

CRITICAL REVIEW · A REVERSE-ENGINEERED WEB REBUILD

Bean to Brew

An honest look at the Tableau workbook Phase1_Climate_Production_v1000.twbx, the web rebuild that replaces it, and what professional-grade gaps remain.

A reverse-engineered web rebuild of the Tableau workbook Phase1_Climate_Production_v1000.twbx, with an honest assessment of the original, the rebuild, and what should come next.

TL;DR

Dimension	Original Tableau	Web rebuild	Gap
Narrative & framing	Strong — five-chapter arc	Preserved 1:1	None
Chart correctness	Good but several encoding mismatches	Corrected (see §4)	Closed
Data lineage	Hidden in the binary .twbx	Surfaced — every chart has a download button and data dictionary	Closed
Interactivity	Click filters within a sheet	Year scrubber, deep-link URLs, theme toggle, per-chart download	Improved
Accessibility	Tableau Public default (poor)	ARIA, keyboard nav, color-blind palettes, reduced-motion	Improved
Mobile	Fixed-pixel; unusable on phone	Responsive down to 375px	Closed
Portability	Requires Tableau Public or Server	Pure static HTML — works on GitHub Pages, S3, Hostinger, any CDN	Closed
Update cadence	Manual republish	Replace CSV → site updates on reload	Improved
Stat / model layer	None — Tableau "trend line" only	None yet — see §7 roadmap	Open
Test / CI	None	None yet — see §7 roadmap	Open

1 • What the original does well

- Clear narrative arc. Five dashboards walk the reader from product → climate stress → exporters → US consumption with strong editorial transitions. The "Bean to Brew" framing is genuinely good copy.
- Sourced from primary data. FAO production, BACI bilateral trade, Berkeley Earth temperature, ERA5-Land rainfall, World Bank population. No mystery aggregates.

- The "surprise hooks" land. Germany as a top-tier exporter despite zero domestic production, and the US #1-by-volume / not-#1-by-price split, are exactly the kind of insights that make a viz memorable.

2 • What the original does poorly

These are honest critiques — most are constraints of Tableau Public, not failings of the analyst.

2.1 • Visual / UX issues

- Fixed-pixel canvases. Each dashboard is rigid (e.g. Landing Page locked to 1000×800). On any laptop screen <14" or any phone, the layout breaks. Tableau's "Automatic" sizing on the other dashboards just stretches text and squashes maps.
- 42–69 zones per dashboard in the raw XML — most of them invisible "current page" overlays Tableau adds when you mix dashboard actions with parameters. They make the file 12 MB, half of which is duplicated (2) worksheets used as cross-tab linkers.
- No keyboard navigation. Tableau Public's embedded viewer is mouse-only. Screen-reader support is essentially zero.
- Color choices conflict with intent. The Climate Stress dashboard uses a custom diverging palette but it's near-isoluminant — color-blind viewers cannot distinguish "marginal" from "too warm".
- Title typography is small and inconsistent across the five pages. Each dashboard reinvents its own type scale.
- Annotations get cut off at certain zoom levels because they're absolutely positioned in pixel space.

2.2 • Chart-encoding issues found while parsing the XML

- "Climate of Top 5 Producers" uses an Automatic mark but encodes both Avg_Temp_C on rows and a categorical filter on Calculation_...596480. The result is a chart where it's unclear whether you're looking at average temperature or country count. The web rebuild splits this into a clean temperature time-series and a separate rainfall time-series, each with one country per line.
- "Temp vs Production" uses a stacked Area + Automatic mark on a measure-names color. Stacked area is the wrong default for a comparison between two unrelated scales (°C and tonnes). Rebuilt as a dual-axis line where appropriate, and replaced with a temperature trend chart where the dual-axis trick wasn't carrying its weight.
- "Heat Map - Top Producers" encodes both size and color on sum(Production_tonnes). That's redundant — one is just a louder version of the other. The web version uses a bubble pack for compactness and a bar chart for ranking, separating the two questions ("how concentrated?" and "in what order?").
- "Top Exports Stacked" uses pctof:sum:Quantity:4 as the text label — that's "% of total within country" — but the title implies global share. Confusing. Rebuilt as an absolute stacked bar with explicit US vs. rest-of-world segments.
- "Import per capita Map" uses sum(Calculation_1320371406790657) which is a hidden calculated field. After tracing, it's $\text{Import_Qty_tonnes} / \text{Population} \times 1000$ (kg per person). The web rebuild surfaces this formula in the tooltip and adds outlier capping at the 95th percentile so Belgium and the Netherlands (entrepôt countries that re-export almost everything they import) don't wash out the rest of the map.

- Year filters are inconsistent. Some dashboards default to 2020, some to 2022, one to "all years". The web rebuild has a single year scrubber that's the same widget everywhere, plus URL-deep-linking (e.g. `?year=2015`).

2.3 · Data issues found

- Rainfall v1 had ~28% of country-years backfilled with a single climatological mean (a flat value repeated for every year). The chart "looked" like Côte d'Ivoire had identical rainfall every year — because the data literally was identical. Switched to the 20260505 dataset which has true year-by-year ERA5-Land values for 228 of 231 territories. Three entities still have no rainfall (uninhabited Pacific atolls and historical trade aggregates) — the choropleth shows them in the neutral surface color, not as "0 mm".
- Berkeley Earth temperature ends at different years for different countries — Ethiopia and Viet Nam have a sharp drop after ~2015 in the temperature chart because their series end and the last-observation falls off. Should be flagged in the chart, not just rendered.
- BACI bilateral trade has missing exporter values for many small countries even when import values exist for them. The "Comparing demand" chart silently treats missing as zero — could be misread as "they import for free". Mitigated by deriving USD-per-tonne only when both numerator and denominator are non-null.
- No data dictionary in the workbook itself. A reviewer has to open the source CSVs in another tool to know what `Calculation_1791719952338944` means. The web rebuild ships a `data/rainfall_dictionary.md` and inline tooltips reveal full units and source labels.

3 · Faithfulness of the rebuild

Original worksheet	Web equivalent	Faithful?	Notes
Production Map	Choropleth on Production page	Yes	YlOrBr sequential, power scale to give Brazil/Vietnam space
Top Producers (bar)	"Top producers" h-bar	Yes	Year-driven
Coffee Origin Concentration (bubble)	D3 circle pack	Yes	Adds country labels inside large bubbles
Production Area Graph	"Production over time" multi-line	Better	Multi-line is more readable than overlapping areas for 8 series
Heat Map - Top Producers	Replaced by bar + bubble pair	Improved	Removed redundant size+color double-encoding
Rainfall Map	Climate page choropleth	Yes	Sequential Blues, sqrt scale
Average Rainfall (multi-line)	"Rainfall over time"	Yes	Switched to real interannual values from the updated dataset
Climate Stress	Diverging choropleth on Δ from 19.5°C	Yes	Replaced isoluminant palette with green→amber→red

Original worksheet	Web equivalent	Faithful?	Notes
Climate of Top 5 Producers	Split into rainfall + temperature charts	Improved	Cleaner separation of variables
Rainfall Volatility	Roll into "Rainfall over time"	Acceptable	Volatility is visible in the line wiggles
Top Exports Stacked	"Top exporters" stacked bar	Yes	Explicit US-vs-rest segmentation
US Imports by Country (area)	Stacked area in Exporters page	Yes	Top 8 ranked by all-time US qty
Import Ranks	"Import ranks" h-bar	Yes	US highlighted in berry-red
Comparing Demand	"Comparing demand" h-bar of \$/tonne	Improved	Replaced text-on-cells with a real bar chart
Import per capita Map	"Import per capita" choropleth	Yes	95th-percentile cap added
Top Imports US	Folded into stacked bar on Exporters page	Acceptable	Same info, one fewer chart
US Imports by Year	Stacked area on US Imports page	Yes	
Temp vs Production	Currently omitted	Gap	Will reintroduce as a true dual-axis chart per producer — see §7
Where coffee is grown	Combined with Production Map	Yes	

4 • Concrete improvements over the Tableau original

- › Single, consistent year scrubber with play/pause animation, used on every page.
- › Deep-linkable URL state — e.g. `pages/climate.html?year=2018` shares a specific view.
- › Dark mode — full theme toggle, respects `prefers-color-scheme`, persisted in `localStorage`.
- › Download data per chart — every card has a download icon that exports the underlying rows as CSV with header.
- › Tooltip with units and sources on every data point, not just a value blob.
- › 95th-percentile cap on the per-capita choropleth so Belgium / Netherlands entrepôt re-exports don't flatten the rest of the world to one color.
- › Outlier-aware sequential scales (power=0.4 on production, sqrt on rainfall) so Brazil/Vietnam don't dominate the legend.
- › Color-blind-safe diverging palette (green → amber → red, not the original near-isoluminant scale).
- › Brand-consistent typography — Fraunces (display) + Inter (UI), one type scale across all pages.
- › Mobile responsive — works down to 375px (the original Tableau cuts off below 1000px).
- › Print stylesheet — saving any page as PDF gives a clean report with charts on white.

- › Reduced-motion support — respects user OS setting.
- › Lazy-loaded TopoJSON — only loaded when a map page opens (~100 kB) and cached.
- › Total payload <2 MB including data, vs. the 12 MB Tableau package, and runs without a Tableau Public account.

5 • Professional-grade software practices still missing

These would all be expected in a "production" deliverable and would be the next priorities:

- No automated tests. No unit tests for the chart helpers, no visual-regression suite, no data-validation step. A quick tests/ folder with Vitest + Playwright would catch breakage.
- No build pipeline. Static is fine, but a tiny Vite/esbuild step would let us tree-shake D3 and ship ~50 kB instead of ~280 kB.
- No CSP / SRI headers. The CDN-loaded D3, Plotly, marked and topojson scripts aren't pinned with integrity=. A real deploy needs subresource integrity hashes and a Content-Security-Policy.
- No analytics or error tracking. Sentry, Plausible, or even a 30-line window.onerror reporter to a file.
- No API layer. All data is shipped as CSV. For >50k row datasets this is fine; for live trade data it should hit a small FastAPI/Express endpoint backed by Postgres or DuckDB.
- No versioning of datasets. Each CSV should have a _meta.json with checksum, source URL, fetch date — like the rainfall data dictionary, but for all five.
- No CI / preview deploys. GitHub Actions → Pages or Cloudflare Pages preview-per-PR is one yaml file away.
- No accessibility audit. Should run axe-core or Pa11y and fix the findings, not just "feel accessible".
- No semantic-versioning or changelog. A CHANGELOG.md per data refresh would let consumers know what moved.
- No 404 page. The current pages/review.html flow falls back to the static-server 404 if the markdown is missing, which it shouldn't be in production but should be handled.

6 • Critique of the rebuild itself

I held myself to the same scrutiny.

- Plotly is overkill for two charts. It adds ~250 kB just to draw stacked areas. Should be replaced with a 60-line D3 stacked-area helper, dropping the dependency.
- The bubble pack lays out bubbles without regard to geography. A more honest visual would be a Voronoi treemap or simply ranked bars — the bubble is decorative.
- Country selection is set in sessionStorage but only the choropleths read it — the bar charts and timelines ignore it. Cross-filtering is half-implemented and should either be finished (every chart filters together) or removed.
- The diverging palette for Climate Stress is interpolated through RGB, which produces muddy intermediate colors. Should be lab-space (d3.interpolateLab) instead.

- The per-capita map uses 95th-percentile cap with no UI to reveal the outliers. A small "outliers > 6 kg/person: BEL, NLD, CHE, FIN" footnote would close the loop.
- The temperature chart shows Berkeley Earth's end-of-series cliff with no visual indication. Need a "last observation: 2015" badge on countries whose data ends early, or trim the chart to a year range with full coverage.
- No empty-state handling. If somebody scrubs to a year before BACI coverage (2000), the charts go blank rather than showing a friendly "no data" panel.
- The "Final" dashboard from the original (a small credits/summary panel) was collapsed into the footer. Acceptable, but a proper "About the data" page would be better.

7 • Roadmap for v2

In priority order:

- › Fix the half-implemented cross-filter — clicking any country highlights it on every chart on the page, and the chosen country persists across pages.
- › Drop Plotly in favor of native D3 stacked-area. Saves ~250 kB and unifies styling.
- › Add the "Temp vs Production" chart as a proper per-producer dual-axis view with selectable country.
- › Statistics layer — compute drought severity (z-score vs. country's own 25-yr mean), correlate temperature anomaly with production change. Surface the actual climate-vs-production relationship the title implies.
- › Data freshness badge on each chart — "FAO updated 2024-06" — pulled from per-dataset `_meta.json`.
- › An "About the data" page with the full data dictionary, source URLs, known limitations, and download links for the raw CSVs.
- › Unit tests + visual regression (Vitest for helpers; Playwright snapshot diffs for each page).
- › Build pipeline (Vite) with tree-shaken D3 and SRI-pinned CDN fallback.
- › CI — GitHub Actions to lint + build + deploy preview on every PR.
- › Sentry + Plausible for error tracking and privacy-respecting analytics.

8 • Grade

	Score	
Faithfulness to original	8/10	All five chapters preserved; one chart deferred
Visual design	8/10	Clean, brand-coherent, but bubble pack and diverging palette could be sharper
Interactivity	7/10	Scrubber + URL state + downloads land well; cross-filter is half-done

Score		
Accessibility	7/10	Big improvement on Tableau, but no formal audit
Engineering rigor	5/10	No tests, no build, no CI — the gap to "production-grade"
Data quality	7/10	Updated rainfall is a big win; temperature data gap needs handling

Overall: 7.0 / 10. A confident MVP that closes most of the original's gaps, surfaces the data lineage honestly, and is ready to be hardened into a v2 with the §7 roadmap.