

Toy Data Global Supply Chain

by Quinn Dyer

Origin Selector (Hold Ctrl/Cmd to select multiple regions and product lines)



Destination Selector



Classic Cars



Motorcycles



Planes



Ships



Trains



Trucks and Buses



Vintage Cars

\$10,032,629

Total Sales

99,067


Total Units Sold

1.38

Supplier Risk Score




Mix of Product Lines




Non-Compliance and Concentration

Metric and Target
(Bubbles and Timeline)
● Sales ● Suppliers
○ Quantity ○ Sales Offices



Timeline Month



With my Toy Data Global Supply Chain dashboard, I wanted to tackle a broader scope rather than just one area. My dashboard connects customers, suppliers, and sales offices by tracking the complete lifecycle of sales, mapping exactly where products are sourced, where they are being shipped, and how effectively internal teams are managing and fulfilling those orders. The dashboard is built around a dark-themed, compartmentalized design, separating different metrics and visualizations into cleanly bordered cards. The layout is interactive, with many filters near the top and several points for selection, all while still being intuitive. Changing a variable in any filter section, like selecting a set of product lines, instantly updates the data displayed across the rest of the board using background Tableau actions.

At the top of the layout, the Origin and Destination Selectors filter the data based on the geographic locations of orders. Users can click on flat country silhouettes to filter by supplier locations (origin) and customer locations (destination). Selected countries are highlighted in purple. Below this is the Product Line filter, featuring stylized icons that represent the different product categories. Toggling these icons filters the data by product type. Hovering over any icon brings up a tooltip showing relevant stats for each location or line, such as sales within and across the filtered lines or countries, along with the calculated average sale price and MSRP (not the faulty given stat capped at \$100) for the specific product lines within the selected regions.

The impact of these filters is summarized in a row of four Key Performance Indicator (KPI) cards. These cards display total sales, total units sold, a pie chart showing the mix of product lines, and a supplier risk score. The dynamically colored risk score represents the average number of missing compliance certifications across the selected subset of suppliers, weighted by quantity. All KPIs include tooltips explaining their function.

The centerpiece of the dashboard is a spider map that visualizes the global flow of the filtered goods. Semi-translucent arcs connect origin points to destination points, with the thickness of each arc indicating the volume of units sold along that specific trade route. Hovering directly over an arc reveals a detailed tooltip listing the origin and destination countries, alongside the top salespeople and customers for that route by total sales, color-coded by product line. Hovering over the specific origin or destination dots provides detailed information about the supplier or customer and their orders.

To the right of the map is a dynamic bubble chart whose subject matter changes entirely based on a parameter selector labeled “Metric and Target” in the upper right corner. When set to

“Suppliers,” the chart is “Non-Compliance and Concentration” and evaluates supplier dependence and reliability. Each bubble represents a single supplier. Its size reflects that supplier’s share of the selected metric (either revenue or units), while the color, ranging from green to red, indicates how many of the four possible compliance certifications that supplier is missing—mirroring the supplier risk KPI for the total selected lines and regions. Hovering over a bubble expands a tooltip showing the supplier’s name, exact missing compliance count, specific failed and missing certifications, percentage of the total metric, markets sorted and sized by sales, and top product lines sorted and sized by sales. This view lets users identify where suppliers’ non-compliance with standards may introduce risk into the supply chain.

When the parameter is switched to “Sales Offices,” the chart changes to “Unresolved Orders and Concentration,” and the two bubbles represent the company’s two sales offices, New York and San Francisco. Here, bubble size still reflects each office’s share of the selected metric, but the color now represents the office’s unresolved order rate—the proportion of orders that are currently cancelled, disputed, on hold, or in process rather than resolved. Hovering over an office bubble reveals a tooltip breaking down those unresolved order counts by status category. It also shows the office's average sales per representative, average discount percentage, exact missing compliance count, specific failed certifications, percentage of the total metric, markets sorted and sized by sales, and top product lines sorted and sized by sales. This view lets users identify whether order fulfillment problems are concentrated in one office or spread evenly across both.

Finally, a stacked area timeline spans the bottom of the dashboard, mapping historical data on revenue or sales quantity to illustrate product line trends at different time granularities, which can be adjusted via the dropdown to the right of the title. This chart also responds to the parameter selector for suppliers and sales offices. When “Suppliers” is active, the color bands reflect the suppliers’ overall current non-compliance. When “Sales Offices” is active, the bands represent the overall current unresolved order percentage. Hovering over any point on the timeline brings up a tooltip showing the relevant time period (based on the dropdown), total sales or quantity, the year-over-year difference in that metric, and a breakdown of unresolved orders for that time period on the sales office graph. Hovering over any bubble in the bubble chart highlights its corresponding supplier or sales office on the timeline, and vice versa.

I think I overcame most limitations when constructing this dashboard, but a few remain. It is difficult, though not impossible, to select smaller routes that are the inverse of larger ones on

the spider graph. It is also tricky to select a country to view its stats as a destination when it is also an origin (relevant for Germany and the United States). In addition, due to the number of features, the dashboard is often slow to load when changing filters or selectors. Because both sales offices frequently have similar unresolved rates for certain filter selections, their colors end up matching; this makes them harder to distinguish on the timeline, although it is an accurate reflection of the data. To create the spider map and define the specific dots and paths for the trade routes, I added two additional CSV files (Supplier_GEO.csv and Customer_GEO.csv) containing coordinates for the nodes of the spider map, as I couldn't figure out a way to accomplish this without those files; hopefully, this won't be an issue when the professor scores my assignment. Finally, due to time constraints, I wasn't able to build out all the features I wanted, such as exploring discount rates across salespeople, customers, and suppliers, or identifying suspicious data—like a salesperson selling items at a 300% markup. I utilized AI to help brainstorm ideas for the dashboard and to assist with some of the calculated field expressions.

