in place of the 3,741.0 of the year before, and total automotive revenue 6,350.8, 3,741.0 and 3,007.0 for 2016, 2015 and 2014 on the company’s own tag; services and other 290.6 and energy generation and storage 14.5 for 2015.
- [23] Tesla, Inc. Form 10-K for the fiscal year ended December 31, 2017. https://www.sec.gov/Archives/edgar/data/1318605/000156459018002956/tsla-10k_20171231.htm, 2018. EDGAR accession 0001564590-18-002956, filed February 23, 2018. Consolidated statements of operations: the revenue lines of the fiscal 2016 report continue under the same tags.
- [24] Tesla, Inc. Form 10-K for the fiscal year ended December 31, 2021. <https://www.sec.gov/Archives/edgar/data/1318605/000095017022000796/tsla-20211231.htm>, 2022. EDGAR accession 0000950170-22-000796, filed February 7, 2022. Consolidated statements of operations: automotive leasing revenue 1,052 for 2020, a fact on the contract revenue tag whose context carries the automotive leasing product member.
- [25] Tesla Motors, Inc. Form 10-K for the fiscal year ended December 31, 2015. https://www.sec.gov/Archives/edgar/data/1318605/000156459016013195/tsla-10k_20151231.htm, 2016. EDGAR accession 0001564590-16-013195, filed February 24, 2016. Consolidated statements of operations: automotive revenue 3,741.0 for 2015.
- [26] U.S. Census Bureau. Relationship files: Census tract relationship files. <https://www.census.gov/geographies/reference-files/time-series/geo/relationship-files.html>. Geography Division. Accessed 2026-09-16.
- [27] U.S. Securities and Exchange Commission. Regulation S-X, 17 CFR part 210: §210.4-02 items not material; §210.5-02 balance sheets; §210.5-03 statements of comprehensive income. <https://www.law.cornell.edu/cfr/text/17/part-210>, . Article 5 lists the line items that “if applicable, and except as otherwise permitted by the Commission, should appear on the face” of the statements; §4-02: “If the amount which would otherwise be required to be shown with respect to any item is not material, it need not be separately set forth. The combination of insignificant amounts is permitted.” Read 2026-09-17.
- [28] U.S. Securities and Exchange Commission. Regulation S-X, 17 CFR part 210: §210.3-01 consolidated balance sheets; §210.3-02 consolidated statements of comprehensive income and cash flows. <https://www.law.cornell.edu/cfr/text/17/210.3-01>, . §3-01(a): “audited balance sheets as of the end of each of the two most recent fiscal years”; §3-02(a): “audited statements of comprehensive income and cash flows for each of the three fiscal years preceding the date of the most recent audited balance sheet”. Read 2026-09-17.
- [29] U.S. Securities and Exchange Commission. Interactive data to improve financial reporting, release nos. 33-9002, 34-59324. <https://www.sec.gov/rules/final/2009/33-9002.pdf>, 2009. Final rule, January 30, 2009, p. 8: interactive data is intended to facilitate “the comparison of

financial and business performance across companies, reporting periods, and industries.” Read 2026-09-18.

- [30] U.S. Securities and Exchange Commission. EDGAR filer manual, volume II: EDGAR filing, version 78. <https://www.sec.gov/files/edgar/filermanual/edgarfm-vol2-v78.pdf>, 2026. September 14, 2026. Appendix A: “Those notices are reported as warnings or errors. Errors cause a submission to be suspended, but warnings do not.”.
- [31] U.S. Securities and Exchange Commission, staff. EDGAR XBRL guide. <https://www.sec.gov/files/edgar/filer-information/specifications/xbrl-guide.pdf>, 2026. August 2026. §5.9.1 on when calculation relationships are required; §9.10: “Calculation inconsistencies have no impact on the validity of EDGAR submissions.”.
- [32] Yan Wang, Lingfei Qian, Xueqing Peng, Yang Ren, Keyi Wang, Yi Han, Dongji Feng, Fengran Mo, Shengyuan Lin, Qinchuan Zhang, Kaiwen He, Chenri Luo, Jianxing Chen, Junwei Wu, Chen Xu, Ziyang Xu, Jimin Huang, Guojun Xiong, Xiao-Yang Liu, Qianqian Xie, and Jian-Yun Nie. FinTagging: Benchmarking LLMs for extracting and structuring financial information. arXiv:2505.20650, 2025.
- [33] XBRL International. Extensible business reporting language (XBRL) 2.1. <https://www.xbrl.org/Specification/XBRL-2.1/REC-2003-12-31/XBRL-2.1-REC-2003-12-31+corrected-errata-2013-02-20.html>, 2013. Recommendation of December 31, 2003, with corrected errata to February 20, 2013. §1.4 defines an abstract element as one that “cannot be used in an XML instance”; §5.2.4.2, Example 51: an abstract concept “exists in the taxonomy only to organise other elements; it MUST NOT appear in an XBRL Instance.” §4 defines the instance: items, each with a context (entity and period) and a unit. §5.2.5.2: “Summation-item arcs MUST represent relationships only between Concepts,” and a contributing item counts only if it “is C-Equal and U-Equal to” the summation item, that is, in the same context. Read 2026-09-17.