Edward R. Tufte

The Visual Display of Quantitative Information

SECOND EDITION

Graphics Press · Cheshire, Connecticut

PUBLISHED BY GRAPHICS PRESS LLC
POST OFFICE BOX 430, CHESHIRE, CONNECTICUT 06410
WWW.TUFTE.COM

Printed in the United States of America   *Second edition, sixth printing, August 2009*

2 Graphical Integrity

For many people the first word that comes to mind when they think about statistical charts is "lie." No doubt some graphics do distort the underlying data, making it hard for the viewer to learn the truth. But data graphics are no different from words in this regard, for any means of communication can be used to deceive. There is no reason to believe that graphics are especially vulnerable to exploitation by liars; in fact, most of us have pretty good graphical lie detectors that help us see right through frauds.

Much of twentieth-century thinking about statistical graphics has been preoccupied with the question of how some amateurish chart might fool a naive viewer. Other important issues, such as the use of graphics for serious data analysis, were largely ignored. At the core of the preoccupation with deceptive graphics was the assumption that data graphics were mainly devices for showing the obvious to the ignorant. It is hard to imagine any doctrine more likely to stifle intellectual progress in a field. The assumption led down two fruitless paths in the graphically barren years from 1930 to 1970: First, that graphics had to be "alive," "communicatively dynamic," overdecorated and exaggerated (otherwise all the dullards in the audience would fall asleep in the face of those boring statistics). Second, that the main task of graphical analysis was to detect and denounce deception (the dullards could not protect themselves).

Then, in the late 1960s, John Tukey made statistical graphics respectable, putting an end to the view that graphics were only for decorating a few numbers. For here was a world-class data analyst spinning off half a dozen new designs and, more importantly, using them effectively to explore complex data.[1] Not a word about deception; no tortured attempts to construct more "graphical standards" in a hopeless effort to end all distortions. Instead, graphics were used as instruments for reasoning about quantitative information. With this good example, graphical work has come to flourish.

Of course false graphics are still with us. Deception must always be confronted and demolished, even if lie detection is no longer at the forefront of research. Graphical excellence begins with telling the truth about the data.

[1] John W. Tukey and Martin B. Wilk, "Data Analysis and Statistics: Techniques and Approaches," in Edward R. Tufte, ed., *The Quantitative Analysis of Social Problems* (Reading, Mass., 1970), 370–390; and John W. Tukey, "Some Graphic and Semigraphic Displays," in T. A. Bancroft, ed., *Statistical Papers in Honor of George W. Snedecor* (Ames, Iowa, 1972), 293–316.

Here are several graphics that fail to tell the truth. First, the case of the disappearing baseline in the annual report of a company that would just as soon forget about 1970. A careful look at the middle panel reveals a negative income in 1970, which is disguised by having the bars begin at the bottom at approximately minus $4,200,000:

Day Mines, Inc., *1974 Annual Report*, 1.

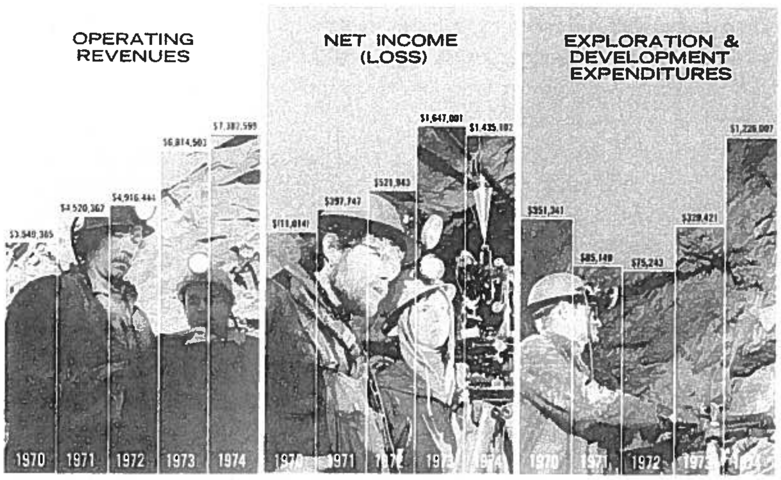

This pseudo-decline was created by comparing six months' worth of payments in 1978 to a full year's worth in 1976 and 1977, with the lie repeated four times over.

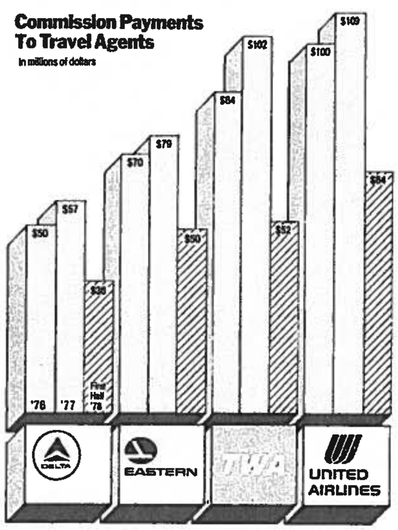

New York Times, August 8, 1978, D-1.

And sometimes the fact that numbers have a magnitude as well as an order is simply forgotten:

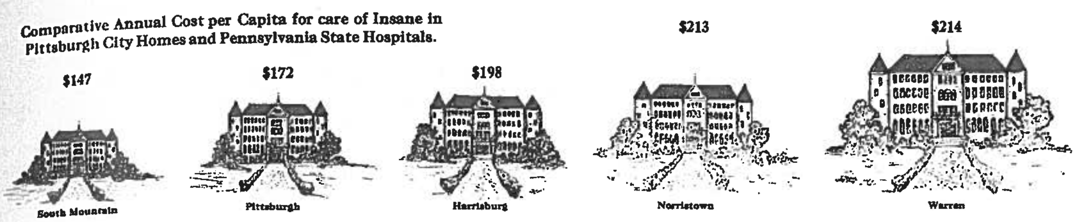

Pittsburgh Civic Commission, *Report on Expenditures of the Department of Charities* (Pittsburgh, 1911), 7.

What is Distortion in a Data Graphic?

A graphic does not distort if the visual representation of the data is consistent with the numerical representation. What then is the "visual representation" of the data? As physically measured on the surface of the graphic? Or the *perceived* visual effect? How do we know that the visual image represents the underlying numbers?

One way to try to answer these questions is to conduct experiments on the visual perception of graphics—having people look at lines of varying length, circles of different areas, and then recording their assessments of the numerical quantities.

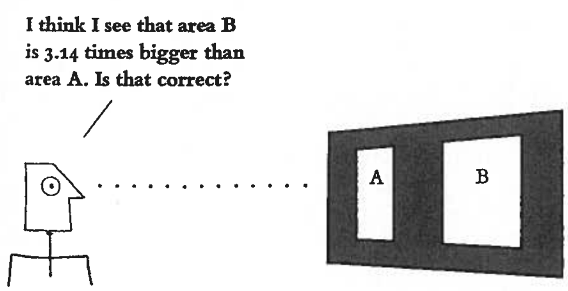

Such experiments have discovered very approximate power laws relating the numerical measure to the reported perceived measure. For example, the perceived area of a circle probably grows somewhat more slowly than the actual (physical, measured) area: the reported perceived area = (actual area)^x, where $x = .8 \pm .3$, a discouraging result. Different people see the same areas somewhat differently; perceptions change with experience; and perceptions are context-dependent.[2] Particularly disheartening is the securely established finding that the reported perception of something as clear and simple as line length depends on the context and what other people have already said about the lines.[3]

Misperception and miscommunication are certainly not special to statistical graphics,

but what is a poor designer to do? A different graphic for each perceiver in each context? Or designs that correct for the visual transformations of the average perceiver participating in the average psychological experiment?

One satisfactory answer to these questions is to use a table to show the numbers. Tables usually outperform graphics in reporting on small data sets of 20 numbers or less. The special power of graphics comes in the display of large data sets.

At any rate, given the perceptual difficulties, the best we can hope for is some uniformity in graphics (if not in the perceivers) and some assurance that perceivers have a fair chance of getting the numbers right. Two principles lead toward these goals and, in consequence, enhance graphical integrity:

> The representation of numbers, as physically measured on the surface of the graphic itself, should be directly proportional to the numerical quantities represented.
>
> Clear, detailed, and thorough labeling should be used to defeat graphical distortion and ambiguity. Write out explanations of the data on the graphic itself. Label important events in the data.

[2] The extensive literature is summarized in Michael Macdonald-Ross, "How Numbers are Shown: A Review of Research on the Presentation of Quantitative Data in Texts," *Audio-Visual Communication Review*, 25 (1977), 359–409. In particular, H. J. Meihoefer finds great variability among perceivers; see "The Utility of the Circle as an Effective Cartographic Symbol," *Canadian Cartographer*, 6 (1969), 105–117; and "The Visual Perception of the Circle in Thematic Maps: Experimental Results," ibid., 10 (1973), 63–84.

[3] S. E. Asch, "Studies of Independence and Submission to Group Pressure. A Minority of One Against a Unanimous Majority," *Psychological Monographs* (1956), 70.

Drawing by CEM; copyright 1961, *The New Yorker*.

Violations of the first principle constitute one form of graphic misrepresentation, measured by the

$$\text{Lie Factor} = \frac{\text{size of effect shown in graphic}}{\text{size of effect in data}}$$

If the Lie Factor is equal to one, then the graphic might be doing a reasonable job of accurately representing the underlying numbers. Lie Factors greater than 1.05 or less than .95 indicate substantial distortion, far beyond minor inaccuracies in plotting. The logarithm of the Lie Factor can be taken in order to compare overstating (log LF > 0) with understating (log LF < 0) errors. In practice almost all distortions involve overstating, and Lie Factors of two to five are not uncommon.

Here is an extreme example. A newspaper reported that the U.S. Congress and the Department of Transportation had set a series of fuel economy standards to be met by automobile manufacturers, beginning with 18 miles per gallon in 1978 and moving in steps up to 27.5 by 1985, an increase of 53 percent:

$$\frac{27.5 - 18.0}{18.0} \times 100 = 53\%$$

These standards and the dates for their attainment were shown:

This line, representing 18 miles per gallon in 1978, is 0.6 inches long.

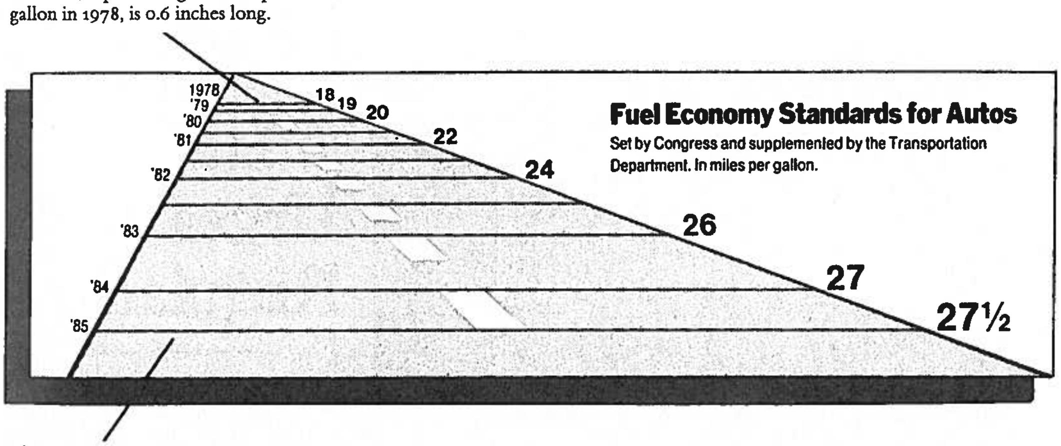

This line, representing 27.5 miles per gallon in 1985, is 5.3 inches long.

New York Times, August 9, 1978, D-2.

The magnitude of the change from 1978 to 1985 is shown in the graph by the relative lengths of the two lines:

$$\frac{5.3 - 0.6}{0.6} \times 100 = 783\%$$

Thus the numerical change of 53 percent is presented by some lines that changed 783 percent, yielding

$$\text{Lie Factor} = \frac{783}{53} = 14.8$$

which is too big.

The display also has several peculiarities of perspective:

- On most roads the future is in front of us, toward the horizon, and the present is at our feet. This display reverses the convention so as to exaggerate the severity of the mileage standards.
- Oddly enough, the dates on the left remain a constant size on the page even as they move along with the road toward the horizon.
- The numbers on the right, as well as the width of the road itself, are shrinking because of two simultaneous effects: the change in the values portrayed and the change due to perspective. Viewers have no chance of separating the two.

It is easy enough to decorate these data without lying:

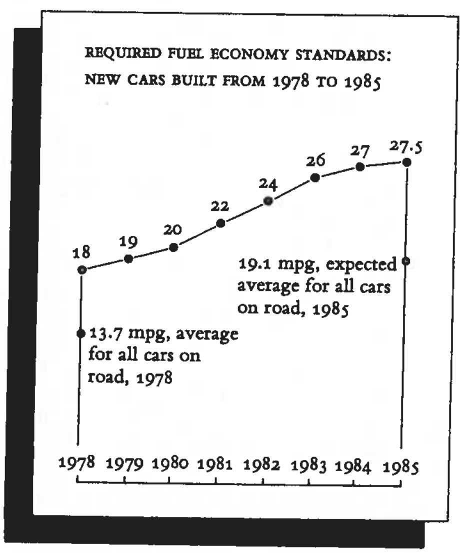

The non-lying version, in addition, puts the data in a context by comparing the new car standards with the mileage achieved by the mix of cars actually on the road. Also revealed is a side of the data disguised and mispresented in the original display: the fuel economy standards require gradual improvement at start-up, followed by a doubled rate from 1980 to 1983, and flattening out after that.

Sometimes decoration can help editorialize about the substance of the graphic. But it is wrong to distort the data measures—the ink locating values of numbers—in order to make an editorial comment or fit a decorative scheme. It is also a sure sign of the Graphical Hack at work. Here are many decorations but no lies:

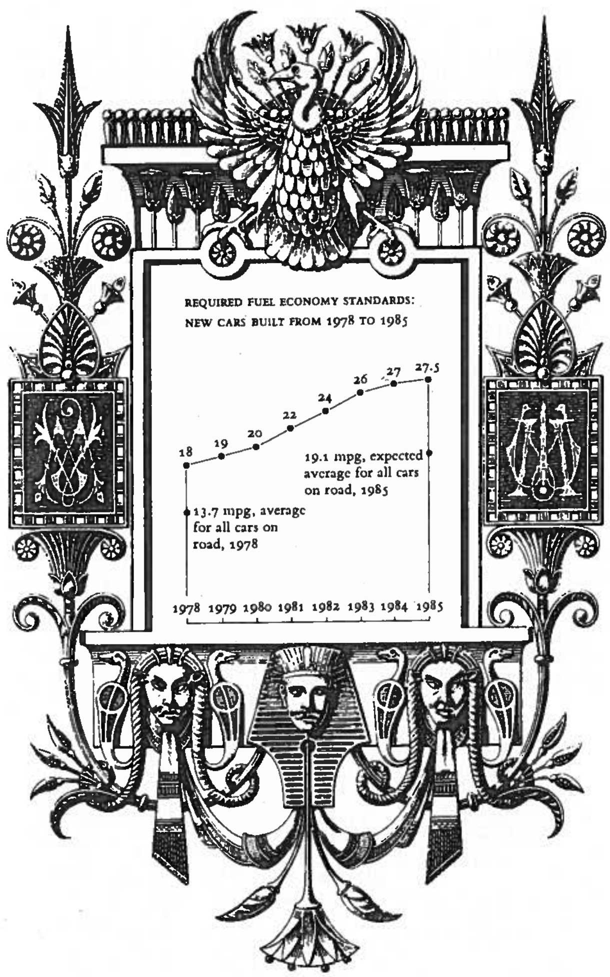

Design and Data Variation

Each part of a graphic generates visual expectations about its other parts and, in the economy of graphical perception, these expectations often determine what the eye sees. Deception results from the incorrect extrapolation of visual expectations generated at one place on the graphic to other places.

A scale moving in regular intervals, for example, is expected to continue its march to the very end in a consistent fashion, without the muddling or trickery of non-uniform changes. Here an irregular scale is used to concoct a pseudo-decline. The first seven increments on the horizontal scale are ten years long, masking the rightmost interval of four years. Consequently the conspicuous feature of the graphic is the apparent fall of curves at the right, particularly the decline in prizes won by people from the United States (the heavy, dark line) in the most recent period. This effect results solely from design variation. It is a big lie, since in reality (and even in extrapolation, scaling up each end-point by 2.5 to take the four years' worth of data up to a comparable decade), the U.S. curve turned sharply upward in the post-1970 interval. A correction, with the actual data for 1971–80, is at the right:

National Science Foundation, *Science Indicators, 1974* (Washington, DC, 1976), 15.

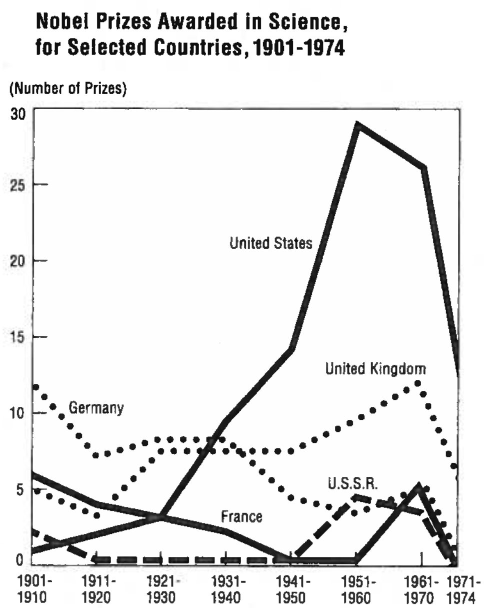

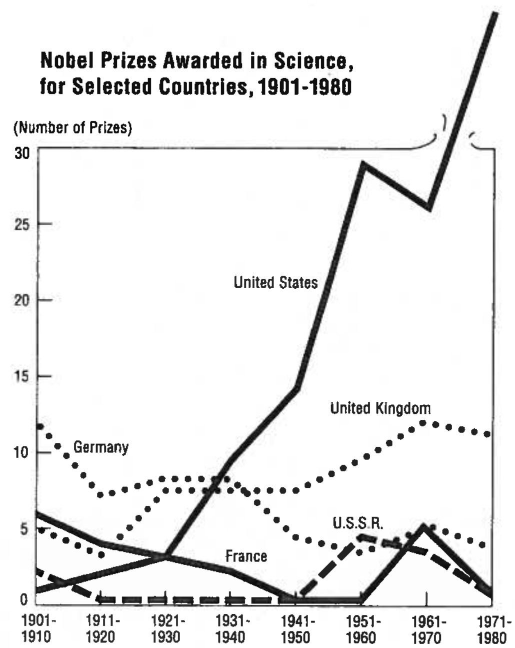

The confounding of *design variation* with *data variation* over the surface of a graphic leads to ambiguity and deception, for the eye may mix up changes in the design with changes in the data. A steady canvas makes for a clearer picture. The principle is, then:

Show data variation, not design variation.

Design variation corrupts this display:

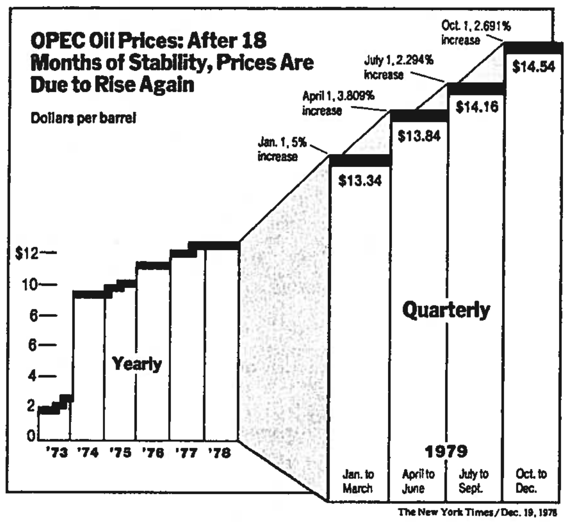

New York Times, December 19, 1978, D-7.

Five different vertical scales show the price:

During this time	one vertical inch equals
1973–1978	\$8.00
January–March 1979	\$4.73
April–June 1979	\$4.37
July–September 1979	\$4.16
October–December 1979	\$3.92

And two different horizontal scales show the passage of time:

During this time	one horizontal inch equals
1973–1978	3.8 years
1979	0.57 years

As the two scales shift simultaneously, the distortion takes on multiplicative force. On the left of the graph, a price of \$10 for one year is represented by 0.31 square inches; on the right side, by 4.69 square inches. Thus exactly the same quantity is 4.69/0.31 = 15.1 times larger depending upon where it happens to fall on the surface of the graphic. *That* is design variation.

Design variation infected similar graphics in other publications. Here an increase of 454 percent is depicted as an increase of 4,280 percent, for a Lie Factor of 9.4:

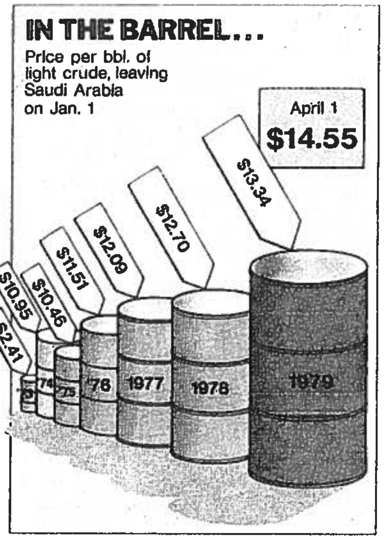

Time, April 9, 1979, 57.

And an increase of 708 percent is shown as 6,700 percent, for a Lie Factor of 9.5:

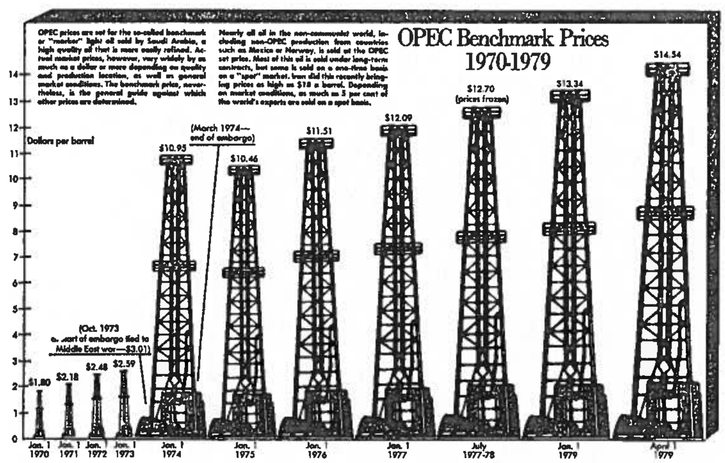

Washington Post, March 28, 1979, A-18.

All these accounts of oil prices made a second error, by showing the price of oil in inflated (current) dollars. The 1972 dollar was worth much more than the 1979 dollar. Thus in sweeping from left to right over the surface of the graphic, the vertical scale in effect changes—design variation—because the value of money changes over the years shown. The only way to think clearly about money over time is to make comparisons using inflation-adjusted units of money. Several distinguished graphic designers did express the price in real dollars—and they also avoided other sources of design variation. The stars were *Business Week*, the *Sunday Times* (London), and *The Economist*.

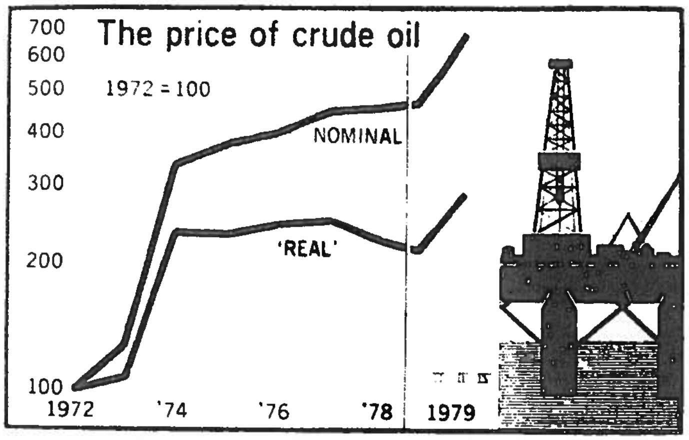

In the graphic we saw first, the two sources of design variation covered up an intriguing, non-obvious aspect of the data: in the four years prior to the 1979-1980 increases, the real price of oil had *declined*. Busy with decoration, the graphic had missed the news.

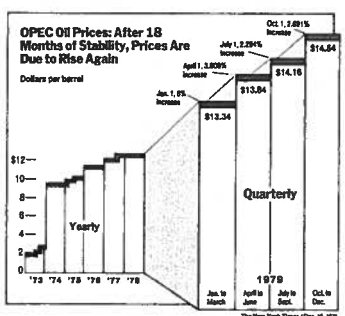

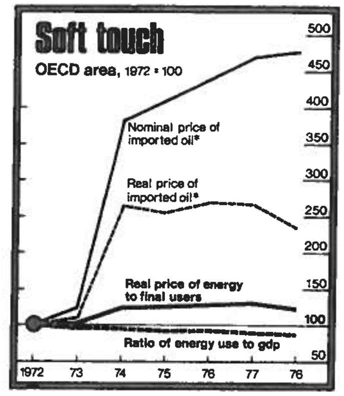

The Economist, December 29, 1979, 41.

Sunday Times (London), December 16, 1979, 54.

Business Week, April 9, 1979, 99.

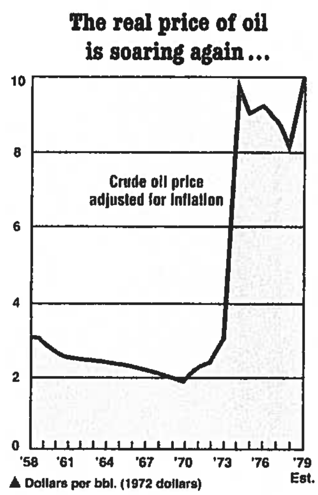

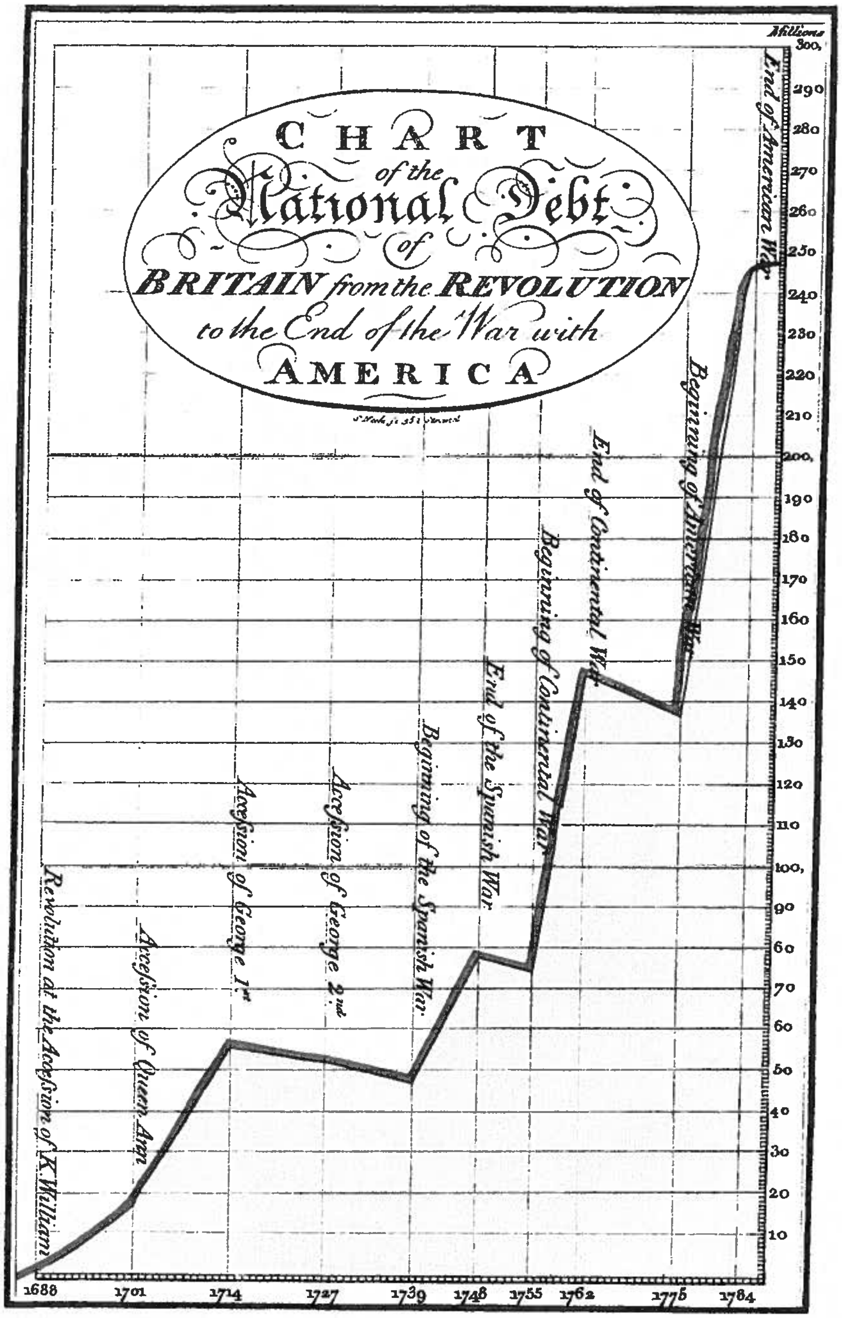

The Divisions at the Bottom are Years, & those on the Right hand Money.

The Case of Skyrocketing Government Spending

Probably the most frequently printed graphic, other than the daily weather map and stock-market trend line, is the display of government spending and debt over the years. These arrays nearly always create the impression that spending and debt are rapidly increasing.

As usual, Playfair was the first, publishing this finely designed graphic in 1786. Accompanied by his polemic against the "ruinous folly" of the British government policy of financing its colonial wars through debt, it is surely the first skyrocketing government debt chart, beginning the now 200-year history of such displays. This is one of the few Playfairs that is taller than wide; less than one-tenth of all his graphics (about 90, drawn during 35 years of work) are longer on the vertical. The tall shape here serves to emphasize the picture of rapid growth. The money figures are not adjusted for inflation.

But Playfair had the integrity to show an alternative version a few pages later in *The Commercial and Political Atlas*. The interest on the national debt was plotted on a broad horizontal scale, diminishing the skyrocket effect. And, furthermore, "This is in real and not in nominal millions" (page 129):

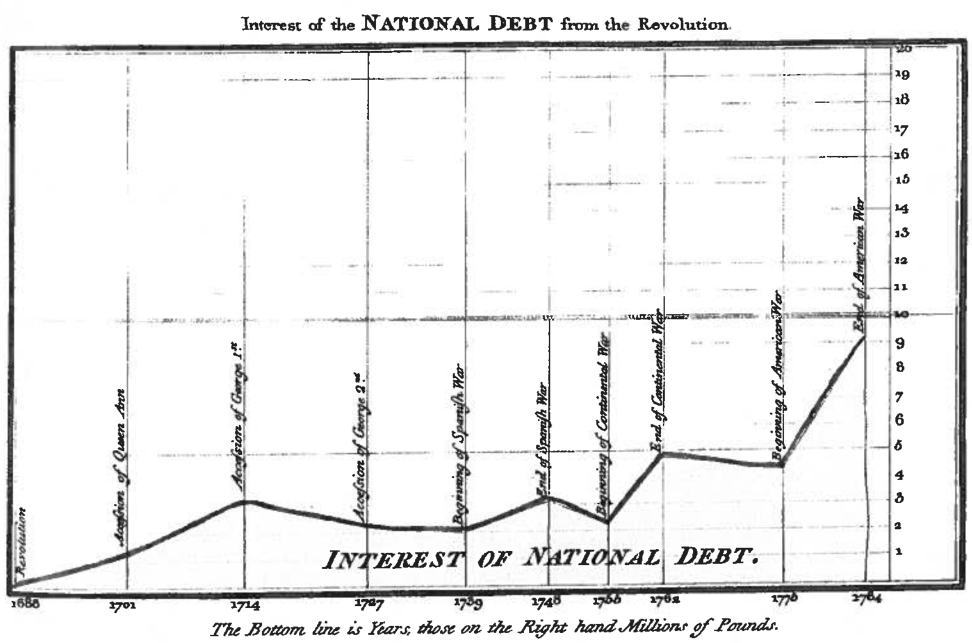

Although Playfair deflated money units over time in his work of 1786, the matter has proved difficult for many, eluding even modern scholars. This display helps its political point along by failing to discount for inflation and population growth and by using a tall and thin shape (the area covered by the data is 2.7 times taller than wide):

Morris Fiorina, *Congress: Keystone of the Washington Establishment* (New Haven, 1977), 92.

Figure A3. The Growth of Government: Federal Spending in Selected Domestic Areas

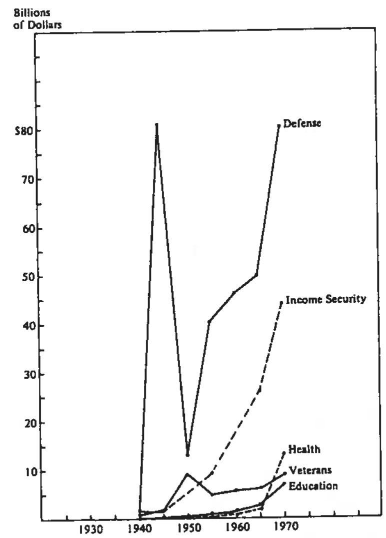

Let us look, in detail, at another graphic on government spending:

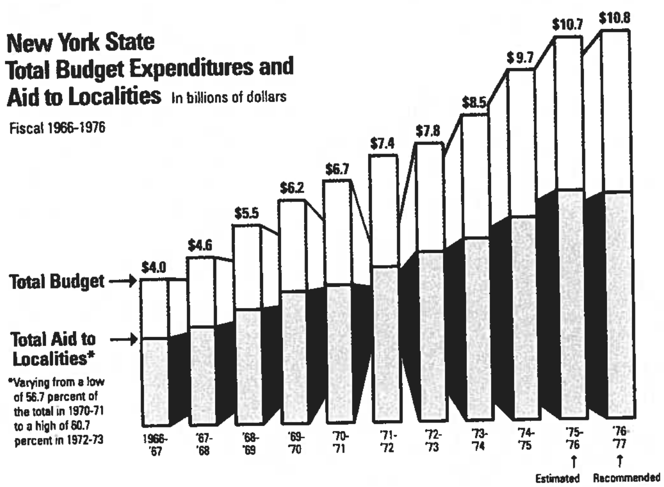

New York Times, February 1, 1976, IV-6.

Despite the appearance created by the hyperactive design, the state budget actually did not increase during the last nine years shown. To generate the thoroughly false impression of a substantial and continuous increase in spending, the chart deploys several visual and statistical tricks—all working in the same direction, to exaggerate the growth in the budget. These graphical gimmicks:

These three parallelepipeds have been placed on an optical plane *in front* of the other eight, creating the image that the newer budgets tower over the older ones.

This cluster of type emphasizes and stretches out the low value for 1966–1967, encouraging the impression that recent years have shot up from a small, stable base. Horizontal arrows provide similar emphasis.

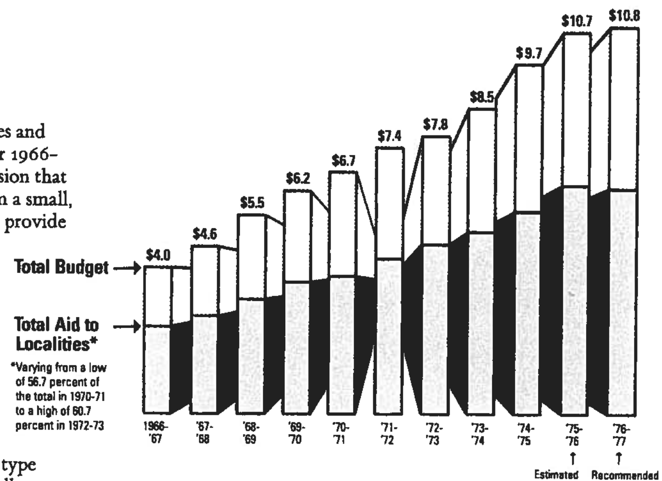

This squeezed-down block of type contributes to an image of small, squeezed-down budgets back in the good old days.

Arrows pointing straight up emphasize recent growth. Compare with horizontal arrows at left.

Leaving behind the distortion in the chartjunk heap at the left yields a calmer view:

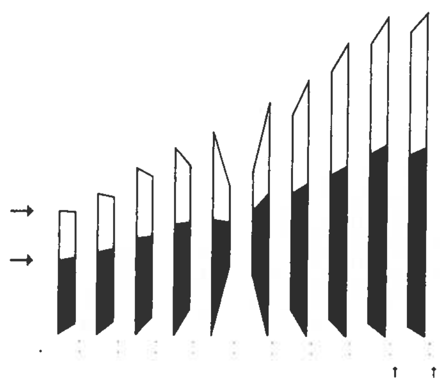

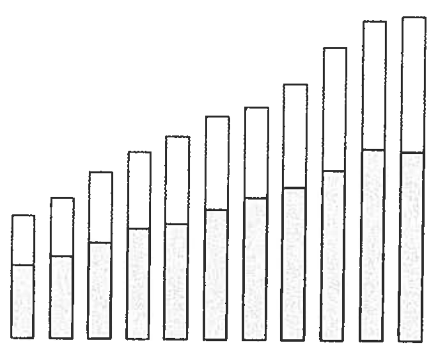

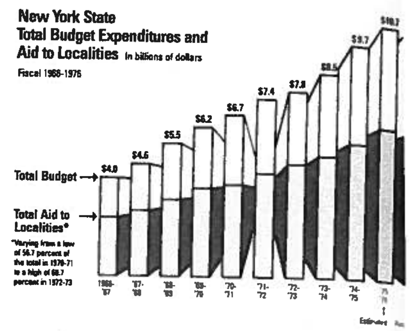

Two statistical lapses also bias the chart. First, during the years shown, the state's population increased by 1.7 million people, or 10 percent. Part of the budget growth simply paralleled population growth. Second, the period was a time of substantial inflation; those goods and services that cost state and local governments $1.00 to purchase in 1967 cost $2.03 in 1977. By not deflating, the graphic mixes up changes in the value of money with changes in the budget.

Application of arithmetic makes it possible to take population and inflation into account. Computing expenditures in *constant (real) dollars per capita* reveals a quite different—and far more accurate—picture:

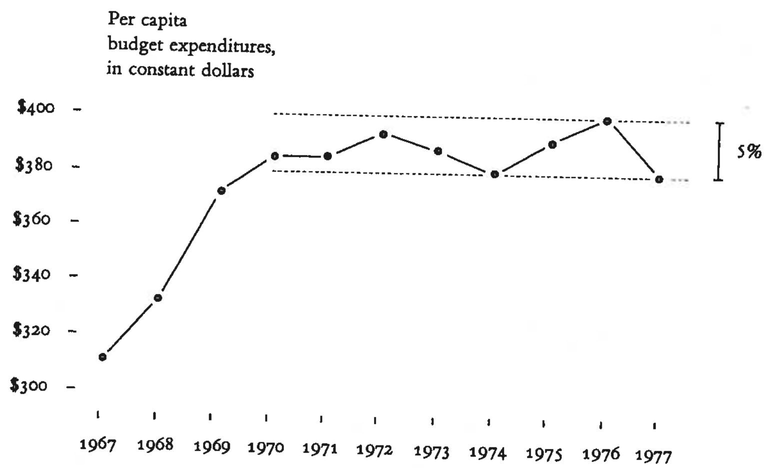

Thus, in terms of real spending per capita, the state budget increased by about 20 percent from 1967 to 1970 and remained relatively constant from 1970 through 1976. And the 1977 budget represents a substantial *decline* in expenditures. That is the real news story of these data, and it was completely missed by the Graph of the Magical Parallelepipeds. Of course no small set of numbers is going to capture the complexities of a large budget—but, at any rate, why tell lies?

The principle:

> In time-series displays of money, deflated and standardized units of monetary measurement are nearly always better than nominal units.

Visual Area and Numerical Measure

Another way to confuse data variation with design variation is to use areas to show one-dimensional data:

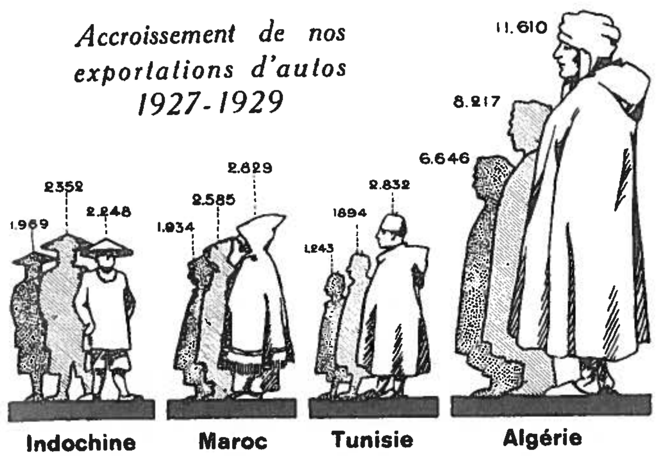

R. Satet, *Les Graphiques* (Paris, 1932), 12.

And here is the incredible shrinking doctor, with a Lie Factor of 2.8, not counting the additional exaggeration from the overlaid perspective and the incorrect horizontal spacing of the data:

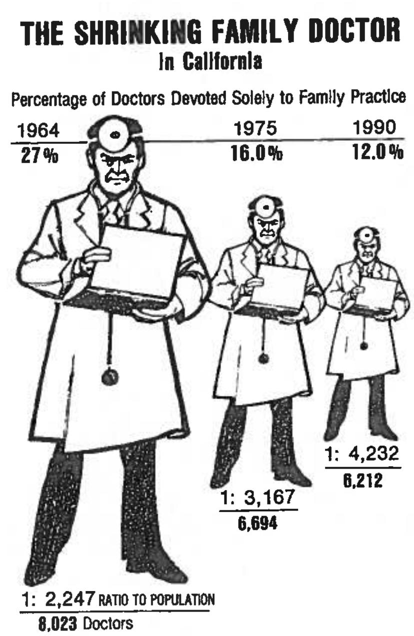

Los Angeles Times, August 5, 1979, 3.

Many published efforts using areas to show magnitudes make the elementary mistake of varying both dimensions simultaneously in response to changes in one-dimensional data. Typical is the shrinking dollar fallacy. To depict the rate of inflation, graphs show currency shrinking on two dimensions, even though the value of money is one-dimensional. Here is one of hundreds of such charts:

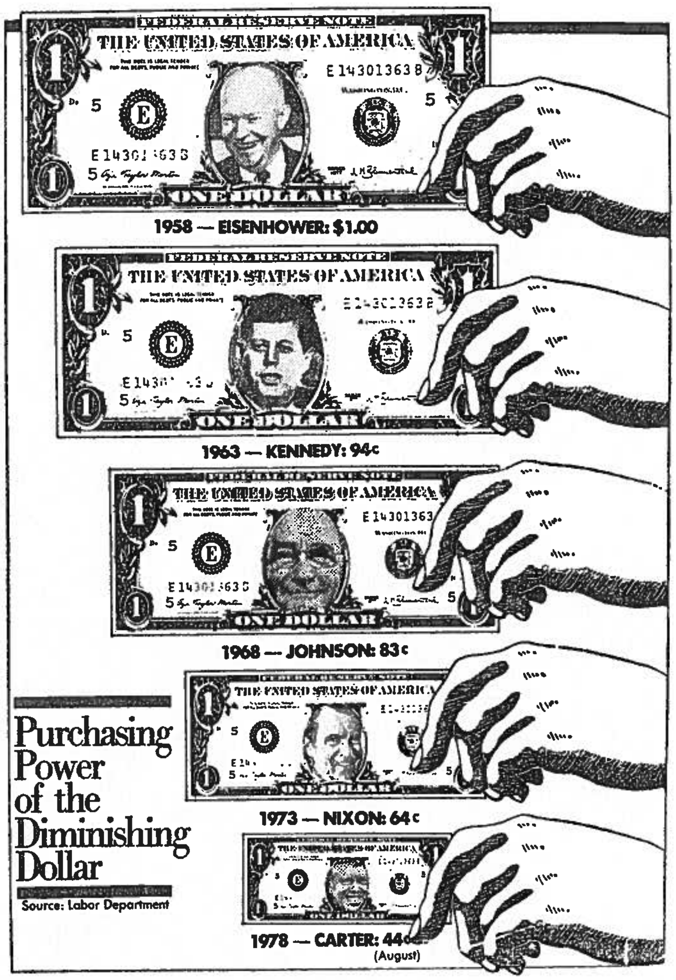

Washington Post, October 25, 1978, 1.

If the area of the dollar is accurately to reflect its purchasing power, then the 1978 dollar should be about twice as big as that shown.

There are considerable ambiguities in how people perceive a two-dimensional surface and then convert that perception into a one-dimensional number. Changes in physical area on the surface of a graphic do not reliably produce appropriately proportional changes in perceived areas. The problem is all the worse when the areas are tricked up into three dimensions:

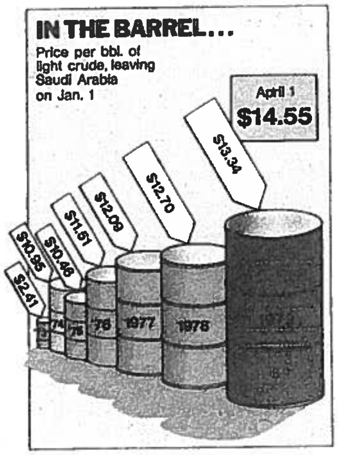

By surface area, the Lie Factor for this graphic is 9.4. But, if one takes the barrel metaphor seriously and assumes that the *volume* of the barrels represents the price change, then the volume from 1973 to 1979 increases 27,000 percent compared to a data increase of 454 percent, for a Lie Factor of 59.4, which is a record.

Similarly, a three-dimensional representation puffing up one-dimensional data:

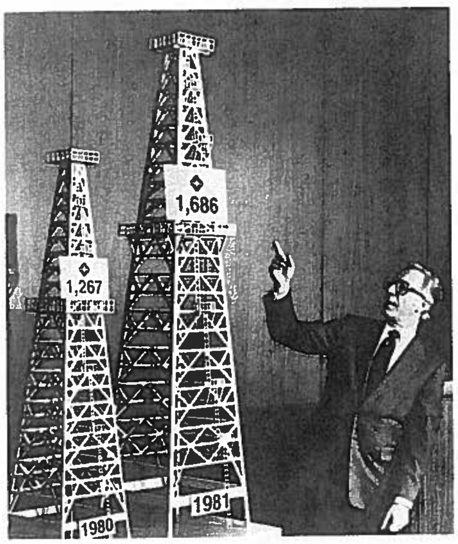

New York Times, January 27, 1981, D-1.

Conclusion: The use of two (or three) varying dimensions to show one-dimensional data is a weak and inefficient technique, capable of handling only very small data sets, often with error in design and ambiguity in perception. These designs cause so many problems that they should be avoided:

> The number of information-carrying (variable) dimensions depicted should not exceed the number of dimensions in the data.

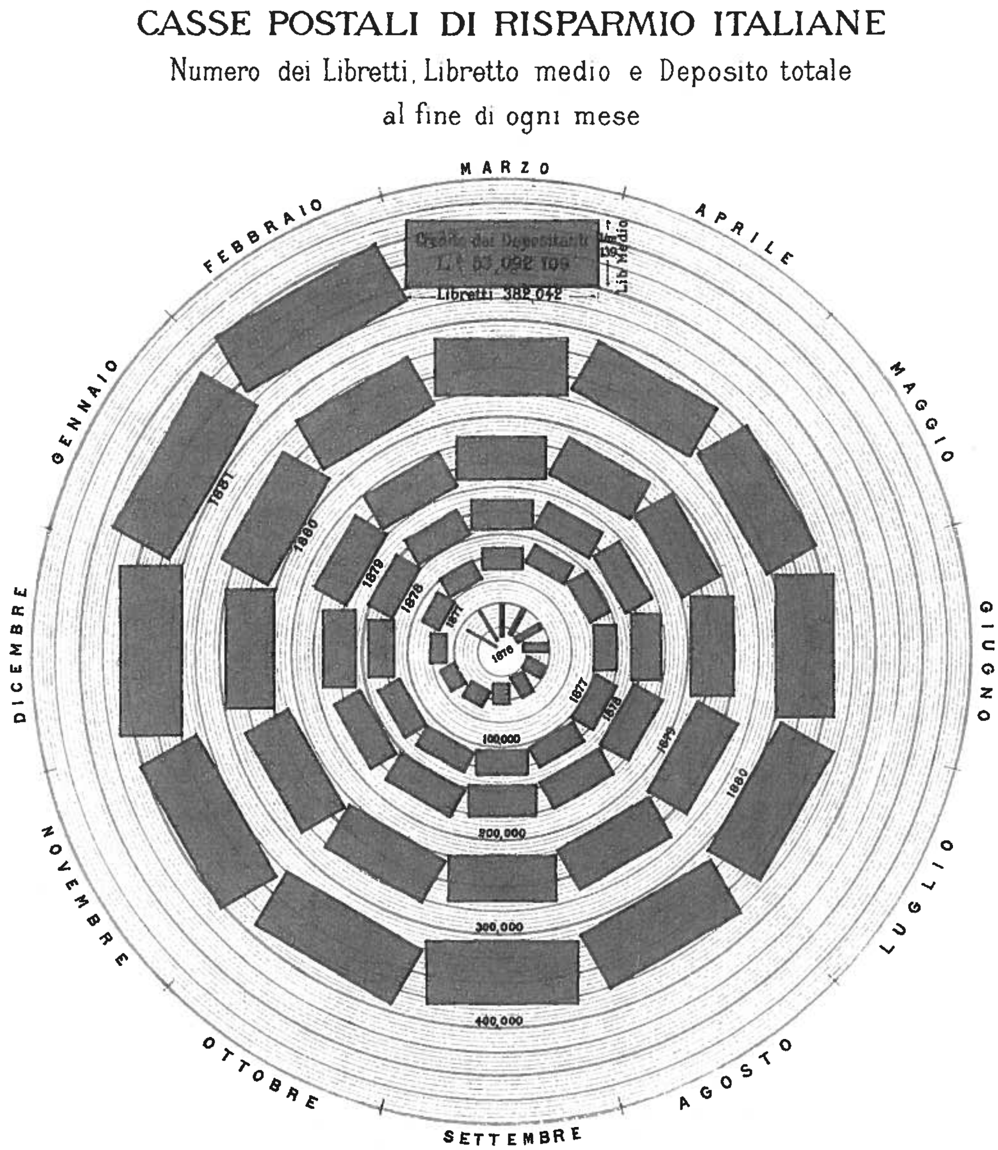

This multivariate history of the Italian post office uses two dimensions in a way nearly consistent with this principle, with the number of postal savings books issued and the average size of deposits multiplying up to total deposits at the end of each month from 1876 to 1881.

Antonio Gabaglio, *Teoria Generale della Statistica* (Milan, second edition, 1888).

But Playfair's circles, an early use of area to show magnitude, are not consistent with the principle, since the one-dimensional data (city populations) are represented by area:

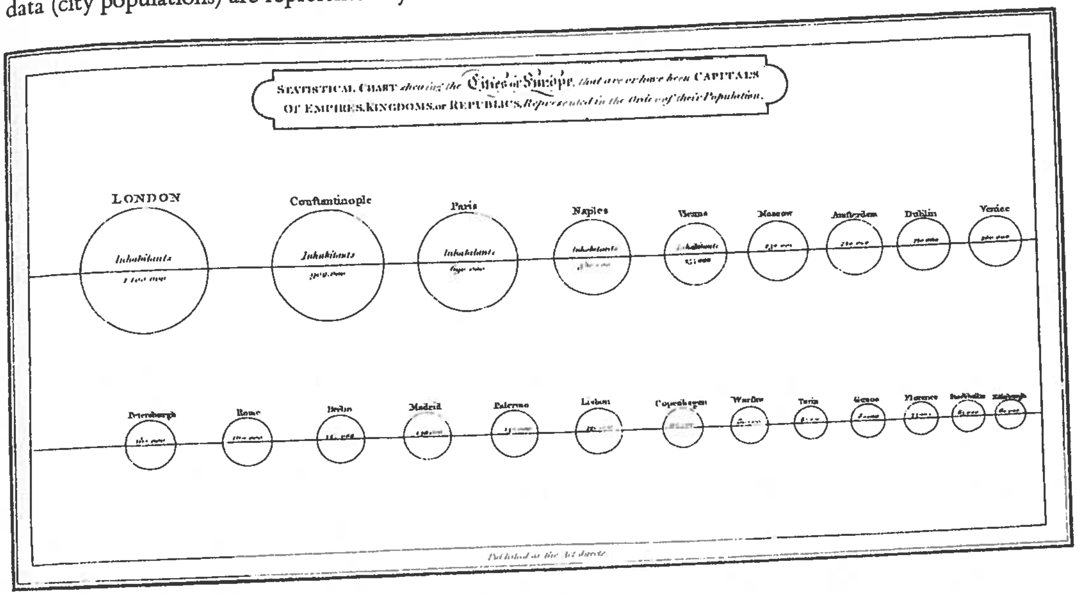

Perhaps graphics that border on cartoons should be exempt from the principle. We certainly would not want to forgo the 4,340 pound chicken:

Scientific American Reference Book (New York, 1909), 280.

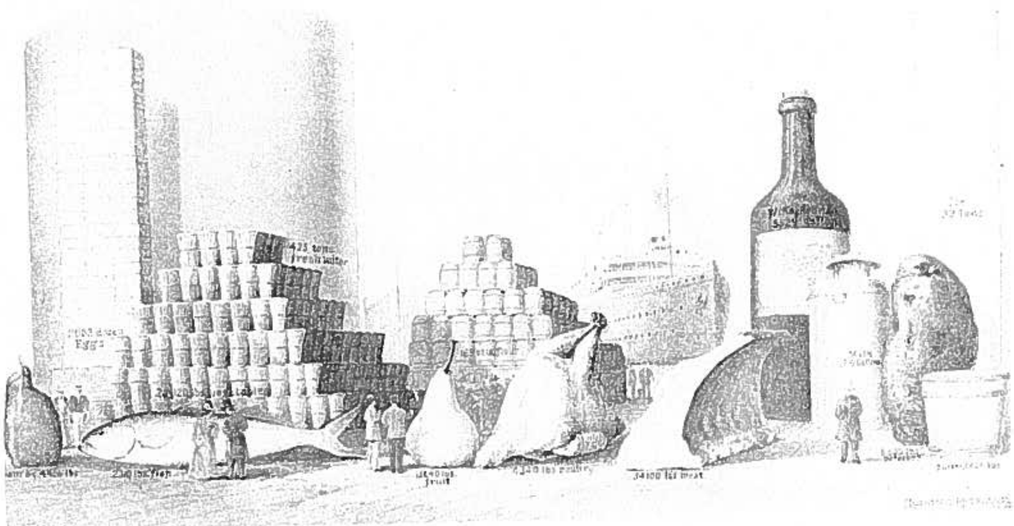

Context is Essential for Graphical Integrity

To be truthful and revealing, data graphics must bear on the question at the heart of quantitative thinking: "Compared to what?" The emaciated, data-thin design should always provoke suspicion, for graphics often lie by omission, leaving out data sufficient for comparisons. The principle:

Graphics must not quote data out of context.

Nearly all the important questions are left unanswered by this display:

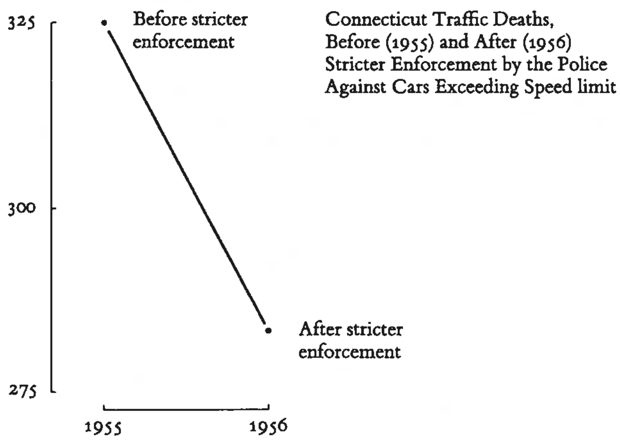

A few more data points add immensely to the account:

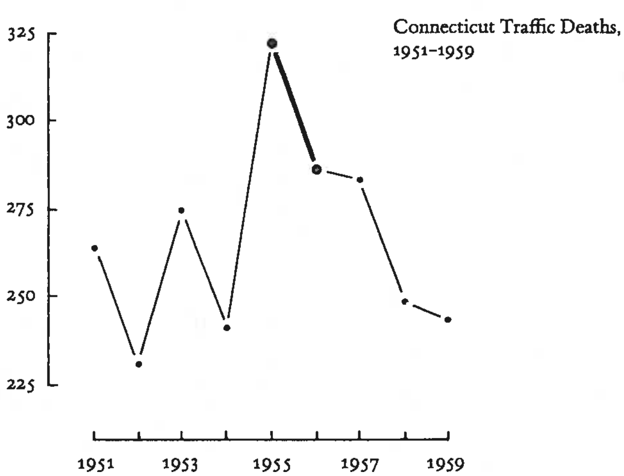

Imagine the very different interpretations other possible time-paths surrounding the 1955–1956 change would have:

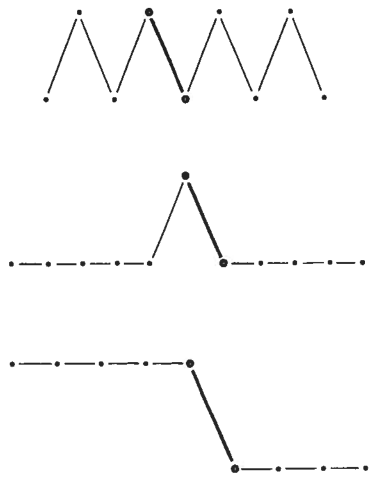

Comparisons with adjacent states give a still better context, revealing it was not only Connecticut that enjoyed a decline in traffic fatalities in the year of the crackdown on speeding:

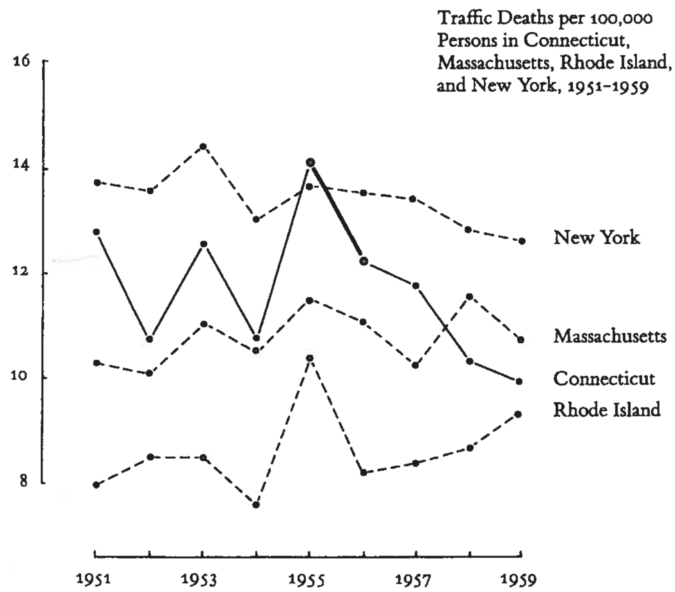

Donald T. Campbell and H. Laurence Ross, "The Connecticut Crackdown on Speeding: Time Series Data in Quasi-Experimental Analysis," in Edward R. Tufte, ed., *The Quantitative Analysis of Social Problems* (Reading, Massachusetts, 1970), 110–125.

Conclusion

Lying graphics cheapen the graphical art everywhere. Since the lies often show up in news reports, millions of images are printed. When a chart on television lies, it lies tens of millions of times over; when a *New York Times* chart lies, it lies 900,000 times over to a great many important and influential readers. The lies are told about the major issues of public policy—the government budget, medical care, prices, and fuel economy standards, for example. The lies are systematic and quite predictable, nearly always exaggerating the rate of recent change.

The main defense of the lying graphic is . . . "Well, at least it was approximately correct, we were just trying to show the general direction of change." But many of the deceptive displays we saw in this chapter involved fifteenfold lies, too large to be described as approximately correct. And in several cases the graphics were not even approximately correct by the most lax of standards, since the falsified the real news in the data. It is the special character of numbers that they have a magnitude as well as an order; numbers measure *quantity*. Graphics can display the quantitative size of changes as well as their direction. The standard of getting only the direction and not the magnitude right is the philosophy that informs the Pravda School of Ordinal Graphics. There, every chart has a crystal clear direction coupled with fantasy magnitudes.

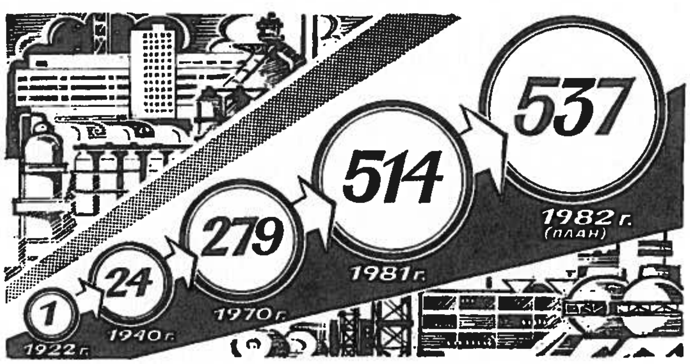

Рост продукции промышленности (1922 г. = 1).

Pravda, May 24, 1982, 2.

A second defense of the lying graphic is that, although the design itself lies, the actual numbers are printed on the graphic for those picky folks who want to know the correct size of the effects displayed. It is as if not lying in one place justified fifteenfold lies elsewhere. Few writers would work under such a modest standard of integrity, and graphic designers should not either.

Graphical integrity is more likely to result if these six principles are followed:

The representation of numbers, as physically measured on the surface of the graphic itself, should be directly proportional to the numerical quantities represented.

Clear, detailed, and thorough labeling should be used to defeat graphical distortion and ambiguity. Write out explanations of the data on the graphic itself. Label important events in the data.

Show data variation, not design variation.

In time-series displays of money, deflated and standardized units of monetary measurement are nearly always better than nominal units.

The number of information-carrying (variable) dimensions depicted should not exceed the number of dimensions in the data.

Graphics must not quote data out of context.

5 *Chartjunk: Vibrations, Grids, and Ducks*

THE interior decoration of graphics generates a lot of ink that does not tell the viewer anything new. The purpose of decoration varies—to make the graphic appear more scientific and precise, to enliven the display, to give the designer an opportunity to exercise artistic skills. Regardless of its cause, it is all non-data-ink or redundant data-ink, and it is often chartjunk. Graphical decoration, which prospers in technical publications as well as in commercial and media graphics, comes cheaper than the hard work required to produce intriguing numbers and secure evidence.

Sometimes the decoration is thought to reflect the artist's fundamental design contribution, capturing the essential spirit of the data and so on. Thus principles of artistic integrity and creativity are invoked to defend—even to advance—the cause of chartjunk. There are better ways to portray spirits and essences than to get them all tangled up with statistical graphics.

Fortunately most chartjunk does not involve artistic considerations. It is simply conventional graphical paraphernalia routinely added to every display that passes by: over-busy grid lines and excess ticks, redundant representations of the simplest data, the debris of computer plotting, and many of the devices generating design variation.

Like weeds, many varieties of chartjunk flourish. Here three widespread types found in scientific and technical research work are catalogued—unintentional optical art, the dreaded grid, and the self-promoting graphical duck. A hundred chartjunky examples from commercial and media graphics have been forgone so as to demonstrate the relevance of the critique to the professional scientific production of data graphics.

Unintentional Optical Art

Contemporary optical art relies on moiré effects, in which the design interacts with the physiological tremor of the eye to produce the distracting appearance of vibration and movement.

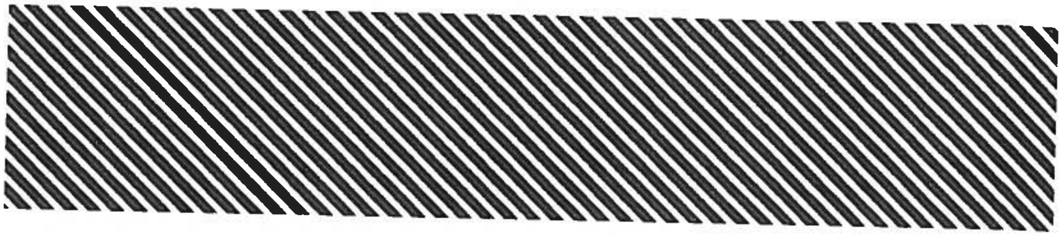

The effect extends beyond the ink of the design to the whole page. When exploited by the experts, such as Bridget Riley and Victor Vasarely, op art effects are undoubtedly eye-catching.

But statistical graphics are also often drawn up so as to shimmer. This moiré vibration, probably the most common form of graphical clutter, is inevitably bad art and bad data graphics. The noise clouds the flow of information as these examples from technical and scientific publications illustrate:

Instituto de Expansão Commercial, *Brasil: Graphicos Economicos-Estatisticas* (Rio de Janeiro, 1929), 15.

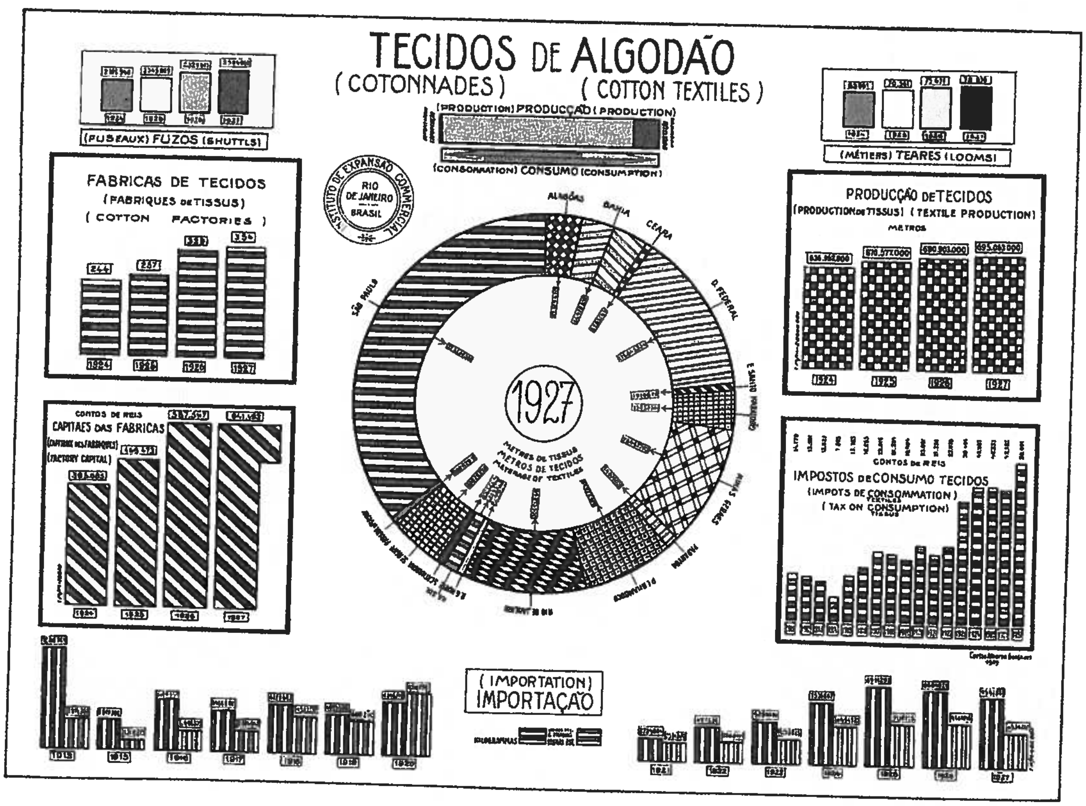

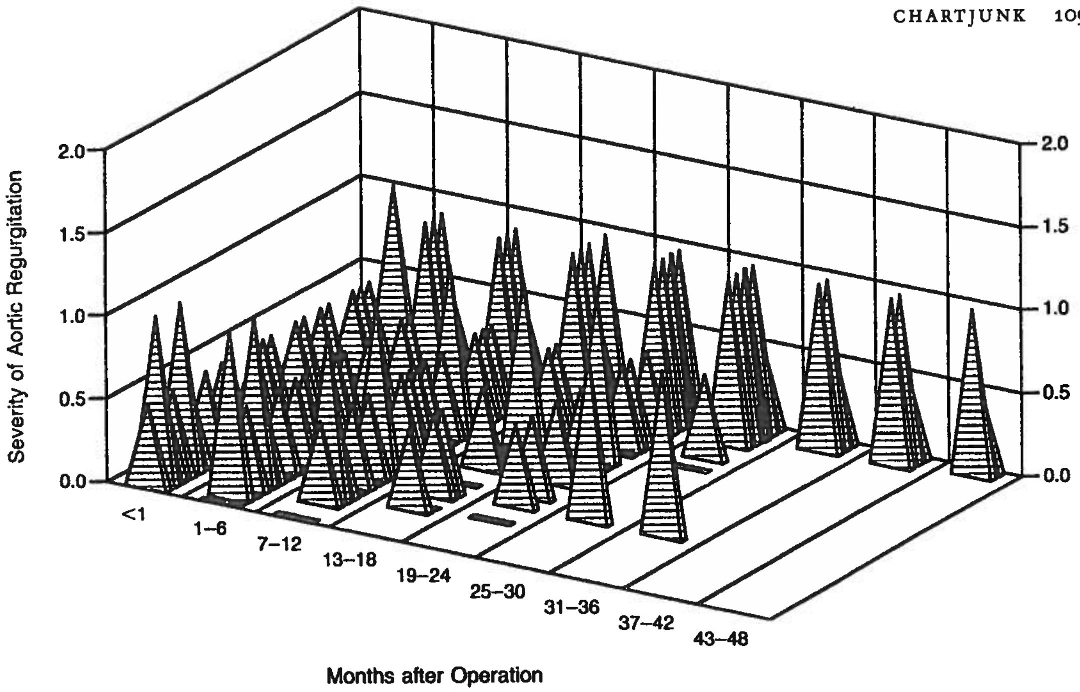

Figure 2. Serial Echocardiographic Assessments of the Severity of Regurgitation in the Pulmonary Autograft in 31 Patients. The numerical grades were assigned according to the severity of regurgitation, as follows: 0, none; 0.5, trivial; 1.0 to 1.5, mild; 2.0, moderate; and 3.0, severe.

On this page, what should have been simple tables are turned into bad graphics published in major scientific journals. Above a duck moiré with an unintentional Necker Illusion, as the two back planes optically flip to the front. Some pyramids conceal others; and one variable (stacked depth of the stupid pyramids) has no label or scale. Below, we learn very little about data, but do discover that moiré vibration may well be at a maximum for equally spaced bars:

Nicholas T. Kouchoukos, *et al.*, "Replacement of the Aortic Root with a Pulmonary Autograft in Children and Young Adults with Aortic-Valve Disease," *The New England Journal of Medicine*, 330 (January 6, 1994), 4.

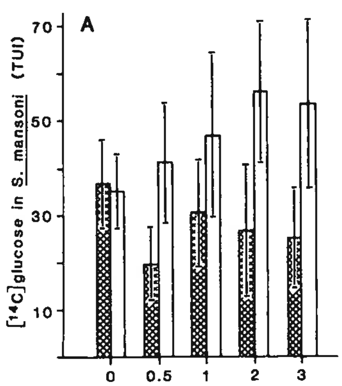

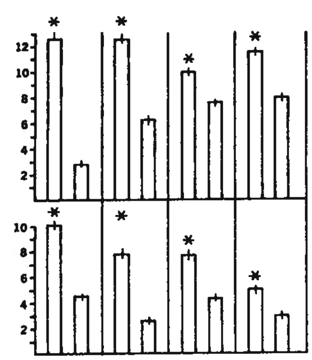

James T. Kuznicki and N. Bruce McCutcheon, "Cross-Enhancement of the Sour Taste on Single Human Taste Papillae," *Journal of Experimental Psychology: General*, 108 (1979), 76.

Eain M. Cornford and Marie E. Huot, "Glucose Transfer from Male to Female Schistosomes," *Science*, 213 (September 11, 1981), 1270.

And, finally, from the style sheet once provided by the *Journal of the American Statistical Association*, a graphic described as "an example of a figure prepared in the proper form":

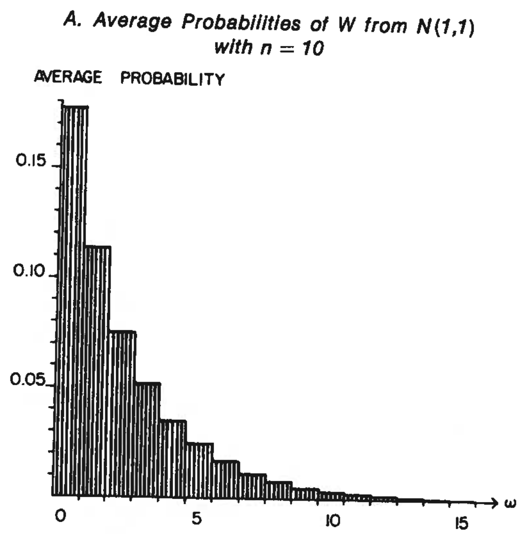

"JASA Style Sheet," *Journal of the American Statistical Association*, 71 (March 1976), 260–261.

The display required 131 line-strokes and 15 digits to communicate its simple information. The vibrating lines are poorly drawn, unevenly spaced, and misaligned with the vertical axis.

Vibrating chartjunk even frequents the graphics of major scientific journals:

The ten most frequently cited (footnoted) scientific journals: random sample of issues published 1980–1982	Percentage of graphics with moiré vibration	Number of graphics in sample
Biochemistry	2%	568
Journal of Biological Chemistry	2%	565
Journal of the American Chemical Society	3%	317
Journal of Chemical Physics	6%	327
Biochimica et Biophysica Acta	8%	432
Nature	11%	225
Proceedings of the National Academy of Sciences, U.S.A.	12%	438
Lancet	15%	364
Science	17%	311
New England Journal of Medicine	21%	338

Moiré effects have proliferated with computer graphics (in programs such as Excel). Such unfortunate patterns were once generated by means of thin plastic transfer sheets; now the computer produces instant chartjunk. Shown here are a few of the many vibrating possibilities. Cross-hatching should be replaced with tint screens of shades of gray. Specific areas on a graphic should be labeled with words rather than encoded with hatching.

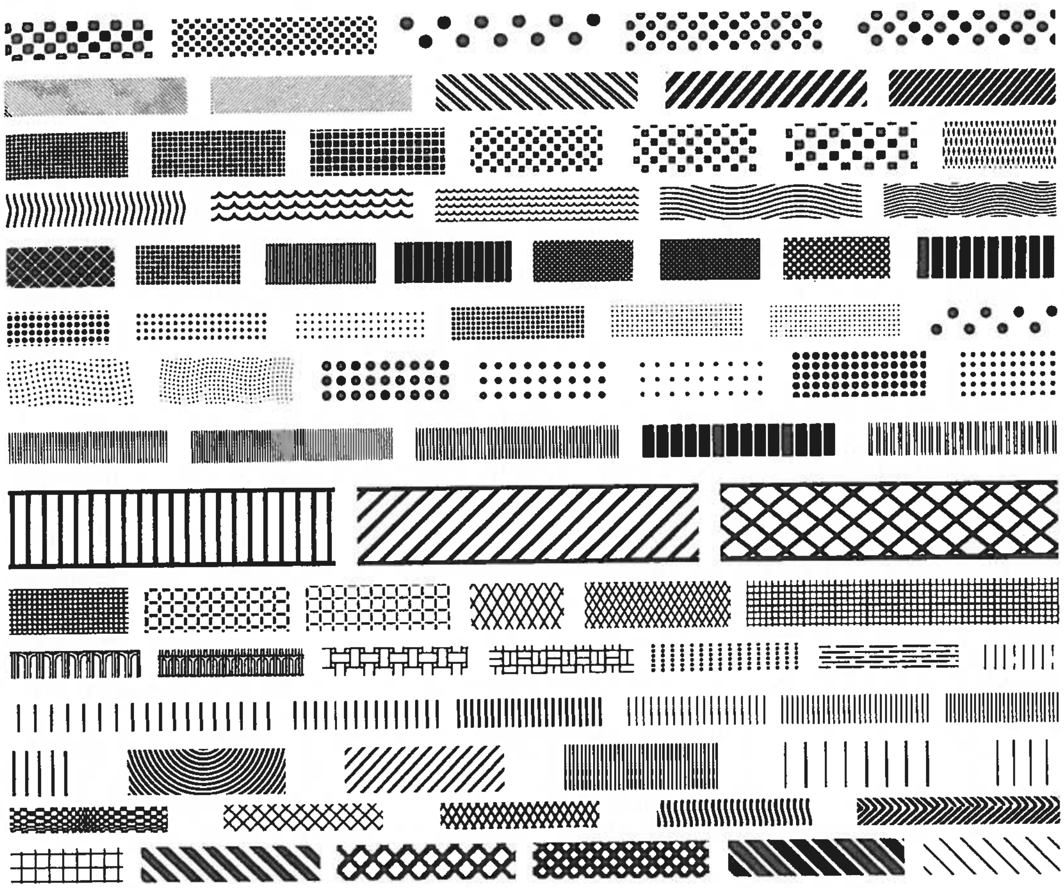

This form of chartjunk is a twentieth-century innovation, and computer graphics are multiplying it more than ever. The handbooks and textbooks of statistical graphics, along with user's manuals for computer graphics programs, are filled up with vibrating graphics, presented as exemplars of design. Note the high proportion of chartjunky graphics in the more recent publications. Computer graphics are particularly active:

Textbooks and handbooks of statistical graphics; and manuals for computer graphics programs (ordered by date of publication)	Percentage of graphics with moiré vibration	Total number of graphics
Willard C. Brinton, *Graphic Methods for Presenting Facts* (New York, 1914)	12%	255
R. Satet, *Les Graphiques* (Paris, 1932)	29%	28
Herbert Arkin and Raymond R. Colton, *Graphs: How to Make and Use Them* (New York, 1936)	17%	95
Mary Eleanor Spear, *Charting Statistics* (New York, 1952)	46%	134
Anna C. Rogers, *Graphic Charts Handbook* (Washington, DC, 1961)	32%	201
F. J. Monkhouse and H. R. Wilkinson, *Maps and Diagrams* (London, third edition, 1971)	14%	322
Calvin F. Schmid and Stanton E. Schmid, *Handbook of Graphic Presentation* (New York, second edition, 1979)	22%	399
A. J. MacGregor, *Graphics Simplified* (Toronto, 1979)	34%	65
The user's manual for a widely distributed computer graphics package: *SAS/GRAPH User's Guide* (Cary, North Carolina, 1980)	68%	28
The manual for a very extensive computer graphics program: *Tell-A-Graf User's Manual* (San Diego, 1981)	53%	459

Can optical art effects ever produce a better graphic? Bertin exhorts: "It is the designer's duty to make the most of this variation; to obtain the resonance [of moiré vibration] without provoking an uncomfortable sensation: to flirt with ambiguity without succumbing to it."[1] But can statistical graphics successfully "flirt with ambiguity"? It is a clever idea, but no good examples are to be found. The key difficulty remains: moiré vibration is an *undisciplined* ambiguity, with an illusive, eye-straining quality that contaminates the entire graphic. It has no place in data graphical design.

[1] Jacques Bertin, *Semiology of Graphics: Diagrams, Networks, Maps* (Madison, Wisconsin, 1983, translated by William J. Berg), 80; this book is the English translation of Bertin's important work, *Sémiologie graphique* (Paris, 1967).

The Grid

One of the more sedate graphical elements, the grid should usually be muted or completely suppressed so that its presence is only implicit—lest it compete with the data. Grids are mostly for the initial plotting of data at home or office rather than for putting into print. Dark grid lines are chartjunk. They carry no information, clutter up the graphic, and generate graphic activity unrelated to data information. This grid camouflages the profile of the data in the age-sex pyramid of the population of France in 1967:

Population of France, by Age and Sex: January 1, 1967

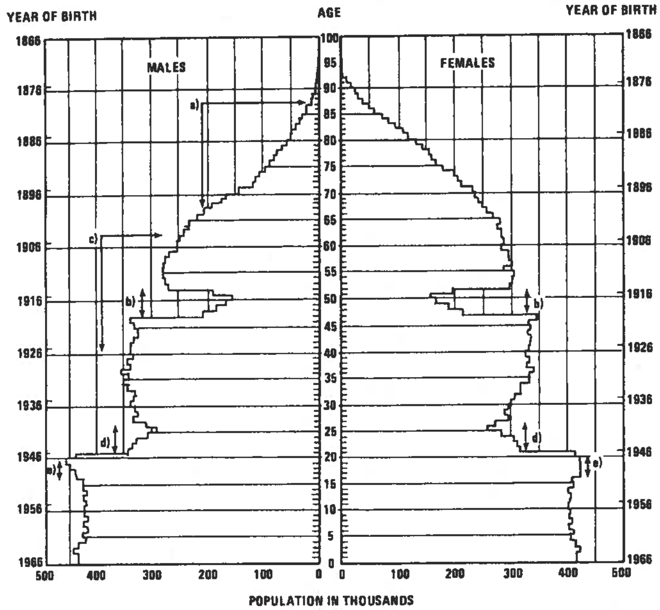

(a) Military losses in World War I
(b) Deficit of births during World War I
(c) Military losses in World War II
(d) Deficit of births during World War II
(e) Rise of births due to demobilization after World War II

A revision quiets the grid and gives emphasis to the data:

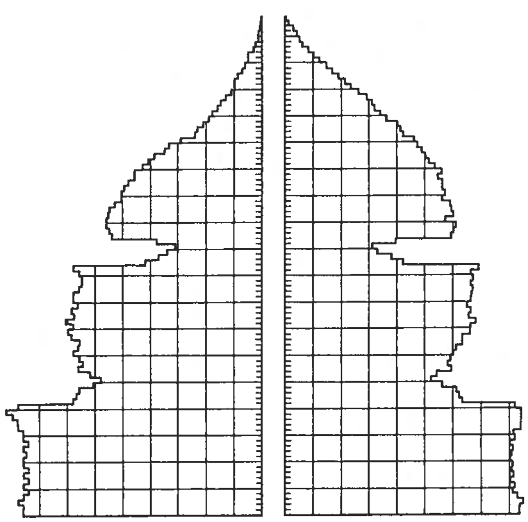

Based on data in Institut National de la Statistique et des Études Économiques, *Annuaire statistique de la France, 1968* (Paris, 1968), 32–33; redrawn in Henry S. Shryock and Jacob S. Siegel, *The Methods and Materials of Demography* (Washington, DC, 1973), vol. 1, 242.

The space occupied by the doubled grid lines consumes 18 percent of the area of this otherwise most ingenious design, a "multiwindow plot." Optical white dots appear at the intersections of the grid lines. (The plot shows the following: The large square contains X_4, X_7 scatterplots for the indicated levels of X_1 and X_3. The marginal plots on the right are conditioned on X_3 and the plots at the top on X_1. The upper right corner shows the unconditional X_4, X_7 scatter.) Redrawing eliminates the vibration:

Paul A. Tukey and John W. Tukey, "Data-Driven View Selection; Agglomeration and Sharpening," in Vic Barnett, ed., *Interpreting Multivariate Data* (Chichester, England, 1981), 231–232.

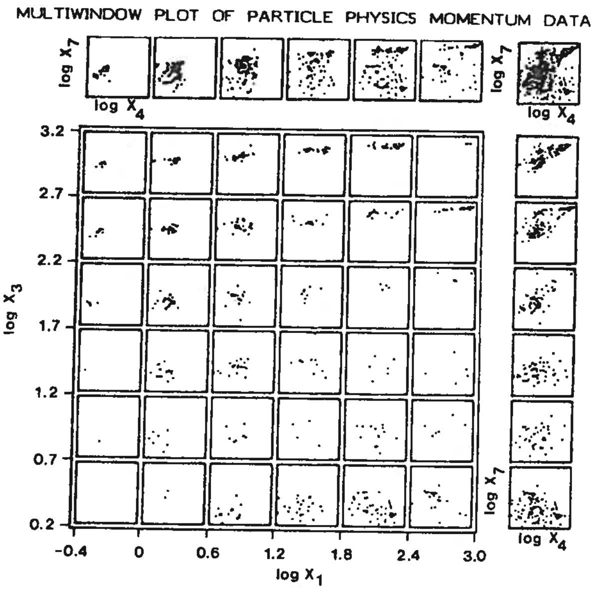

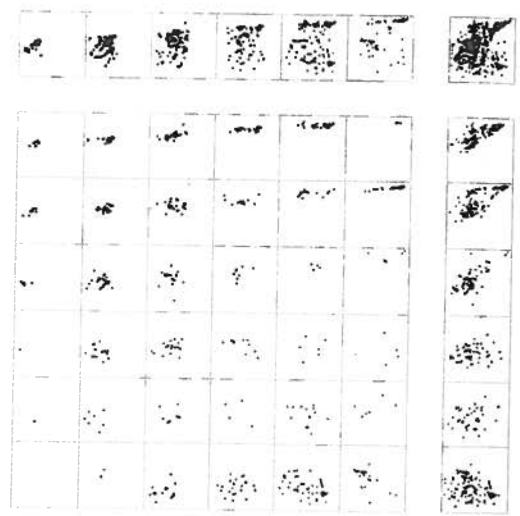

The grid in the classic Marey train schedule is very active:

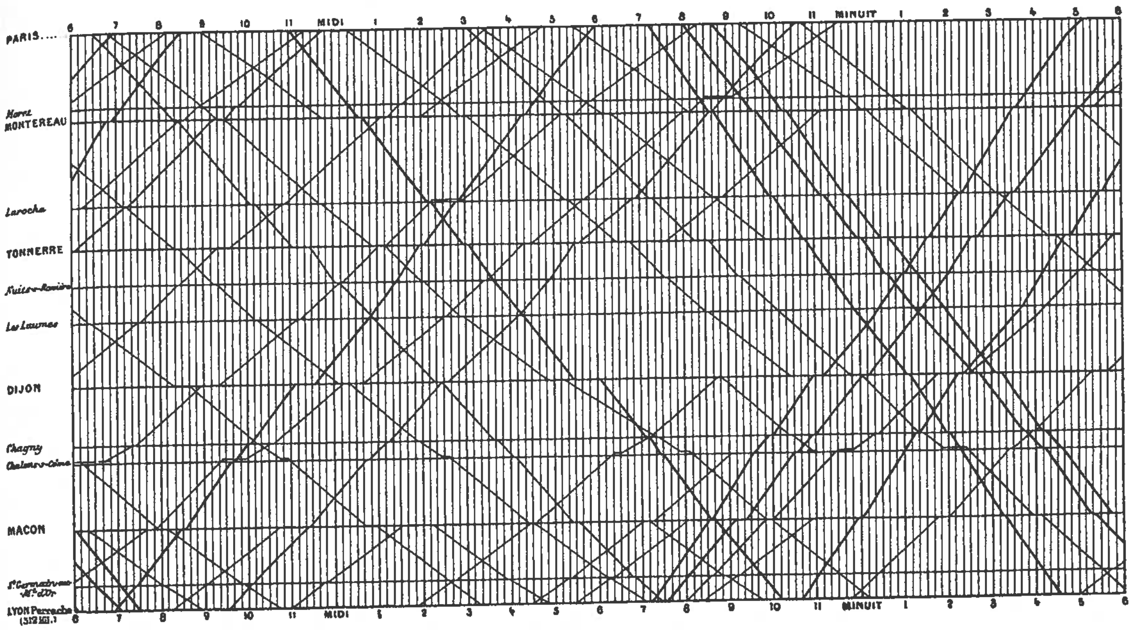

Thinning the grid lines helps a little bit:

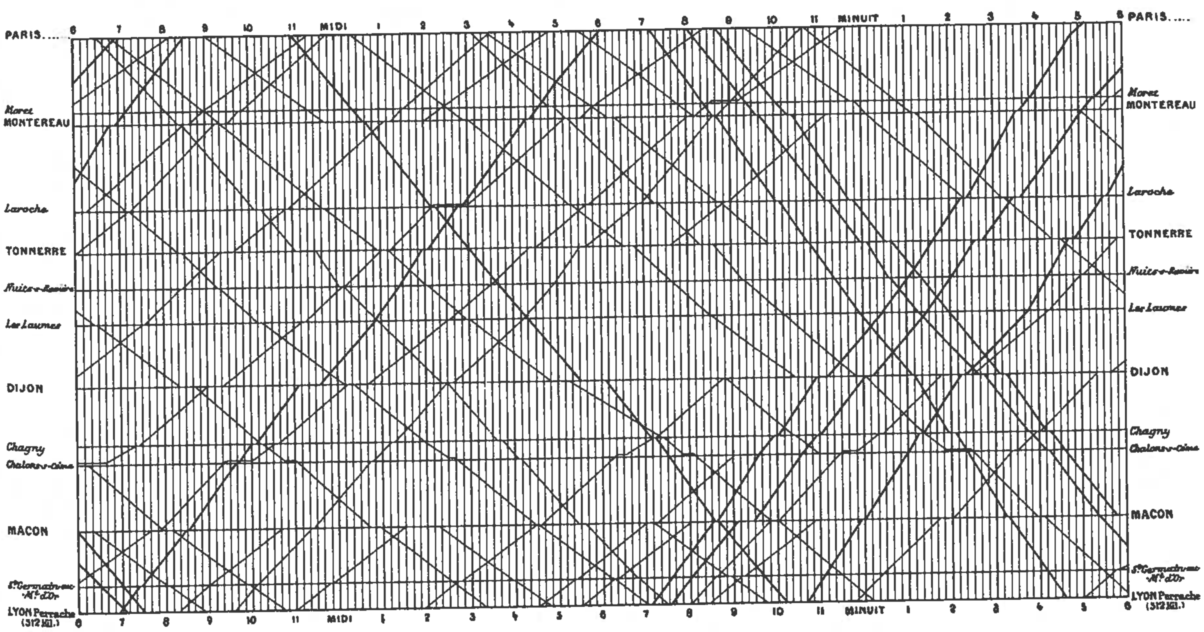

A better treatment, however, is a *gray grid:*

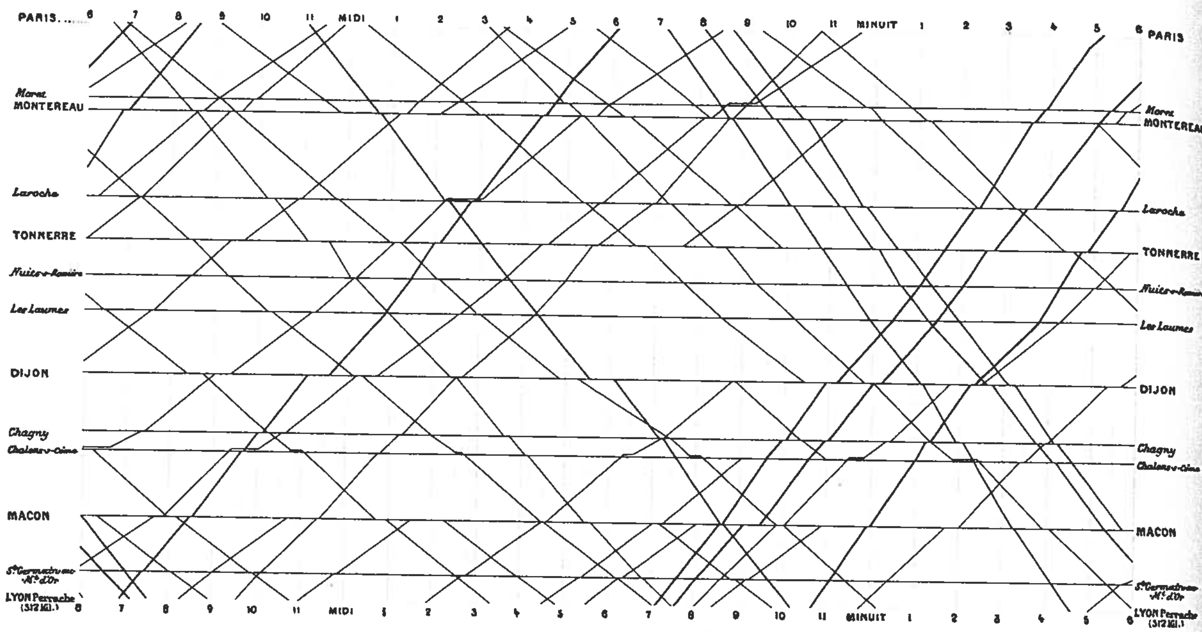

When a graphic serves as a look-up table, then a grid may help in reading and interpolating. But even in this case the grids should be muted relative to the data. A gray grid works well and, with a delicate line, may promote more accurate data reconstruction than a dark grid.

Most ready-made graph paper comes with a darkly printed grid. The reverse (unprinted) side should be used, for then the lines show through faintly and do not clutter the data. If the paper is heavily gridded on both sides, throw it out.

Self-Promoting Graphics: The Duck

When a graphic is taken over by decorative forms or computer debris, when the data measures and structures become Design Elements, when the overall design purveys Graphical Style rather than quantitative information, then that graphic may be called a *duck* in honor of the duck-form store, "Big Duck." For this building the whole structure is itself decoration, just as in the duck data graphic. In *Learning from Las Vegas*, Robert Venturi, Denise Scott

Brown, and Steven Izenour write about the ducks of modern architecture—and their thoughts are relevant to the design of data graphics as well:

> When Modern architects righteously abandoned ornament on buildings, they unconsciously designed buildings that *were* ornament. In promoting Space and Articulation over symbolism and ornament, they distorted the whole building into a duck. They substituted for the innocent and inexpensive practice of applied decoration on a conventional shed the rather cynical and expensive distortion of program and structure to promote a duck. . . . It is now time to reevaluate the once-horrifying statement of John Ruskin that architecture is the decoration of construction, but we should append the warning of Pugin: It is all right to decorate construction but never construct decoration.[2]

[2] Robert Venturi, Denise Scott Brown, and Steven Izenour, *Learning from Las Vegas* (Cambridge, revised edition, 1977), 163. The initial statement of the duck concept is found on 87–103.

Big Duck, Flanders, New York; photograph by Edward Tufte, July 2000.

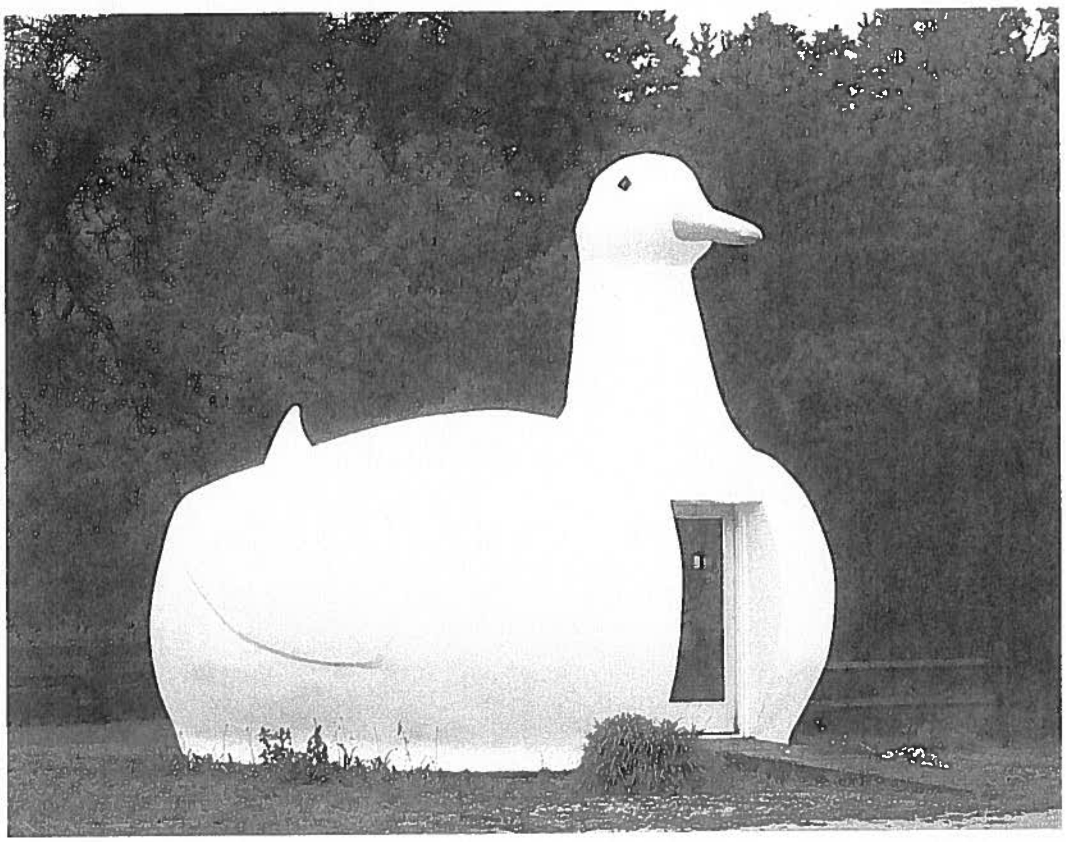

The addition of a fake perspective to the data structure clutters many graphics. This variety of chartjunk, now at high fashion in the world of Boutique Data Graphics, abounds in corporate annual reports, the phony statistical studies presented in advertisements, the mass media, and the more muddled sorts of social science research.

A series of weird three-dimensional displays appearing in the magazine *American Education* in the 1970s delighted connoisseurs of the graphically preposterous. Here five colors report, almost by happenstance, only five pieces of data (since the division within each year adds to 100 percent). This may well be the worst graphic ever to find its way into print:

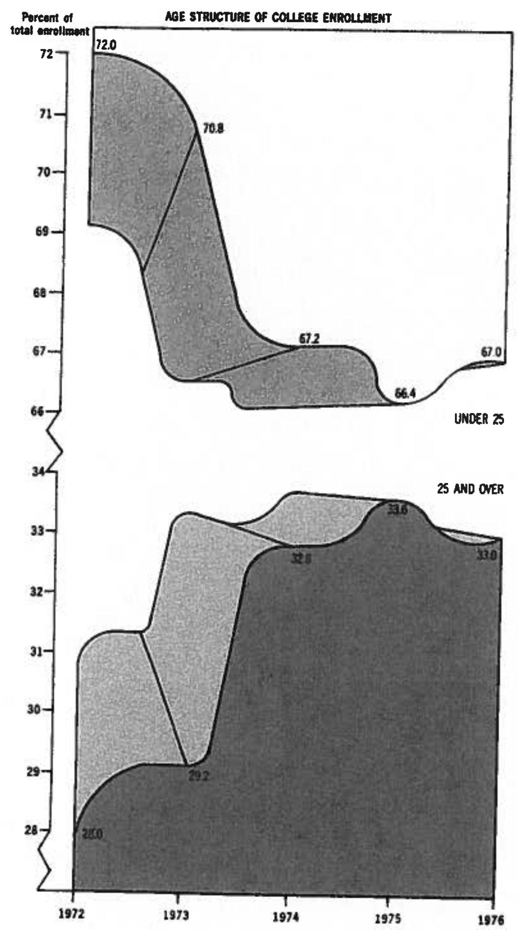

There are some superbly produced ducks:

William L. Kahrl, et al., *The California Water Atlas* (Sacramento, 1978, 1979), 55.

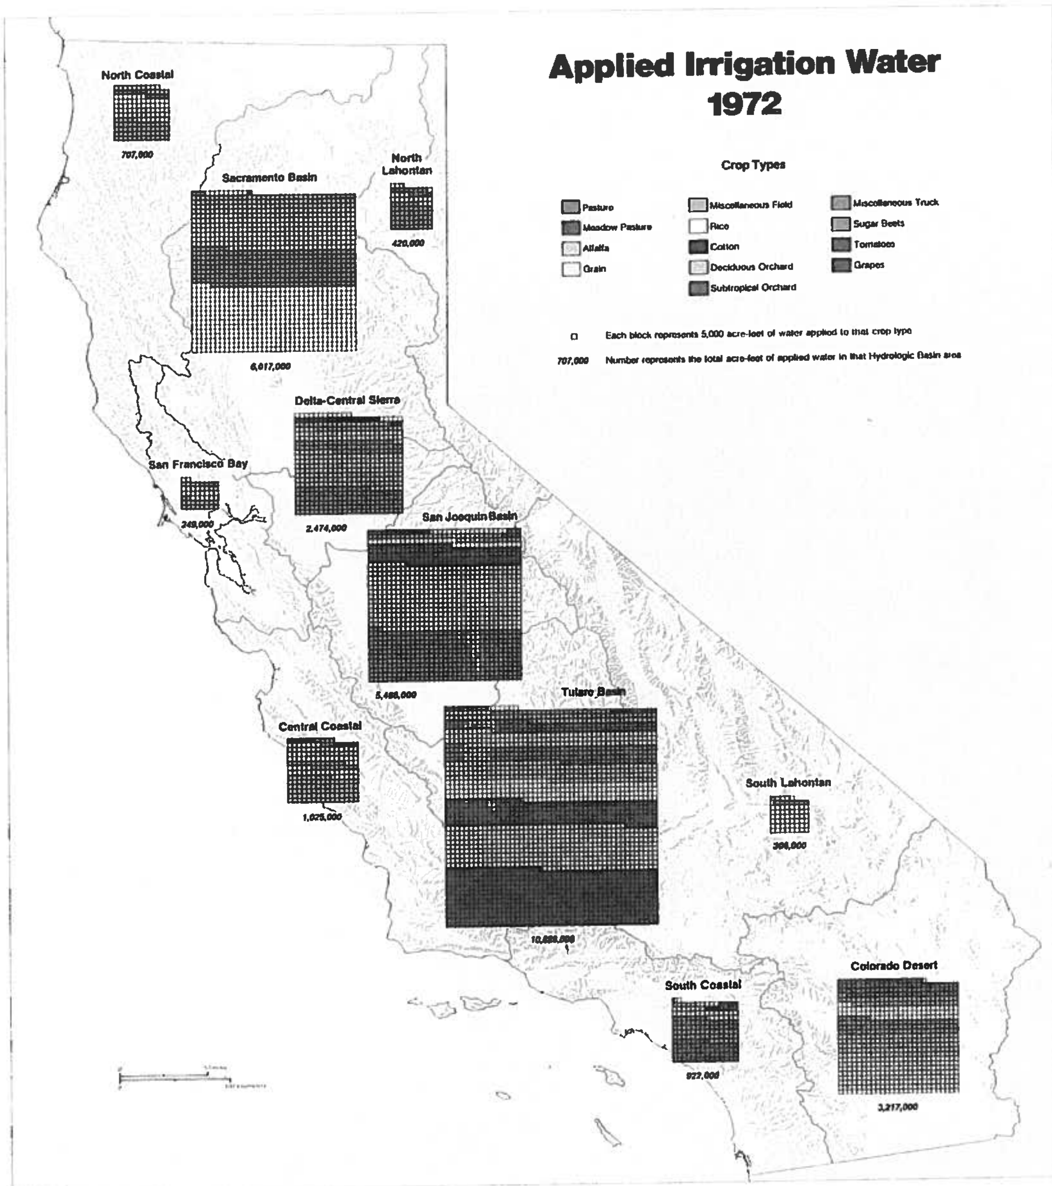

Occasionally designers seem to seek credit merely for possessing a new technology, rather than using it to make better designs. Computers and their affiliated apparatus can do powerful things graphically, in part by turning out the hundreds of plots necessary for good data analysis. But at least a few computer graphics only evoke the response "Isn't it remarkable that the computer can be programmed to draw like that?" instead of "My, what interesting data."

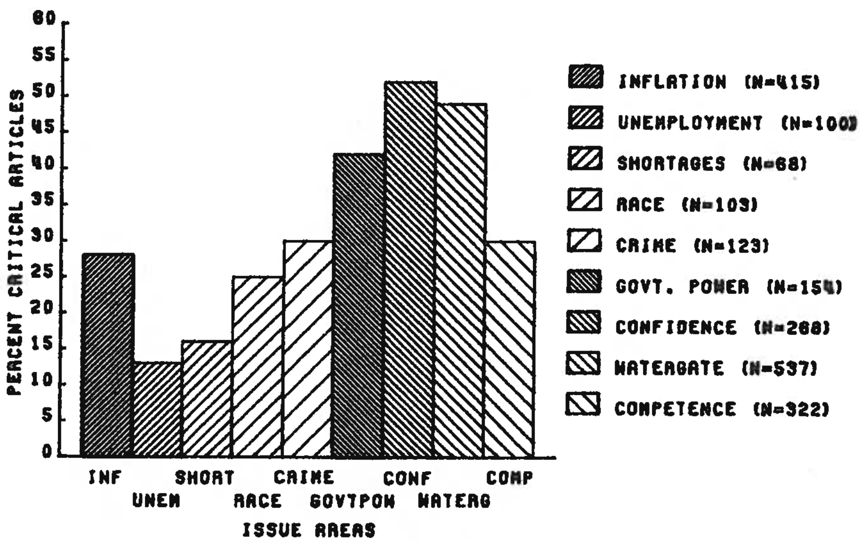

The symptoms of the We-Used-A-Computer-To-Build-A-Duck Syndrome appear in this display from a professional journal: the thin substance; the clotted, crinkly lettering all in upper-case sans serif; the pointlessly ordered cross-hatching; the labels written in computer abbreviations; the optical vibration—all these the by-products of the technology of graphic fabrication. The overly busy vertical scaling shows more percentage markers and labels than there are actual data points. The observed values of the percentages should be printed instead. Since the information consists of a few numbers and a good many words, it is best to pass up the computerized graphics capability this time and tell the story with a table:

Arthur H. Miller, Edie N. Goldenberg, and Lutz Erbring, "Type-Set Politics: Impact of Newspapers on Public Confidence," *American Political Science Review*, 73 (1979), 67–84.

Content and tone of front-page articles in 94 U.S. newspapers, October and November, 1974	Number of articles	Percent of articles with negative criticism of specific person or policy
Watergate: defendants and prosecutors, Ford's pardon of Nixon	537	49%
Inflation, high cost of living	415	28%
Government competence: costs, quality, salaries of pubic employees	322	30%
Confidence in government: power of special interests, trust in political leaders, dishonesty in politics	266	52%
Government power: regulation of business, secrecy, control of CIA and FBI	154	42%
Crime	123	30%
Race	103	25%
Unemployment	100	13%
Shortages: energy, food	68	16%

Conclusion

Chartjunk does not achieve the goals of its propagators. The overwhelming fact of data graphics is that they stand or fall on their content, gracefully displayed. Graphics do not become attractive and interesting through the addition of ornamental hatching and false perspective to a few bars. Chartjunk can turn bores into disasters, but it can never rescue a thin data set. The best designs (for example, Minard on Napoleon in Russia, Marey's graphical train schedule, the cancer maps, the *Times* weather history of New York City, the chronicle of the annual adventures of the Japanese beetle, the new view of the galaxies) are *intriguing and curiosity-provoking*, drawing the viewer into the wonder of the data, sometimes by narrative power, sometimes by immense detail, and sometimes by elegant presentation of simple but interesting data. But no information, no sense of discovery, no wonder, no substance is generated by chartjunk.

Forgo chartjunk, including
moiré vibration,
the grid, and the duck.

Painting is special, separate, a matter of meditation and contemplation, for me, no physical action or social sport. As much consciousness as possible. Clarity, completeness, quintessence, quiet. No noise, no schmutz, no schmerz, no fauve schwärmerei. Perfection, passiveness, consonance, consummateness. No palpitations, no gesticulation, no grotesquerie. Spirituality, serenity, absoluteness, coherence. No automatism, no accident, no anxiety, no catharsis, no chance. Detachment, disinterestedness, thoughtfulness, transcendence. No humbugging, no button-holing, no exploitation, no mixing things up.

Ad Reinhardt, statement for the catalogue of the exhibition, "The New Decade: 35 American Painters and Sculptors," Whitney Museum of American Art, New York, 1955.