CHAPTER 11

Data Maps: A Thicket of Thorny Choices

A single set of numerical data can yield markedly dissimilar maps. By manipulating breaks between categories of data to be shaded on a choropleth map, for instance, a mapmaker can often create two distinctly different spatial patterns. A single map is thus just one of many maps that might be prepared from the same information, and the map author who fails to look carefully at the data and explore cartographic alternatives easily overlooks interesting spatial trends or regional groupings.

Wary map users must watch out for statistical maps carefully contrived to prove the points of self-promoting scientists, manipulating politicians, misleading advertisers, and other propagandists. Meanwhile, this is an area in which the widespread use of mapping software has made unintentional cartographic self-deception inevitable. How many software users know that using area-shading symbols with magnitude data produces misleading maps, or that size differences between areal units such as counties and census tracts can radically distort map comparisons?

This chapter uses several simple hypothetical examples featuring a fictional electronic device we'll call a "gizmo" to examine the effects of areal aggregation and data classification on mapped patterns. Anyone interested in public-policy analysis, marketing, social science, or disease control needs to know how maps based on numbers can yield useful information as well as flagrant distortions.

Aggregation, Homogeneity, and Areal Units

Most quantitative maps display data collected for areas such as counties, states, and countries. When displayed on a map, present-

Number of Gizmos

1,000	100	50	100	50	100	50
200	100	200	100	200	100	200
100	200	100	4,000	100	200	100
200	400	200	400	200	400	3,000

Number of Households

2,000	200	100	200	100	200	100
200	100	200	100	200	100	200
100	200	100	4,000	100	200	100
100	200	100	200	100	200	1,500

Gizmos per Household

0.5	0.5	0.5	0.5	0.5	0.5	0.5
1.0	1.0	1.0	1.0	1.0	1.0	1.0
1.0	1.0	1.0	1.0	1.0	1.0	1.0
2.0	2.0	2.0	2.0	2.0	2.0	2.0

Figure 11.1. Town-unit number tables showing number of gizmos (top left), number of households (top right), and average number of gizmos per household (bottom) for twenty-eight hypothetical towns.

Number of Gizmos

2,300	5,700	4,150

Number of Households

3,100	5,500	2,600

Gizmos per Households

0.74	1.04	1.60

Figure 11.2. County-unit number tables of number of gizmos (left), number of households (middle), and average number of gizmos per household (right) for a three-county aggregation of the twenty-eight hypothetical towns in figure 11.1.

ed on a statistical plot, or analyzed using correlation coefficients or other measures, geographic data produce results that reflect the type of areal unit. Because different areal aggregations of the data might yield substantially different patterns or relationships, the analyst should qualify any description or interpretation by stating the type of geographic unit used. Noting that values generally increase from north to south "at the county-unit level" warns the reader (and the mapmaker as well!) that a different trend might arise with state-level data, for instance.

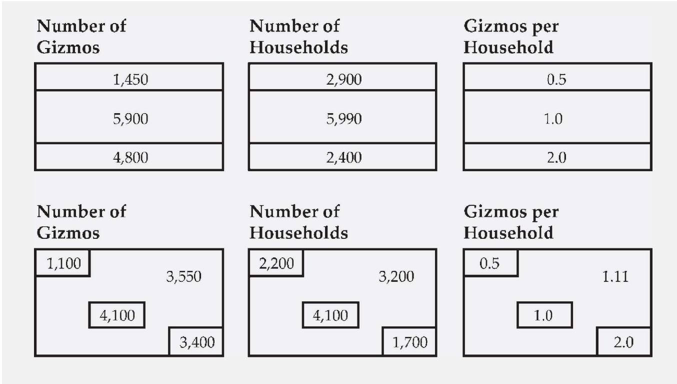

Figure 11.3. County-unit number tables based on other aggregations of the twenty-eight towns into counties.

Areal aggregation can have a striking effect on the mapped patterns of rates and ratios. A ratio such as the average number of gizmos per household might, for example, produce radically different maps when the data are aggregated separately by counties and by the towns that make up these counties. The three town-level maps in figure 11.1 are spatially ordered number tables, without graphic symbols, so that we can see how rate calculations depend on what boundaries are used and how they are drawn. The upper left-hand map shows the number of gizmos in each of twenty-eight towns, the upper right-hand map represents the number of households, and the lower map portrays the gizmo-ownership rate. Note the straightforward top-to-bottom pattern of the rates: low in the upper tier of towns, average in the two middle tiers, and high in the lower tier. Note also that three towns in the upper left, lower right, and just below the center of the region have relatively high numbers of households. These variations in household density underlie the markedly different left-to-right trend in gizmo-ownership rates in figure 11.2, based on the same data aggregated by county.

Spatial pattern at the town-unit level of aggregation depends on how somewhat arbitrary political boundaries group towns into counties. Figure 11.3 uses two additional aggregations of these

Number of Gizmos

190	285	200	350	350	210	890
455	450	1,085	960	895	520	1,260
355	315	525	480	595	360	700
130	120	80	100	80	110	100

Number of Households

100	150	100	200	100	50	100
350	300	700	600	500	200	300
500	450	700	600	700	400	500
650	600	400	500	400	550	500

Gizmos per Household

1.90	1.90	2.00	1.75	3.50	4.20	8.90
1.30	1.50	1.55	1.60	1.79	2.60	4.20
0.71	0.70	0.75	0.80	0.85	0.90	1.40
0.20	0.20	0.20	0.20	0.20	0.20	0.20

Figure 11.4. Patterns of the number of gizmos, the number of households, and the gizmo-ownership rate radically different from those in figure 11.1 could yield county-unit patterns identical to those in figure 11.2.

twenty-eight towns to demonstrate the possible effect of historical accident. The upper row of maps shows an alternative aggregation of towns into three horizontal counties that reflect the town-level top-to-bottom trend. In contrast, the lower series of maps shows an equally plausible aggregation into four counties, three based on the concentrations of households and one comprising the balance of the region. The gizmo-ownership map for this lower set isolates what might be more urban counties from a single much larger, more rural county with an average of slightly more than one gizmo per household. Graytone area symbols would yield very different choropleth maps for the three sets of rates shown in the right-hand maps of figures 11.2 and 11.3.

Another example illustrates how areal aggregation can affect geographic pattern. Whereas figure 11.3 demonstrates that different aggregations of towns into counties can yield markedly different county-level patterns, figure 11.4 illustrates how a single aggregation can produce the same county-level pattern from markedly different town-level patterns. Note that the town-level maps in figure 11.4 reflect a pattern of gizmo-ownership rates very different from that in figure 11.1. Note in particular the progression of rates from a tier of low-ownership towns across the bottom of the region to a peak of much higher rates at the upper right. Yet when aggregated according to the county boundaries in figure 11.2, these data will yield similar county-unit rates. Comparing this trio of spatial number tables with those in figure 11.1 demonstrates the importance of stating clearly the data units used and of not assuming that a trend apparent at one level of aggregation exists at other levels as well.

The counties in these examples obviously are not homogeneous. But can we assume homogeneity even within the towns? What spatial variations in the distribution and density of these 11,200 households lie hidden in the network of town boundaries? Figure 11.5 presents one of many plausible point patterns that could produce the aggregated town-level counts and rates in figure 11.1. Three types of point symbols represent groups of ten, one hundred, and five hundred households. Each symbol represents a group of households owning an average of zero, one, or two gizmos. The small, ten-household symbols represent rural residences, which might lack gizmos because of a lack of connectivity, less spare time, or a low opinion of digital gadgetry. Because of rough terrain, swamps, park- or forestland, and undeveloped federal land, large parts of the region are uninhabited. Of the six large villages, with four hundred or more households, two have two-gizmo households on the average, two have one-gizmo households, and two have gizmo-free households. Although figure 11.5 contains elements of both the top-to-bottom town-level trend in figure 11.1 and the left-to-right county-unit trend in figure 11.2, its pattern of gizmo ownership is more similar to the lower right of figure 11.3, where county boundaries segregate three large population clusters from the balance of the region. Yet even here the differences are striking, again demonstrating how the configuration of areal units can hide interesting spatial detail and present a biased view of a variable's geography.

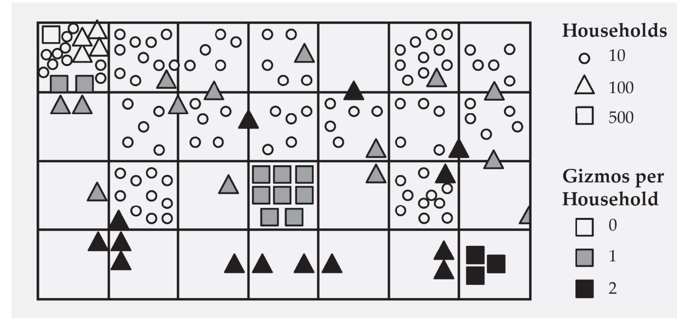

Figure 11.5. Detailed map of gizmo ownership for villages and rural households illustrates one possible spatial structure that could yield the town-unit and county-unit maps in figures 11.1 and 11.2.

Aggregation's effects become even more serious if the careless analyst or naive reader leaps from a pattern based on areal units to conclusions based on individual households. Consider, for instance, the large village toward the lower right-hand corner of figure 11.5. The average gizmo-ownership rate here of 2.0 need not mean that each of the village's 1,700 houses has two gizmos. Some households might have none while others might have three or four or five. One or two residents might even be compulsive collectors—hoarders masquerading as hobbyists—so that more than half the homes have one or none.

If households collecting old gizmos seems far-fetched, consider average household income, an index used frequently by social scientists and marketing analysts. Because of one or two innovative, unscrupulously manipulative, or otherwise successful residents, a small village might have an enormous *mean* household income. More of a statistical quirk than a realistic reflection of overall local prosperity, this high average income might mask the employment of most villagers as household servants, gardeners, or security guards. Because nondisclosure rules prohibit a more precise publication of individual incomes, aggregated census data are the most refined information available. They provide an average for the place but say little about individual residents.

Are areally aggregated data bad? Surely not. In many cases, particularly in public-policy analysis, towns and counties are the truly relevant units for which state and federal governments allocate funds and measure performance. And even more highly aggregated data can be useful, for instance, when governors and senators want to compare their states with the other forty-nine. Local officials and social scientists concerned with differences between neighborhoods readily acknowledge the value of geographic aggregation. Moreover, nondisclosure regulations, which are needed to ensure cooperation with censuses and surveys, require aggregation, and areally aggregated data are better than no data at all. Thus persons who depend on local-area data encourage the Census Bureau to modify boundaries to preserve the homogeneity of census tracts and other reporting areas. And when tract data are not adequate, they sometimes pay for new aggregations of the data to more meaningful areal units.

What else can the conscientious analyst do? Very little aside from the obvious: know the area and the data, experiment with data for a variety of levels of aggregation, and carefully qualify all conclusions.

And what should the skeptical map user do? Look for and compare maps with different levels of detail, and be wary of cartographic manipulators who choose the level of aggregation that best proves their point.

Aggregation, Classification, and Outliers

Choropleth mapping further aggregates the data by grouping all areas with a range of data values into a single category represented by a single symbol. This type of aggregation addresses the difficulty of displaying more than six or seven visually distinct colors or graytones in a consistent light-to-dark sequence. Often the mapmaker prefers only four or five categories, especially when the area symbols available do not afford an unambiguous graded series. (For aesthetic reasons or to avoid confusion with interior

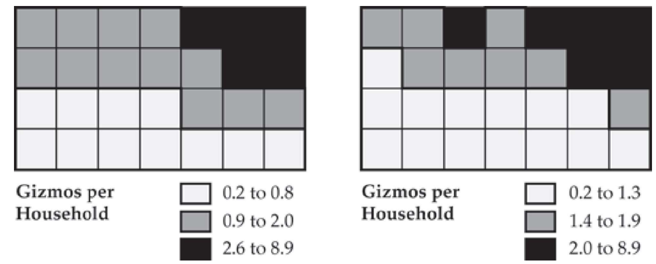

Figure 11.6. Different sets of categories yield different three-category choropleth maps for the data in figure 11.4.

lakes or areas without data, opaque white and solid black are not good graytone symbols for choropleth maps.)

But classification introduces the risk of a mapped pattern that distorts spatial trends. Arbitrary selection of breaks between categories might mask a clear, coherent trend with a needlessly fragmented map or oversimplify a meaningfully intricate pattern with an excessively smoothed view. Figure 11.6 illustrates the influence of class breaks on the appearance of choropleth maps of the town-level gizmo-ownership rates in figure 11.4. Note that the map on the left presents a clear, straightforward, readily remembered upward trend toward a peak at the upper right of the region, whereas the map at the right offers a more fractured view of the same data.

Classification raises many questions. Which map, if either, is right? Or if “right” sounds too dogmatic, which provides a better representation of the data? Don’t both maps hide much variation in the broad third category, represented by the darkest symbol? Shouldn’t the seven towns with rates of 0.2 occupy a category by themselves? Is a difference of, say, 0.1 at the lower end of the overall range of data values more important than a similar difference at the upper end? Can a three-class map provide even a remotely adequate solution?

These questions are vital not only to map users but also to map authors, particularly those using graphics software but untrained

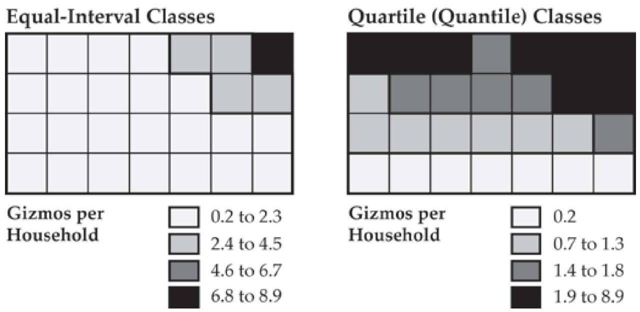

Figure 11.7. Two common classing schemes used as "defaults" by choropleth mapping software yield radically different four-category patterns for the data in figure 11.4.

in cartography. Software applications usually provide a few options for "automatic" classification, and naive mapmakers often settle for one of the easier options. Sometimes the software even provides a map instantly, without offering a choice of classification strategies. Called a *default option*, this automatic choice of class breaks is a marketing ploy that gives the hesitant prospective purchaser an immediate success.

But does the default give you a good map? Figure 11.7 shows four-category mapped patterns produced by two common default classing options for the same town-level gizmo-ownership data used in figure 11.6. The *equal-intervals* scheme, on the left, divides the range (8.7) between the lowest and highest data values (from 0.2 to 8.9) into four equal parts (each spanning 2.175 units). Note, though, that this classification assigns most of the region to a single category and that the third category (from 4.6 to 6.7) is empty. Of possible use when data values are uniformly distributed across the range, the only consistent asset of equal-interval classification is ease of calculation.

By contrast, the *quartile* scheme, on the right, ranks the data values and then divides them so that all categories have the same number of areal units. Of course, only an approximately equal balance is possible when the number of areas is not a multiple of four or when

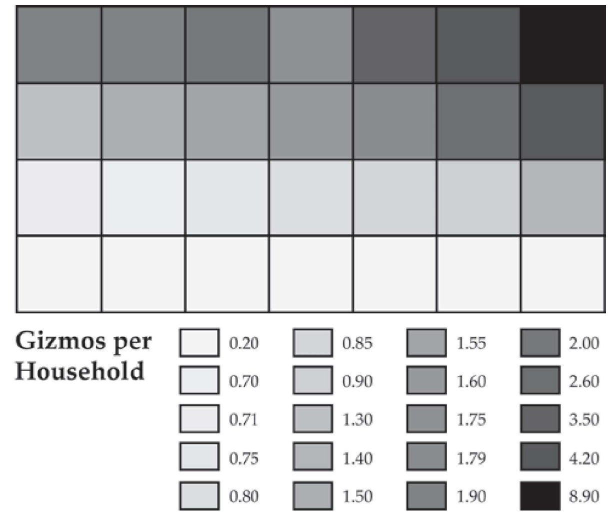

Figure 11.8. A continuous-tone, non-classed choropleth map for the data in figure 11.4.

a tie thwarts an equal allocation (as occurs here at the upper left, where the highest category receives both of the towns with rates of 1.9). Although the map pattern is more visually balanced, the upper category is broad and highly heterogeneous, and the break between the second and third categories falls between two very close values (1.3 and 1.4). Yet the map based on these four quartile categories does have meaning for the user interested in the locations of towns in the highest and lowest quarters of the data values. Called *quintiles* for five categories and *quantiles* more generally, this rank-and-balance approach can accommodate any number of classes.

Some mapping applications offer the option of a "no-class" or "class-less" choropleth map, on which each unique data value (perhaps up to fifty of them) receives a unique graytone. In principle this might seem a good way to sidestep the need to set class breaks. But as figure 11.8 illustrates, the graytones might not form a well-ordered series, and the map key is either abbreviated or cumbersome. Moreover, assigning each unique value its own category can

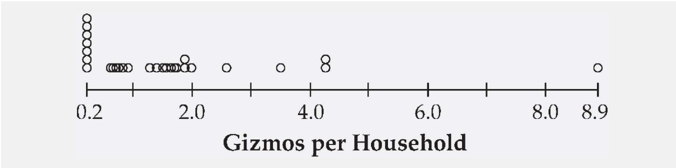

Figure 11.9. Number line for the town-level gizmo-ownership rates in figure 11.4.

destroy a clear, easily remembered picture of a strong, meaningful spatial trend. This ideal solution might not be so ideal after all.

Eschewing defaults and panaceas, the astute map author begins by asking two basic questions: How are the data distributed throughout their range? And what, if any, class breaks might have particular meaning to the map user? The answer to this second question depends on the data and on whether the map author deems useful a comparison with the national or regional average. On state-level maps, for instance, a break at the United States average would allow governors and senators to compare their constituents' or their own performance with that of the rest of the nation. Of course, the map key would have to identify this break to make it truly meaningful.

After addressing the question of meaningful breaks, the conscientious map author might then plot a *number line* similar to that in figure 11.9. A horizontal scale with tick marks and labels represents the range of the data. Each dot represents a data value, and identical values plot at the same position along the scale, one above the other. The resulting graph readily reveals natural breaks, if any occur, and distinct clusters of homogeneous data values, which the classification ought not subdivide. Number lines allow the map author to visualize the distribution of data values and to choose an appropriate number of categories and appropriate positions for class breaks. Computer algorithms can also search the data distribution for an optimal set of breaks, but in many cases the computer-determined optimum is not significantly better

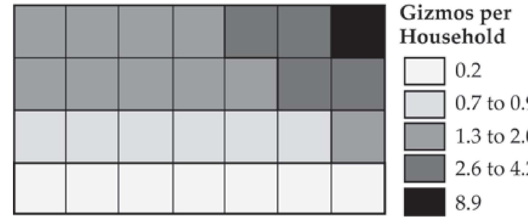

Figure 11.10. Choropleth map based on the number line in figure 11.9 and the character of the data.

than a visually identified suboptimal grouping. Rounded breaks and a more balanced allocation of places among categories can be important secondary factors in choropleth mapping.

Extremely high or extremely low values isolated from the rest of the distribution can confound both human cartographers and sophisticated mapping software. Should these *outliers* be grouped with markedly more homogeneous clusters higher or lower on the number line? Should each be accorded its own category? Can two or three widely separated data values at either end of the distribution be grouped into a single, highly heterogeneous category? Or should each outlier be treated as its own category, with its own symbol, at the risk of reducing graphic differentiation between area symbols? Or might the map author treat outliers as outcasts—errors or deviants that "don't belong"—and either omit them or give them a special symbol?

No simple, standard solution addresses all outliers. The map author should know the data, know whether these deviant values are real or improbable, and know whether a large difference between outliers really matters. Also important is the relation of outliers to the theme of the map and the interests of map users. For the gizmo-ownership data in figure 11.9, an average of 8.9 gizmos per household surely is not only exceptional but probably significantly higher than its neighboring values at 4.2. If not an error, it deserves special treatment in a category of its own. The next four lower values, 4.2 (twice), 3.5, and 2.6, might then constitute a single category; all are above the more plausible rate of 2.0, and yet 4.2 gizmos per household is not improbable, especially in an affluent area.

Other breaks seem warranted between 0.9 and 1.3, a gap that includes the inherently meaningful rate of one gizmo per household, and between 0.2 and 0.7, to separate the seven technophobic towns at the lower end of the distribution. The resulting five-category map in figure 11.10 provides not only an honest, meaningful representation of the data values and their statistical distribution, but a straightforward portrayal of the spatial trend as well. An arbitrary classification, such as a computer program's default categories, is unlikely to do as well, even with six or more categories.

Classification, Correlation, and Visual Perception

Choropleth maps readily distort geographic relationships between two distributions. Hastily selected or deliberately manipulated categories can diminish the visual similarity of two essentially identical trends or impose an apparent similarity between two very different patterns.

Consider as a case in point figure 11.11, a spatial-data table and number line for the mean number of children per household, which has a strong town-level relationship to gizmo ownership. Although the range of data values is not as broad for this index of family size, the highest values are at the upper right and the lowest values occur across the bottom of the region. Towns toward the right and toward the top of the region generally have more children in the home than do towns toward the bottom or left edge of the map. That the pair of maps in figure 11.12 shows identical spatial patterns for children and gizmos is thus not surprising.

Statistical analysts commonly depict correlation with a two-dimensional scatterplot, with data values for one variable measured along the vertical axis and those for the other scaled along the horizontal axis. A dot represents each place, and the density and orien-

Average Number of Children per Household

2.25	2.30	2.31	2.35	2.52	2.41	3.20
2.04	2.15	2.18	2.20	2.27	2.42	2.53
1.62	1.53	1.63	1.63	1.71	1.92	2.11
0.37	0.39	0.31	0.36	0.49	0.44	0.30

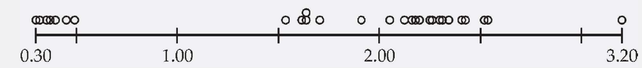

Figure 11.11. Spatial-data table and number line for average number of children per household.

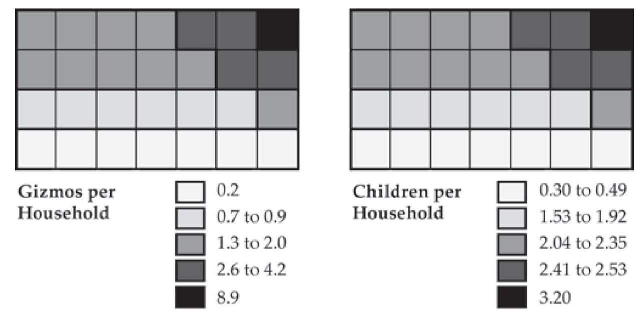

Figure 11.12. Choropleth maps with identical patterns for gizmo-ownership rate and average number of children per household.

tation of the point cloud indicates the strength and direction of the correlation. Figure 11.13 is a pair of scatterplots, both showing the strong positive association between the household rates for children and gizmos. The perpendicular lines extending from the scales of the left-hand scatterplot into the scatter of points represent the class breaks in figure 11.12. These two sets of four lines each divide the scatterplot into an irregular five-by-five grid. Because all dots on the left-hand scatterplot lie within one of the five diagonal cells, the two five-category maps in figure 11.12 have identical patterns, enhancing the impression of a strong correlation.

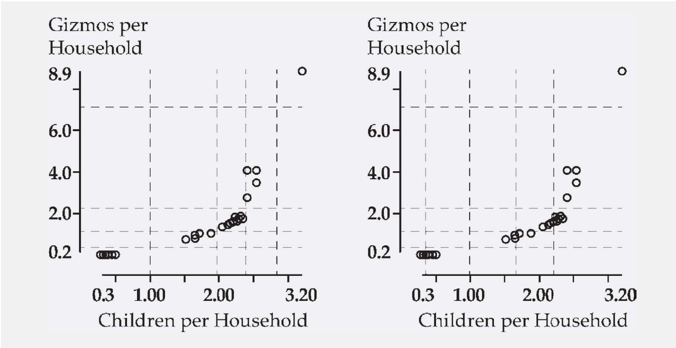

Figure 11.13. Scatterplots for the town-level gizmo-ownership rate and average number of children per household. Additional lines on the left-hand scatterplot represent class breaks for the pair of maps in figure 11.12. Additional lines on the right-hand scatterplot show breaks used in figure 11.14.

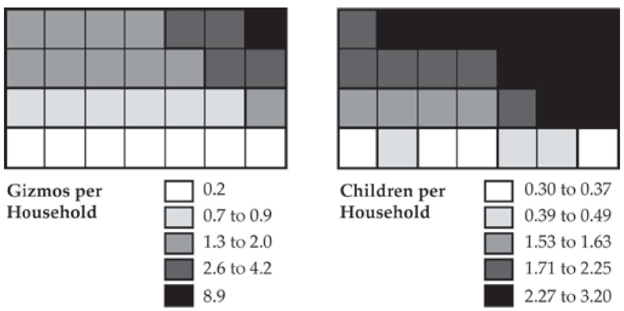

Figure 11.14. Distinctly different choropleth maps suggest minimal correlation between gizmo ownership and family size.

Figure 11.13's right-hand scatterplot adds some cartographic skulduggery. As before, the perpendicular lines from the scales into the point cloud represent class breaks and form a five-by-five grid. But note that this configuration of breaks places all but four dots in an off-diagonal cell so that few towns will belong to the same category on both maps. Figure 11.14 demonstrates the resulting dissimilarity in map pattern and suggests a mediocre correlation at best. Similar tactics might make a weak relationship appear strong,

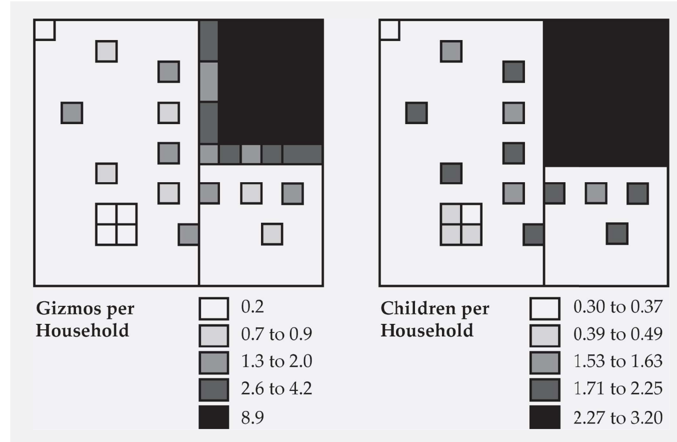

Figure 11.15. Similarity between large areas can distort visual estimates of correlation by masking significant dissimilarity between small areas. Numerical data and mapping categories are identical to those for the more obviously dissimilar pair of maps in figure 11.14.

especially if the maps are identical for the highest category, with the darkest symbol. Indeed, the spatial correspondence of the darkest, most eye-catching symbols strongly influences judgments of map similarity by naive map users. Some will even regard as similar two maps with roughly equal amounts of the darkest symbol—even if the high areas are in different parts of the region! Different area symbols for the two maps and different numbers of categories are other ways of tricking the map user or deluding oneself.

Another visual distortion might lie in the base map the data are plotted on. Not all sets of areal units are as uniform and visually equivalent as the square towns in the preceding examples. Figure 11.15 demonstrates this point with a deceptively similar-looking pair of maps based on the numerical data and class breaks of the visually dissimilar maps in figure 11.14. These twenty-eight towns vary markedly in size, and similarity is high because the largest towns belong to the same category. Towns not in the same category on both maps are smaller and less visually influential.

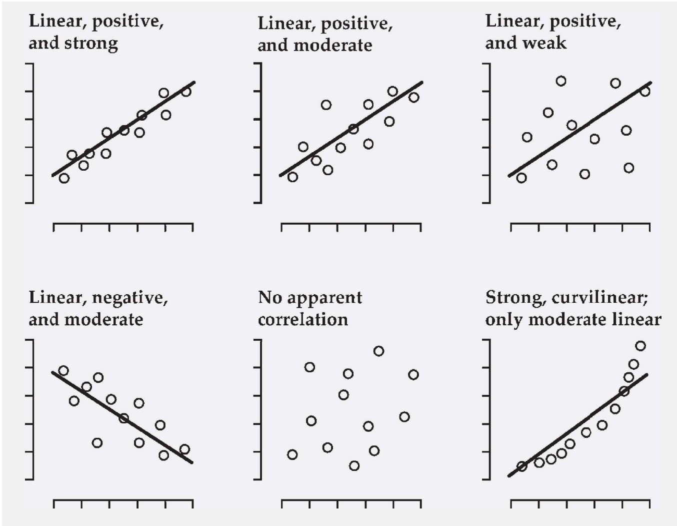

Figure 11.16. Scatterplots and trend lines for various types of correlation.

Although this example involving gizmo ownership is contrived, it is not atypical. Wards, census tracts, congressional districts, and other areal units designed to have similar populations often vary widely in area because of variations in population density. Disparities are even worse on county-unit maps, where populous metropolitan counties often are much smaller than rural counties with few inhabitants. The careful map user never judges numerical correlation by the similarity in map pattern alone and is especially cautious when some data areas are much bigger than others.

To avoid estimates of correlation biased by the size of areal units, the astute analyst will inspect the more egalitarian scatterplot, on which identical dots represent each area. As figure 11.16 illustrates, the density and orientation of the point cloud reflect the strength and direction of the correlation. If a straight line provides a good generalization of the point cloud, the correlation is called linear and the scatter of points around the line indicates the strength of the *linear correlation*. Positive relationships slope upward to the

right, negative relationships slope downward to the right, and a point cloud without a discernible relationship has no apparent slope. Weak correlations have a wide, barely coherent scatter about the trend line, whereas for strong linear correlations most points are near or on the line. Not all correlations are linear, though; a strong *curvilinear correlation* has a marked curved trend, which a curved line fits better than a straight line.

Statisticians use a single number, the *correlation coefficient*, to measure the strength and direction of a linear correlation. Represented by the symbol r, the correlation coefficient shows the direction of the relationship by its sign and the strength of the relationship by its absolute value. The coefficient ranges from +1.00 to –1.00; r would be 0.9 or higher for a strong positive correlation, –0.9 or lower for a strong negative correlation, and close to zero for an indeterminate or very weak correlation. (As a rule of thumb, squaring r yields the proportion of one variable's variation accounted for by the other variable. Thus, if r is –0.6, the correlation is negative and one variable might be said to "explain" 36 percent of the other variable. A correlation coefficient measures only association, not causation, which depends on logic and supporting evidence.)

Maps, scatterplots, and correlation coefficients are complementary, and the analyst interested in correlation relies on all three. The correlation coefficient, which provides a concise comparison for a pair of variables, measures only linear correlation. Yet a scatterplot quickly reveals a strong curvilinear relationship, with a mediocre value of r. Scatterplots also show outliers, which can greatly bias the calculation of r. But reliance on visual estimation makes scatterplots poor for comparing strengths of relationships. Moreover, scatterplots and correlation coefficients tell us nothing about the locations of places, whereas maps, which present spatial trends, can offer unreliable estimates of correlation.

Maps also show a different kind of correlation, a *geographic correlation* distinct from the statistical correlation of the scatterplot and the correlation coefficient. Statistical correlation is aspatial

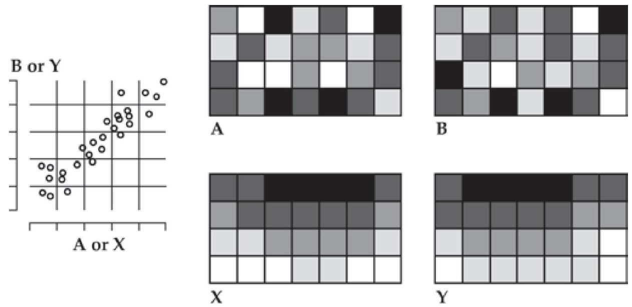

Figure 11.17. Two pairs of variables with identical scatterplots, correlation coefficients (r = .93), and class breaks, yet distinctly different map patterns.

and reveals nothing about spatial trends. Figure 11.17 demonstrates this difference with two map pairs distinct in spatial pattern yet identical in scatterplot and correlation coefficient. Variables A and B, which share a comparatively chaotic, fragmented pattern, clearly differ in geographic correlation from variables X and Y, which have a distinct common trend with higher values toward the top of the region and lower values toward the bottom. Although not identical, the maps for X and Y suggest the influence of a third, underlying geographic factor, such as latitude, ethnicity, soil fertility, or proximity to a major source of pollution. Despite the problems posed by areal aggregation, the analyst of geographic data who explores correlation without also checking for spatial pattern is either ignorant, careless, or callous. And the nonskeptical reader is easily misled.

Whether expressed numerically or with maps, correlations based on spatially aggregated data are vulnerable to the *ecological fallacy*, whereby a relationship demonstrated for one level of areal aggregation—say, with county units—is presumed to hold for other aggregations (such as states) as well as for individuals. (Regardless of their size, areal units are considered *ecological* units, rather than individuals.) For example, a finding that areas above average in number of years of education tend to be above average

in income does not mean that people with master's degrees are necessarily well paid—graduate students pursuing a doctorate are a case in point.

Places, Time, and Small Numbers

Areal data can yield particularly questionable patterns when choropleth maps show rates based on infrequent events, such as deaths from a rare type of cancer. Yet disease maps based on small numbers are a common tool of the epidemiologist, who uses mapping to explore the possible effects on human health of radon-rich soils, incinerators, chemical-waste dumps, and drinking water supplied through lead pipes. But one question arises whenever the map shows a trend or cluster: Is the pattern real?

The problem is one of small numbers. Pandemics are rare, and seldom is the association between disease and an environmental cause so overwhelming that the link is easily identified and unchallenged. Clusters of deaths or diagnosed cases usually are few and unspectacularly small, perhaps no more than three deaths in a town or two in the same neighborhood. Epidemiologists map these cases both as points, to get a sense of patterning, and by areal units, to adjust for spatial differences in the number of people at risk. After all, an area with half the region's cases is not remarkable if it has half the region's population. But what is the significance of a small area with two or three cases and a rate several times above the national or regional rate? Could this pattern have arisen by chance? Would one or two fewer cases make the area no longer a "hot spot"? If one more case were to occur elsewhere, would this other area also have a high rate? To what extent does the pattern of high rates reflect arbitrary boundaries, drawn in the last century to promote efficient government or decades ago to expedite delivery of mail? Might another partitioning of the region yield a markedly different pattern? Might another level of aggregation—larger units or smaller units—alter the pattern? Is the mapping method inflating the significance of some clusters? And is it possibly hiding others?

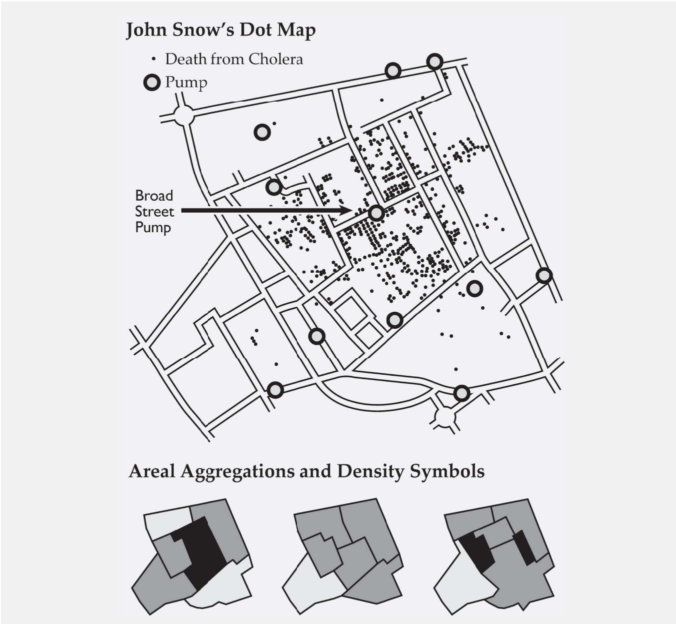

Figure 11.18. A reconstruction of John Snow's famous dot map of cholera (above) and three choropleth maps (below) produced by different areal aggregations of this part of London.

Consider, for example, the maps in figure 11.18. At the top is a reconstruction of John Snow's famous map showing cholera deaths clustered around the Broad Street pump. A physician working in London during the cholera epidemic of 1854, Snow suspected drinking water as the source of infection. At that time homes did not have running water, and people carried buckets from a nearby pump. According to legend, Snow's map confirmed the waterborne transmission of cholera, and when authorities removed the pump's handle, new cases in this part of the city plummeted. Truth be told, the epidemic had run its course, and Snow made his map months later when he revised his book on cholera.

But what might have happened had Snow not worked with point data? The three maps at the bottom of figure 11.18 show how var-

ious schemes of areal aggregation might have diluted the Broad Street cluster. If addresses are available, as on most death certificates, aggregation to census tracts or other areal units larger than the city block increases the risk of missing intense, highly local clusters.

Aggregation involves not only areal units but also time, disease classification, and demography. One solution to the question of significance is to get more data by collecting information over a longer time span. Adding together several years of data, or even several decades, dampens the effect of chance occurrences but risks involving a wider range of causal agents. Aggregation over time might, for instance, mask important temporal trends, dilute the impact of new or abated environmental contaminants, or incorporate difficult-to-measure effects of population mobility. Likewise, combining several disease categories or the mortalities of diverse demographic groups promotes stability and significance by increasing the number of cases and broadening the set of causes.

Clearly one map is not sufficient, although one good map can signal the need for a more detailed investigation. It is then up to a variety of scientific researchers to explore further the effects of geography and environment by examining employment and residential histories, characteristics of residence and neighborhood, and hereditary factors; by carefully studying maps at various levels of spatial, temporal, and demographic aggregation; through computer simulation to test the stability of known clusters; through automated pattern recognition to identify new clusters; and through related clinical and laboratory studies. Although maps can indeed lie, they can also hold vital clues for the medical detective.

Indexes, Rates, and Rates of Change

Another danger of one-map solutions is a set of measurements that presents an unduly positive or negative view. Often the map author has a single theme in mind and has several variables to choose from. Usually some variables are markedly more optimistic in tone or pattern than others, and the name of the index can cast a favorable or unfavorable impression in the map title. "Labor Force Participation," for instance, sounds optimistic, whereas "Job Losses" clearly is a pessimist's term. An appropriately brazen title offers a good way to overstate economic health or industrial illness.

If the picture is bleaker or brighter than suits your politics, try a rate of change rather than a mere rate. After all, minor downturns often interrupt a run of good years, and depressions do not last forever. If unemployment is high now but a bit lower than a year, six months, or a month ago, the optimist in power would want a map showing a significant number of areas with declining unemployment. Conversely, the pessimist who is out of power will want a map depicting conditions at least as bad as before the current scoundrels took over. A time interval that begins when proportionately fewer people were out of work will make the opposition party's point, especially if unemployment has become worse in large, visually prominent, mostly rural regions.

A useful index for the optimist is one with relatively low values, such as the unemployment rate, if conditions have improved, or an index with comparatively high values, such as employment level, if conditions are worse. Thus a drop of one percentage point from a base of 4 percent unemployment yields an impressive 25 percent improvement! Yet a substantial increase in the unemployment rate from 4 to 6 percent can be viewed more optimistically as a drop in labor force participation from 96 to 94 percent—a mere 2 percent drop in employment.

Point symbols and counts, rather than rates, can be useful too. If the economy has been improving in all regions, the current government might want a map with graduated circles or bars showing actual counts beneath the title "Employment Gains." If the country is in a widespread recession, the opposition would use similar point symbols with the title "New Job Losses."

The cartographic propagandist is also sensitive to spatial patterns. Favorable symbols should be large and prominent, and unfa-

Area	Labor Force (000s)	Unemployment (000s)			Unemployment Rate			Percentage Change
		t_1	t_2	change	t_1	t_2	change	
1	3,000	120	180	+60	4.0%	6.0%	+2.0%	+50.0%
2	16,000	640	800	+160	4.0%	5.0%	+1.0%	+25.0%
3	2,500	125	113	-12	5.0%	4.5%	-0.5%	-9.6%
4	800	56	48	-8	7.0%	6.0%	-1.0%	-14.3%
5	500	40	35	-5	8.0%	7.0%	-1.0%	-12.5%

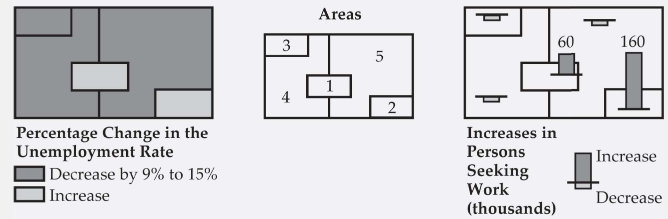

Figure 11.19. Unemployment data (top) for a hypothetical region (bottom center) yield different maps, supporting an optimistic view (bottom left) and a pessimistic view (bottom right) of recent temporal trends.

vorable ones small and indistinct. Thus the optimist might present the unemployment data in figure 11.19 with the map at the lower left, to focus attention on improved conditions in larger areas, whereas the pessimist would prefer the map at the lower right, to emphasize the much greater number of unemployed persons in more urban areas. Note as well how the titles and keys in these examples reinforce cartographic manipulation.

Labor economists, who commonly adjust unemployment data for seasonal effects, discourage some manipulation of time intervals. After all, more people are seeking work in early summer, when many high-school and college graduates enter the labor force for the first time. And more people find at least temporary work in November and December, the peak holiday shopping season. Local seasonal effects, such as tourism and the temporary hiring of field and cannery workers in agricultural areas, also require seasonal adjustment.

Mortality, fertility, and other phenomena that do not affect all segments of the population equally also require adjustment. Figure 11.20, a comparison of the age-adjusted death rate with the crude

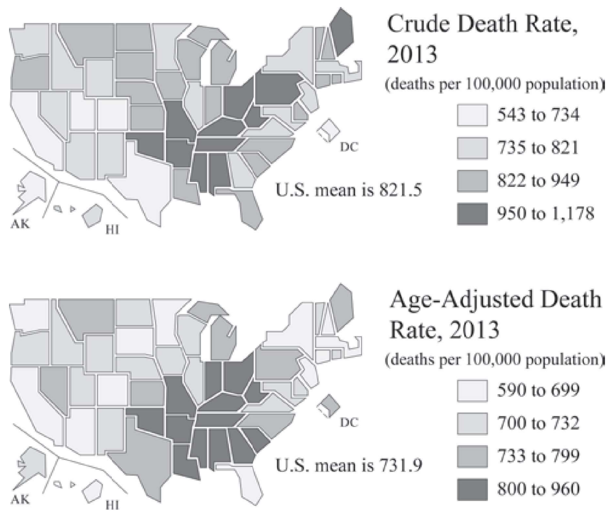

Figure 11.20. Maps of the crude death rate (top) and the age-adjusted death rate (bottom) can present markedly different geographic patterns of mortality.

death rate, illustrates the wisdom of mapping demographically adjusted rates. The map at the top is a simple rate, which does not consider, for example, Maine's relatively older population. When the rates portrayed in the upper map are adjusted for age differences, Indiana and several southeastern states emerge as high-rate areas whereas those in the Northeast slip to a lower category. Age-adjustment allows the map at the bottom to reveal the effects of relatively good health care and a higher socioeconomic status in the New England, Middle Atlantic, and North Central states, in contrast to greater poverty and less accessible health care in the South.

Be skeptical of maps based on numbers. Because a single variable can yield many different maps, don't be silenced by the argument that more than one map would cause needless confusion. You might ask to see several maps, or be given the opportunity to experiment with categories and symbolization online or using software.

Be wary of not only the known cartographic manipulator but also the careless map author unaware of the effects of aggregation and classification. Also question the definitions, measurements, shortcuts, and motives of a government agency, research institute, or polling firm that generously provides its data—even the most conscientious mapping effort is undermined by flawed data.