U.S. Demographic Quantities Visualized by Cartographic Treemaps

William H. Press
The University of Texas at Austin

December 25, 2010

1 Introduction

This is a paper about politics. However, to shake off all but the most intrepid politically-minded readers, I begin with some background on cartograms, some mathematics about maps, and an introduction to so-called treemaps. If you are merely political, and neither intrepid nor mathematical, you can skip all the way to Section 5 and just start looking at the pictures.

Cartograms are maps that show statistical information in a diagrammatic form. A simple shaded map, where the intensity of shading represents some quantity that depends on location, can be termed a cartogram. An early, now disturbing, example, is the U.S. Coast & Geodetic Survey's 1861 map of the slave population of the United States, shown in Figure 1 (see also reference [1]).

As now used, however, the term cartogram usually means a *distorted map*, a map whose scale is locally altered so that areas represent not geographical area, but instead some other quantity of interest.

Ideally, a distorted map would alter *only* the local map scale, preserving the shapes of all regions, that is, the topology of boundaries and the angles at which boundaries meet. However this ideal is not mathematially possible. To see this, suppose that (x_1, x_2) is a Cartesian coordinate system over the original, undistorted, map. Then describe the distorted map by its Cartesian coordinates (y_1, y_2), with

$$y_1 \equiv y_1(x_1, x_2), \qquad y_2 \equiv y_2(x_1, x_2) \tag{1}$$

The four scalar quantities in the Jacobian matrix

$$\mathbf{J}(x_1, x_2) = \begin{pmatrix} \frac{\partial y_1}{\partial x_1} & \frac{\partial y_1}{\partial x_2} \\ \frac{\partial y_2}{\partial x_1} & \frac{\partial y_2}{\partial x_2} \end{pmatrix} \tag{2}$$

can be decomposed into one parameter of magnification, μ; one parameter of rotation, θ; and two parameters of shear, $\sigma_{1,2}$ [2].

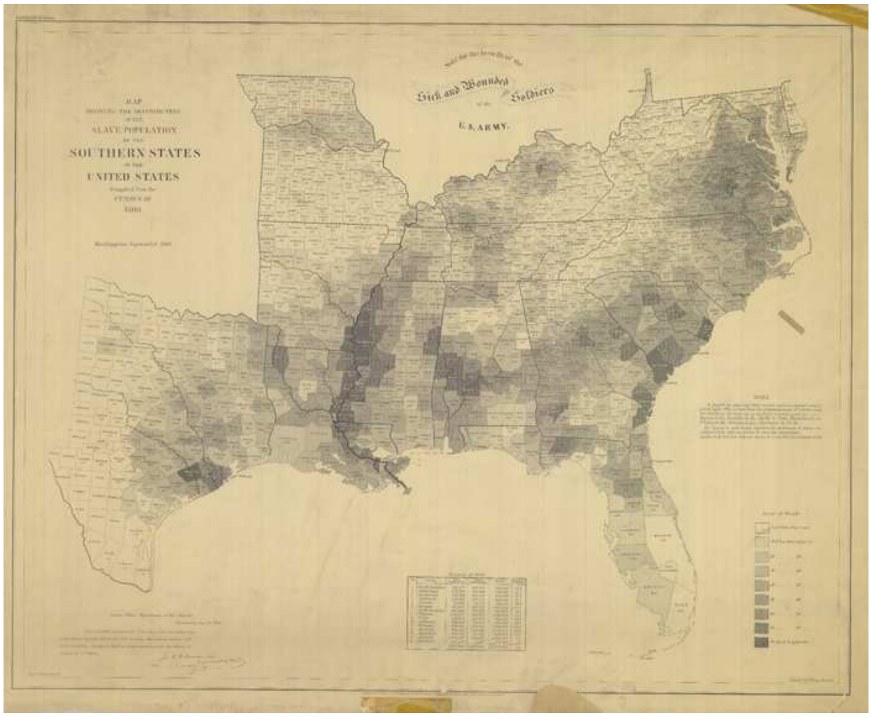

Figure 1: On an geographical base map, this 1861 map shows U.S. slave population density as intensity of shading, an early form of cartogram. (map source: NARA C&GS Collection)

In constructing a desired distorted map we thus have two free functions (the two equations (1)) minus one constraint specifying the desired magnification,

$$\det[\mathbf{J}(x_1, x_2)] = \mu(x_1, x_2) \tag{3}$$

for some specified $\mu(x_1, x_2)$ under our control. We thus have one net functional degree of freedom remaining. Can we use it to set the shear to zero? If so, we would have a conformal map [3] that preserves angles and thus, in general terms, the shapes of geographical areas. No, we can't, because there are *two* components of shear. To zero both would require that both Cauchy-Riemann equations be satisfied,

$$\frac{\partial y_1}{\partial x_1} = \frac{\partial y_2}{\partial x_2}, \qquad \frac{\partial y_2}{\partial x_1} = -\frac{\partial y_1}{\partial x_2} \tag{4}$$

It is thus impossible to preserve shapes while specifying magnifications, and, thus, all cartograms must be compromises of one kind or another. This is entirely analogous to the well-known fact that there is no map projection from the globe to the page that preserves both areas and angles, although one can choose to preserve one or the other.[4]

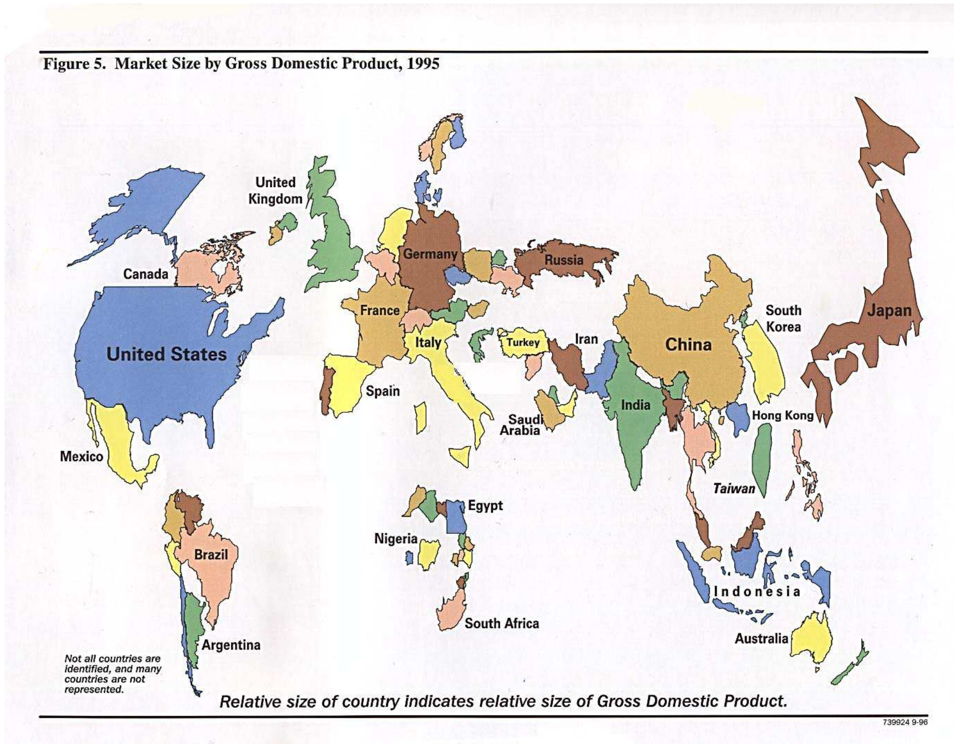

Figure 2: World GDP, preserving shape, but not connectivity. New oceans are created where needed. (map source: Handbook of International Economic Statistics, 1996)

2 World Cartograms

Before turning to the Unites States as our main subject, let us look at the world. The world is about 70% oceans. The oceans provide many benefits, including one to would-be constructors of cartograms: Their audience rarely cares about the shape or size of the oceans, so they are free to let distorted land masses encroach on the oceans at will. Figure 2 illustrates this in the extreme. Here, in a cartogram showing gross domestic product (GDP), every country's shape has been preserved perfectly (within the limits of a flat map), simply by the device of creating new bits of fictitous ocean wherever needed! The connectivity or topology of national boundaries has of course been lost.

If we want to do a better job of preserving topology, we must in effect use our remaining one degree of freedom to find an approximate compromise of the two Cauchy-Riemann equations. Frequently this is done by intutition and craftsmanship rather than by mathematics. Figure 3 is a cartogram showing world population, where each country is represented as a connected set of Cartesian

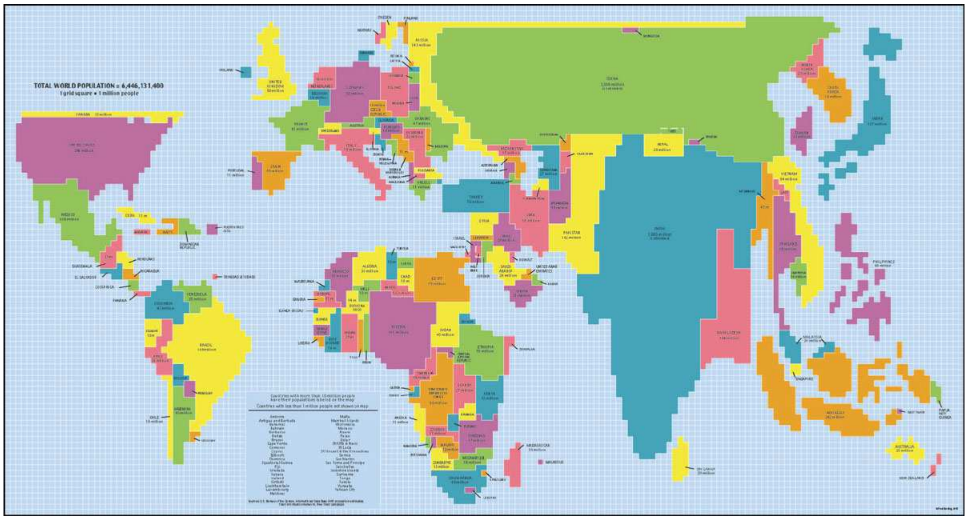

Figure 3: World population, with an artistic compromise between shape and topology. (map source: Paul Breding et al., www.odt.org)

grid squares in the appropriate integer number. The shapes and positions of the countries have been artistically arranged to meet the viewer's expectations of a world map. As in Figure 2, a lot of bad karma is absorbed into the oceans. Thus, landlocked countries, for example in Eastern Europe, don't fare as well, shapewise, as island nations, for example Australia and Indonesia. There are also cultural biases, for example preserving the shape of the U.S. at the expense of that of Canada. (Canada's shape could have been better preserved at the expense of the U.S. by "pinching" the northern U.S. border to a much smaller width.) Nevertheless, assuming an intended U.S. audience, Figure 3 is a remarkable piece of cartographic craftsmanship.

We would like a way to proceed by mathematics, rather than by craftsmanship. Gastner and Newman [5] review a number of previous mathematical approaches and propose an elegant way of mathematizing, and thus automating, some of the general notions that are only intuitive in Figure 3. Roughly, they take a geographical map as a starting point, and imagine that each country contains a specified quantity of "stuff". Thus, some countries have a higher density of stuff than others. With this starting point, the time-dependent diffusion equation is now applied. After some time, an equilibrium will have been reached where all countries have the same density of stuff. Therefore, the areas of countries are now equal to the original quantities of stuff assigned. An important point is that the oceans must be assigned some quantity of stuff *in toto*; otherwise land mass diffusion would continue until the oceans disappeared. But since the ocean's stuff is shared across all oceans, the shapes of the oceans may vary freely.

Figure 4 shows a cartogram of world population (compare Figure 3) produced

Figure 4: Population by country, diffusion method of Gastner and Newman. (map source: Mark Newman, University of Michigan)

by the Gastner-Newman algorithm. The topology of national boundaries is exactly preserved. Shapes are best preserved where the gradient of population density (which drives the diffusion) is not too large, but are distorted where the density gradient is initially large (as around Nigeria), or where a country has "ballooned" into the ocean (as India). Note how Canada, with low population density, has become "stringy", but not one-dimensional.

None of the techniques discussed thus far are entirely satisfactory for displaying U.S. demographics to an American audience. The main problem is that we are trained from elementary school to an outline of the U.S. and its states (shown in Figure 5) that, when distorted, quickly loses geographical meaning. We lose our bearings. (Doubtless the Frenchman has a similar conditioning to the shape of France, the Aussie to Australia, and so on.) As an example, Figure 6 shows Gastner and Newman's cartogram of U.S. population by state (also colored by results of the 2000 U.S. Presidential election). It is an elegant piece of cartography, but somehow not very, well, American! To put it less emotionally, the distortions in Figure 6, especially the loss of familiar outline shape of the U.S., distract the viewer from the data presented.

3 Introduction to Treemaps

Turn now to so-called *treemaps*. Treemaps [6] are pictorial representations that, like cartograms, represent a set of data values as corresponding areas on a two-dimensional map. However, the underlying data that is shown in treemaps are not usually geographical. While cartograms have the underlying structure of a full two-dimensional topology, the underlying data on a treemap has only

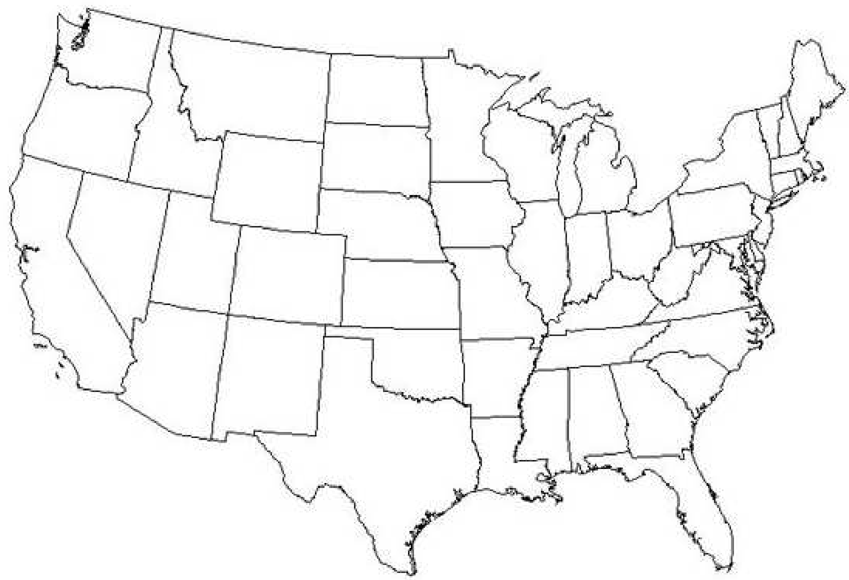

Figure 5: Outline map of the United States, typical map projection. Most Americans are drilled on this image from primary school onwards.

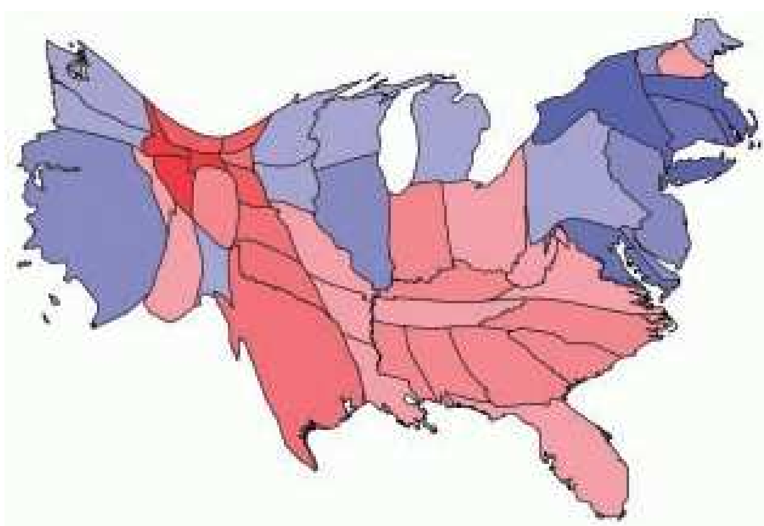

Figure 6: Population by state, diffusion method of Gastner and Newman, colored by 2000 Presidential election results. Although the map is locally undistorted, its global distortion is unappealing and distracts from the data. (source: [5])

Figure 7: Australian web shopping, 2009, as a treemap. (source: Sandra Hanchard, Hitwise Australia)

a branching tree structure. That is, a data point (a "country") is contained in a "parent", and contains "children", but no other geometrical relationship is assumed. Whereas the underlying structure of a cartogram is too rich to be represented without distortion, the underlying structure of a treemap is so spare that it can be represented perfectly by areas, and in multiple ways.

An example, Figure 7, will make things clearer. The map represents, by areas, the proportion of web visits to various types of shopping sites. You can figure out what is the underlying tree by looking at the structure of horizontal and vertical lines, starting with the longest ones. For example, the vertical lines to the left and right of Rewards/Computers can be seen to separate sibling nodes at the highest level, while the horizontal line between Rewards and Computers divides a node at the next level. Continuing in this fashion, one reconstructs the full tree, as listed in Figure 8 (with indented lines representing child nodes of the preceding less-indented parent).

In the case of Figure 7, the topology of the tree was evidently chosen not by the meaning of the data items, but by an algorithm that puts smaller areas deeper in the tree. The purpose is to generate a treemap with the largest blocks towards the upper left, smallest towards the lower right, and with the shapes of the blocks as close to square as possible. Several such tiling schemes have been proposed, [6] embodying different tradeoffs among order (how accurately the order of the data is preserved), aspect ratios (how square are the boxes), and stability (how drastic are the rearrangements if the areas are perturbed slightly).

Often, the higher levels of a treemap are chosen to partition by a logical

```
root
    Auctions
    node
        Rewards
        Computers
    node
        node
            Department Stores
            Classifieds
            Electronics
        node
            node
                Apparel
                House and Garden
            node
                node
                    Video and Games
                    Music
                    Sports
                node
                    node
                        node
                            Toys and Hobbies
                            Books
                        node
                            Ticketing
                            Health and Beauty
                    node
                        node
                            Automotive
                            Wholesale
                        node
                            node
                                Grocery
                                Flowers
                            node
                                Office Supplies
                                Intimate Apparel
```

Figure 8: The tree underlying Figure 7 can be reconstructed from that Figure's pattern of horizontal and vertical divides.

structure in the data, while lower levels may default to a scheme such as that of the preceding figures. Figure 9, for example, is a treemap of the S&P 100 corporations, where size represents market capitalization. The upper levels of the tree encode industry sector, while the lower tree levels are arranged to put larger companies towards the upper left within each sector. Color can be used to represent up or down trends. A popular, interactive treemap along these lines is "Map of the Market," originally designed by Martin Wattenberg.[7]

4 Cartographic Treemaps

The main idea of this paper is that we can define a tree whose nodes are the U.S. states in such a way that the resulting treemap gives an immediately recognizable, if stylized, representation of a U.S. outline map (Figure 5) over a large range of areas assigned to individual states. Figure 10 illustrates this idea as a treemap where each state is assigned its actual land area.

It is worth noting that, analogously with the artistic Figure 3, there are cultural as well as cartographic choices involved in picking a tree topology. For

Figure 9: Treemap of the S&P 100 companies by capitalization. The upper tree levels partition by industry sector. (source: unknown)

example, the tree underlying Figure 10 honors the standard groupings "East (West) of the Mississippi", "Pacific States", "Intermountain West", "New England", and (with some liberties) "North (South) of the Mason-Dixon line", "Deep South", "Upper Midwest", and "Rust Belt". Constraints like these tend to improve the fidelity of geographical boundaries over a wide range of possible cartograms, since states with similar characteristics will tend to remain connected. However, already in Figure 10, some disparities are evident. For example, Louisiana does not actually have a border with Tennessee; nor does New York actually border Michigan – Ontario and three Great Lakes (can you name them?) are factually interposed. And don't plan to visit Four Corners in Figure 10. Still, for so simple a data structure and so stylized a representation, the overall geographic fidelity seems pretty good.

All of the U.S. cartographic treemaps in the rest of this paper have the same underlying tree as Figure 10. Three different types of nodes are used: North-South Stacked, West-East Stacked, and Northeast Corner (this approximating the relationship of Nevada to California, Rhode Island to Connecticut, and DC to Virginia). Hawaii, Alaska, and West Virginia are special cases with their own node types. California, Nevada, Texas, and Florida are cosmetically altered when drawn, but are actually represented by underlying rectangles (with areas adjusted so that the as-drawn map is correct). One could easily do a better job of this for these states, and similarly for Alaska, Hawaii, Maine, and Massachusetts (Cape Cod). The program that generated all of the following figures was written in C++ to produce simple PostScript graphical output. A more complicated implementation could produce better looking results.

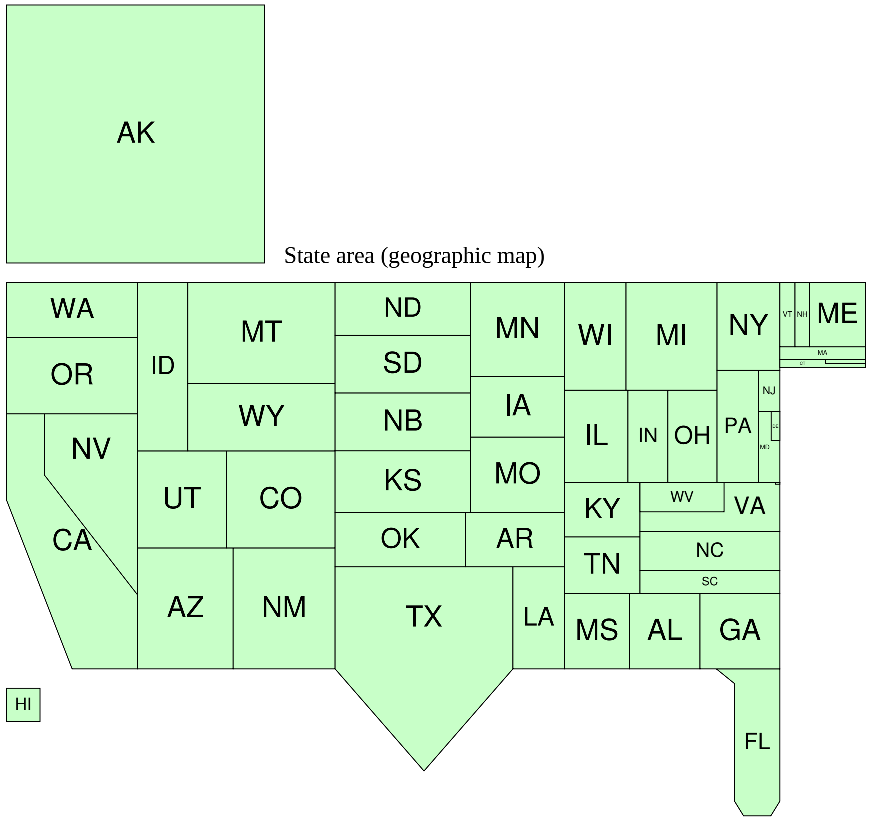

Figure 10: Cartographic treemap of the United States, preserving true geographic areas of the states.

5 U.S. Demographic Data

5.1 Population, House and Senate Representation

While Figure 10 shows the states in proportion to their true geographical area, the top map in Figure 11, by contrast, shows states according to population. It should be compared to the diffusion method's Figure 6. Both maps are of course highly distorted. Figure 6 is more map-like, while Figure 11 is more diagrammatic, but also arguably easy to read.

The bottom map in Figure 11 shows each state with an identical area, analogously to how each state is represented in the U.S. Senate. A lot of U.S. politics is implicit in the two maps in Figure 11!

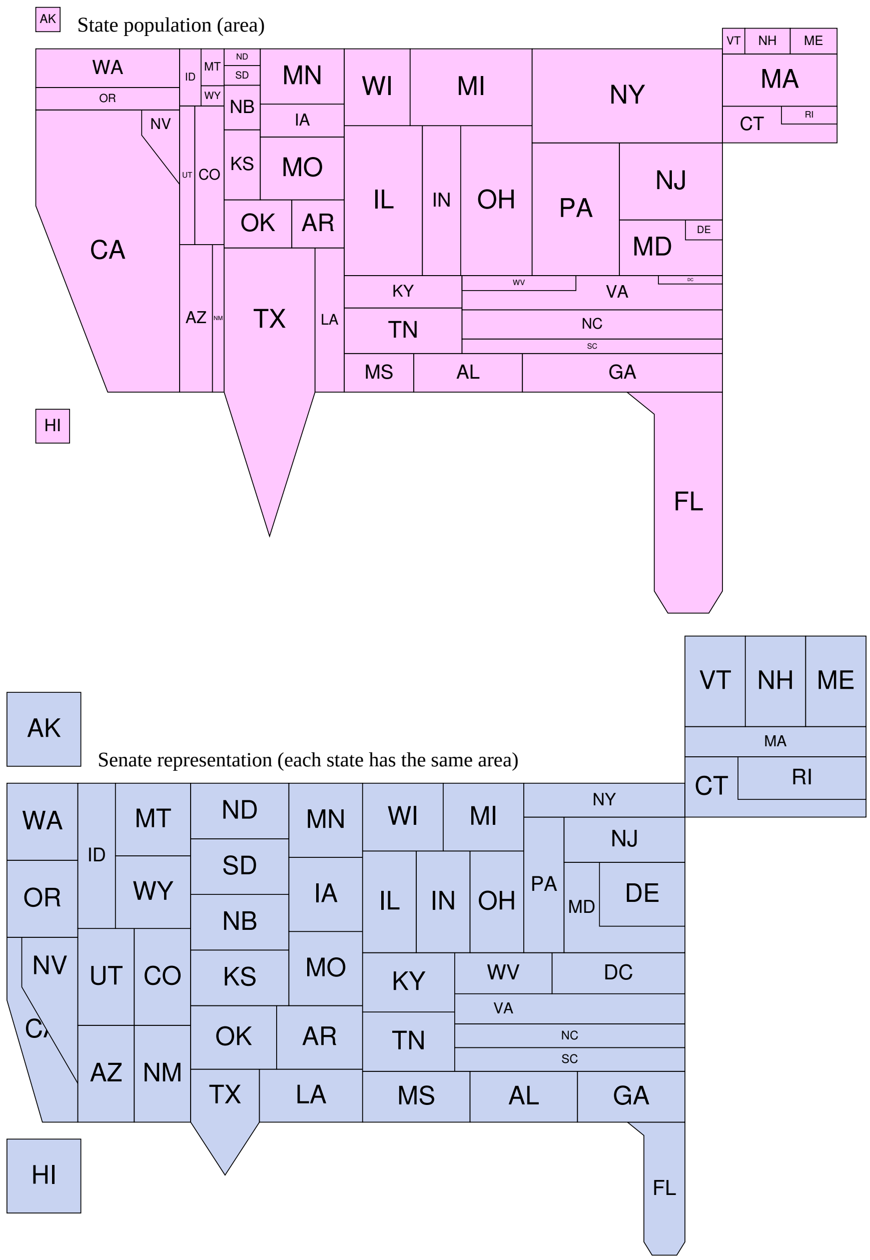

Figure 11: Top: States by population (2000 Census), approximately representation in U.S. House of Representatives. Bottom: Each state is plotted with the same area, as in proportion to its representation in the U.S. Senate.

5.2 Red and Blue Political Maps

For many kinds of data, the treemap with area proportional to population is the logical base map, on top of which we can use color to show additional quantities of interest. For example, in Figure 12) we show the 2010 midterm election swing from a Democrat (blue) majority to a Republican (red) majority in the U.S. House of Representatives, by coloring the population map on a red-blue intensity scale according to the fraction of a state's Congressional districts represented by each party before and after the election.

One can learn a lot about what happened in the 2010 election from Figure 12. In essence, the Midwest turned from blue to red. Harder to grasp in this Figure, however, are the actual numerics by which the Democrats lost their House majority. The problem is that the eye does not easily translate color intensities into integer districts gained or lost. For this purpose, we need a cartogram that shows states in proportion to the Republican swing, i.e., shows where, at the margin, the new Republican majority came from. This is shown in Figure 13.

One sees the central importance of the Rust Belt states. Secondary, but important, are the more traditionally Democratic southern states of Arkansas, Tennessee, and Floria. Some states (Utah, Delaware, most of New England), which had no Republican swing, are omitted from the map completely. The omission of states is readily accommodated by the treemap method. (No states in the 2010 election had a Democratic swing.)

Figure 13 uses area, not shading, to quantify the swing (as measured by number of Congressional districts). Shading is used for something else entirely. That usage is the second main idea of this paper and requires a bit of explanation. North Dakota and New Jersey each had a swing of one district, as indicated by their area. North Dakota is darkly shaded because its swing represents 100% of its delegation of one, Rick Berg (R) defeating Earl Pomeroy (D). New Jersey is lightly shaded because its swing was one out of 13 Congressional districts, that is, less than 8%. We call this use of shading *per capita shading* because (in the common case that the base map represents population) it indicates not total quantity, but quantity *per capita*.

(If you find the diffusion-method map format more appealing than the treemap format, see [8] and [9] for some beautiful representations of the 2008 and 2010 elections, respectively.)

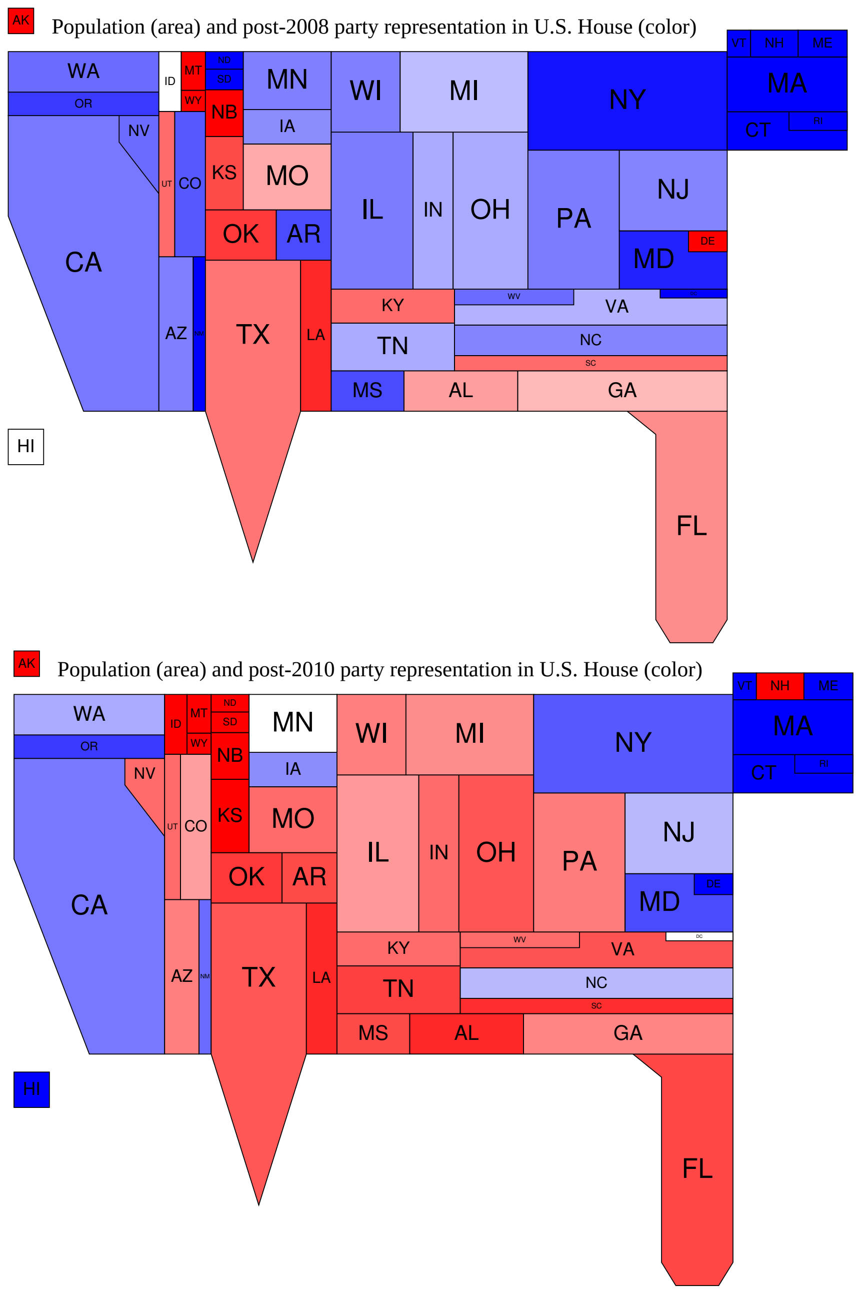

Figure 12: U.S. House representation before (top) and after (bottom) the 2010 mid-term election. States are shown proportional to their population and colored in intensity according to their proportion of Democrat (blue) vs. Republican (red) Congressional districts.

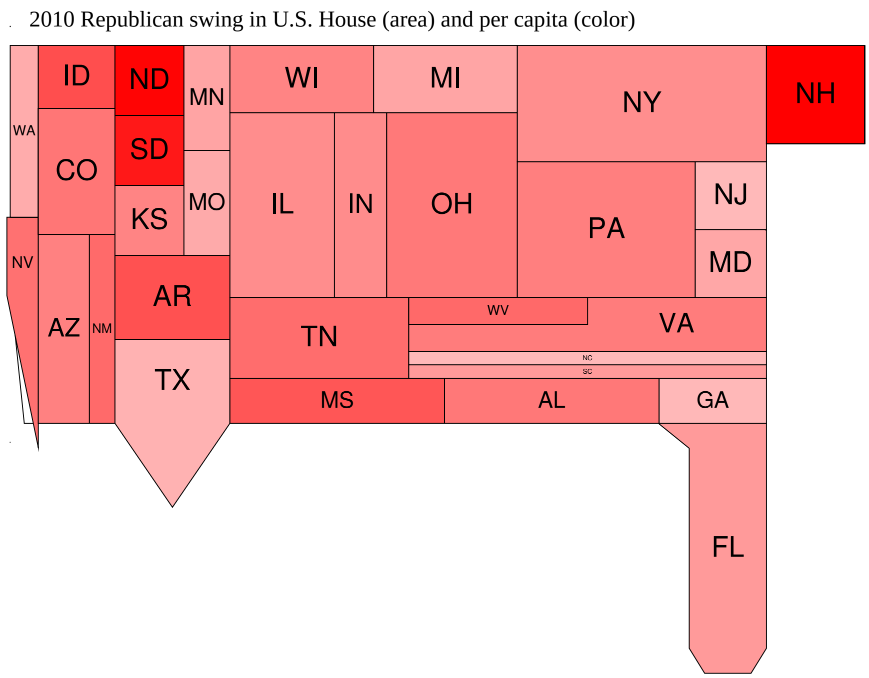

Figure 13: Swing to Republican U.S. Representatives in the 2010 mid-term election. Area indicates the number of Congressional districts, while intensity of shading indicates the fraction of the state's delegation that swung.

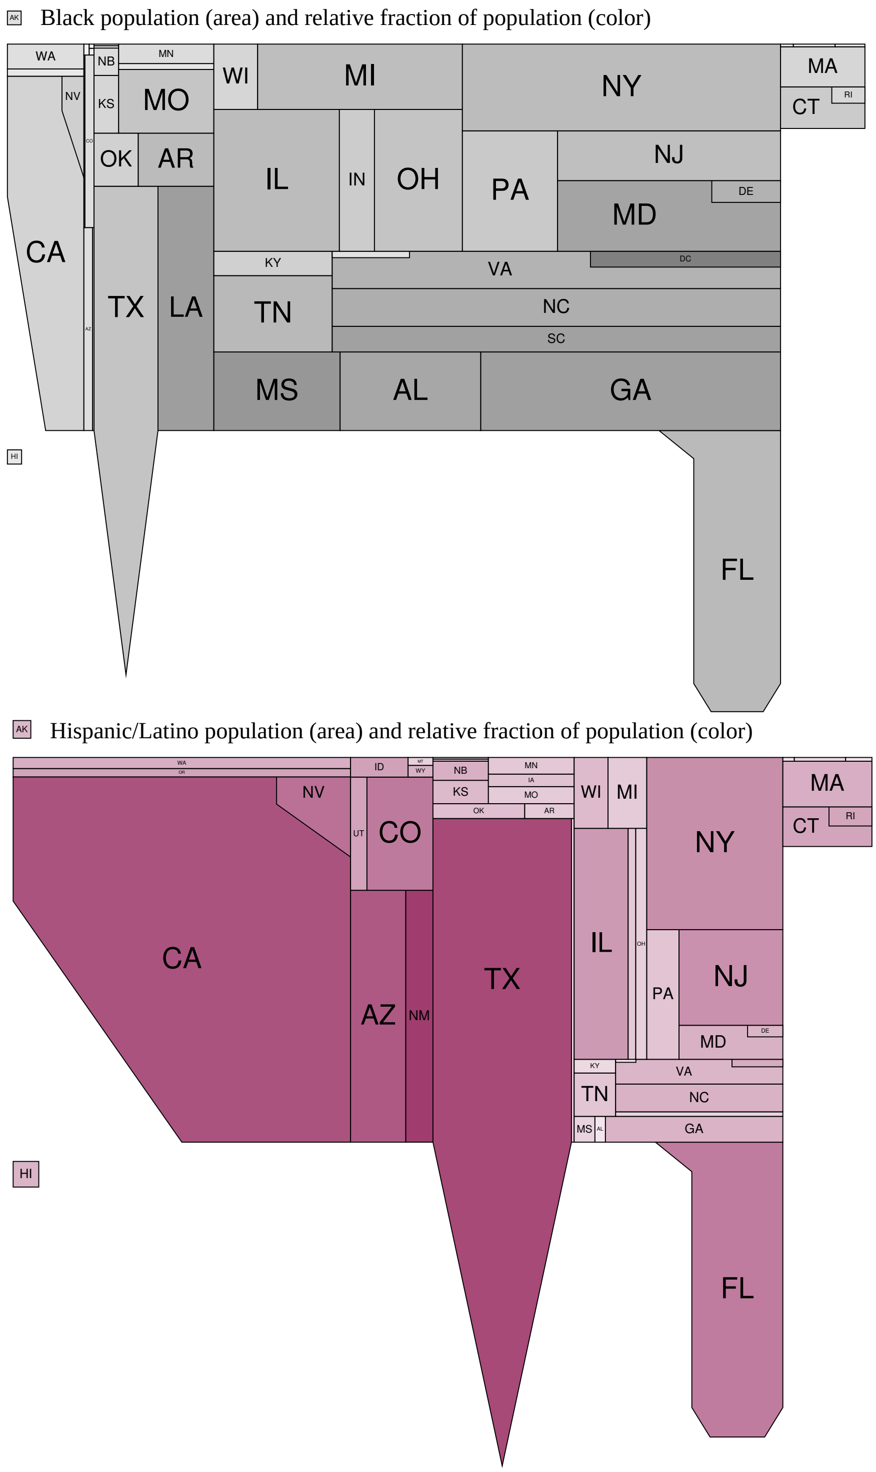

Figure 14: Top: State areas correspond to 2004 total Black population. Shading represents the relative density of the Black population within the state's total poulation. Bottom: Same, for Hispanic/Latino population. (Data source: www.statemaster.com)

5.3 Minority Population Demographics

For examples in which per capita shading is useful, we turn to demographic statistics on minority populations in the U.S. Figure 14 shows the geographical distribution of, respectively, African-Americans and Hispanics/Latinos (2004 data). The combination of cartographic treemap and shading enable the data to be grasped in various ways: The Intermountain West is in effect nonexistent for the Black population, as is the Upper Midwest and Pacific Northwest for Hispanics. Texas and California have both large, and fractionally large, Hispanic populations. New Mexico's Hispanic/Latino population is fractionally large, but numerically smaller than Illinois'. And so on. It is interesting to compare Figure 14, top, to Figure 1, 150 years earlier, a cartographic juxtaposition that compresses into two snapshots the suffering of slavery, the Civil War, and the migrations of African-Americans to cities of the north, but not (with Calfiornia as the sole exception) the west.

Two smaller groups, the Native American and Jewish populations, are respectively shown in Figure 15. That North Carolina's Native American population rivals that of New Mexico and Oklahoma is perhaps a surprise to many, although North Carolina's population density of Native Americans hardly rivals that of New Mexico. Alaska, normally only a blip for most demographics, is comparable to South Dakota by numbers, and comparable to New Mexico by density. While the overall Native American population is fractionally small in the U.S., it is striking that there are no "missing" regions. In virtually every state, people identify themselves as Native American.

For the Jewish population, New York is unsurpassed in both number and density, while the South and Intermountain West (apart from Arizona and Colorado) hardly exist. One must look hard to find any trace of Montana, Wyoming, or the Dakotas on this map; Mississipi and Arkansas also fare badly in this demographic. The cartogram uncannily resembles Saul Steinberg's famous 1976 New Yorker cover (Figure 16), perhaps not completely coincidentally. It should be noted that the U.S. Census does not ask people about their religion, so that the geographical distribution of Jews (as also Catholics, Methodists, or any other denomination) is known only from estimates by nongovernmental organizations, usually those associated with the corresponding denomination.

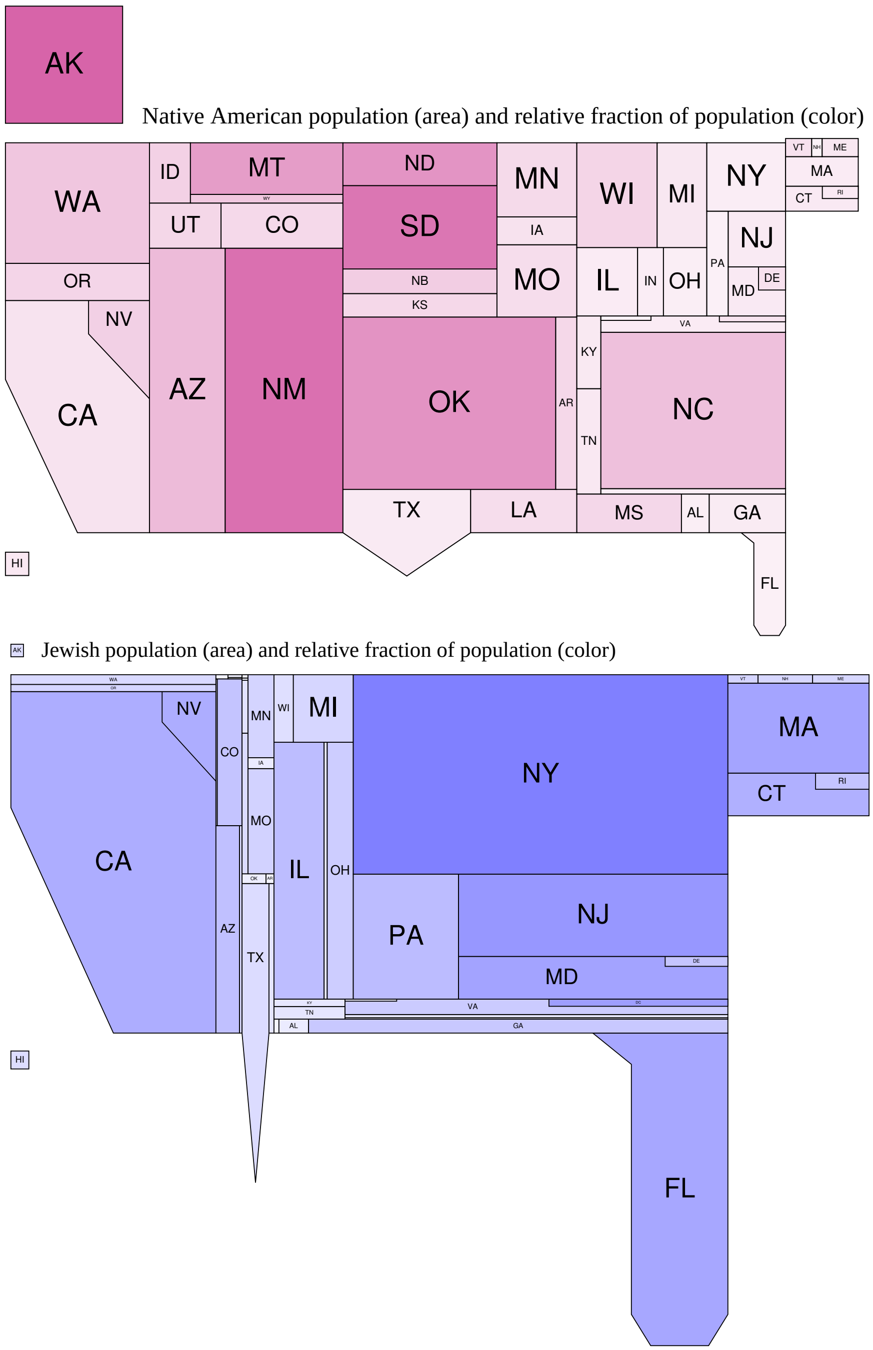

Figure 15: Top: State areas correspond to Native American population, with shading proportional to density in the total population. (data source: www.statemaster.com) Bottom: Same, for Jewish population. (data source: www.jewishvirtuallibrary.org)

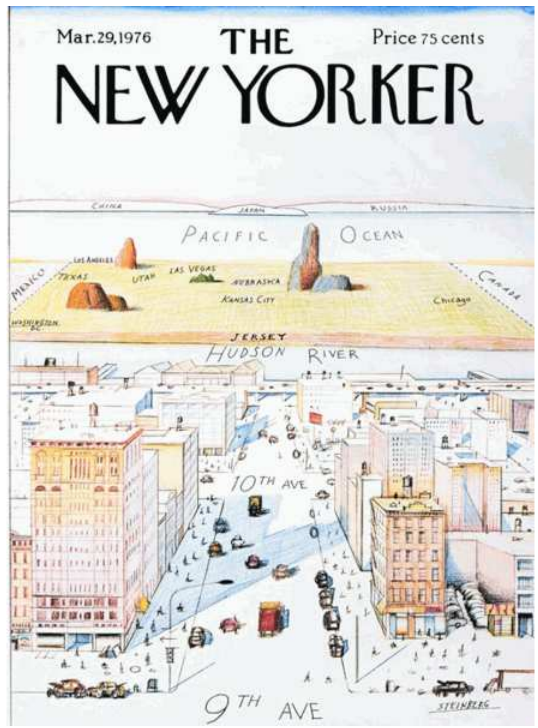

Figure 16: Cover of The New Yorker, March 29, 1976. (Copyright The Saul Steinberg Foundation/Artists Rights Society)

5.4 Gross Domestic Product (GDP)

In a society with perfect income equality, and perfect equality of productivity, GDP would perfectly follow the population, and a cartogram of GDP would be identical to that of population, Figure 11 (top). Actually, the gross deviations from this ideal are not, graphically, as great as one might suppose. Excluding the five top and five bottom states, GDP per capita in the U.S. varies only between about $30,000 and $40,000 (2005 data). As important as this difference is to a struggling family, it is not a large enough range to be immediately evident in the corresponding cartogram, Figure 17. Indeed, one learns more from the shading by GDP per capita (which could have been applied to any kind of base map) than from the cartographic treemap itself. Things would be quite different on a world map, where GDP per capita varies over more than two orders of magnitude.

The cartographic treemap (or other area-distorted cartogram) shows itself as more useful when we compare different industry or economic sectors. Figure 18 shows cartographic treemaps for GDP from manufacturing and from information industries, the latter category including broadcasting, publishing, motion pictures, data processing, and computer-related industries.

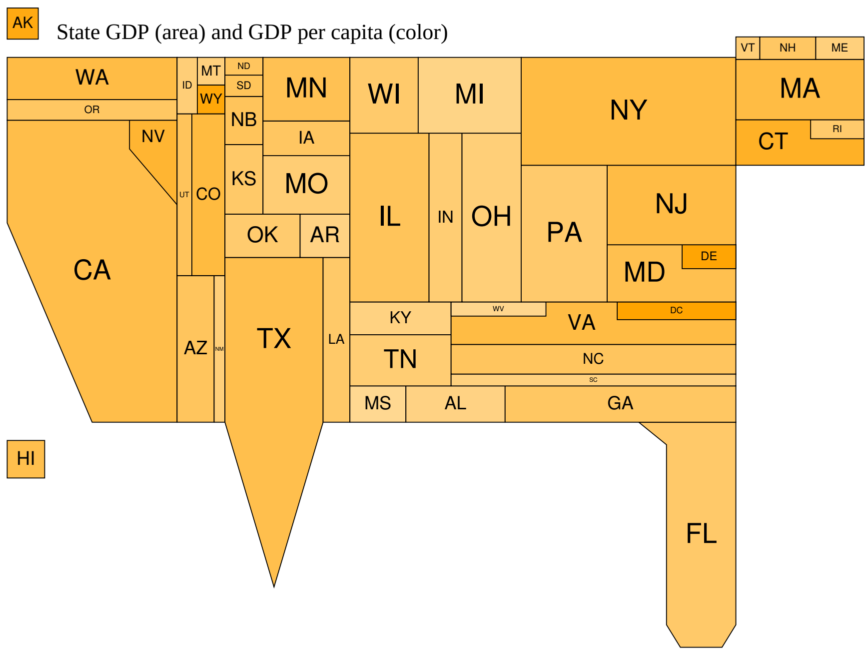

Figure 17: A state's area is here proportional to its Gross Domestic Product (GDP). Shading intensity shows GDP per capita. This map generally follows the distribution of population, Figure 11, with only subtle differences. (data source: U.S. Department of Commerce, Bureau of Economic Analysis, 2005)

The map of manufacturing is dominated by the Rust Belt states plus Wisconsin and Michigan, together with the Middle South, Kentucky, Tennessee, and North Carolina. By contrast, the information industries are heavily "bicoastal", plus Colorado, Texas, and Illinois as additional hubs. If we excluded the contributions of local media, which generally scale with population, the effect would be even more pronounced. (Saul Steinberg again?)

The District of Columbia is noteworthy in Figure 18 as producing a lot of (dis?) information for its size, but not much in the way of tangible goods.

Another interesting comparison is between GDP from legal services (Figure 19, top) and GDP from finance and insurance (Figure 19, bottom). There are a lot of lawyers in the District of Columbia! For legal services, the pattern is largely bicoastal, plus Texas and Illinois. While generally similar, the pattern for finance and insurance has a few notable differences. Delaware's role in banking and corporate finance stands out, as does the insurance business in Iowa and the Midwest generally, states whose residents are known for their probity.

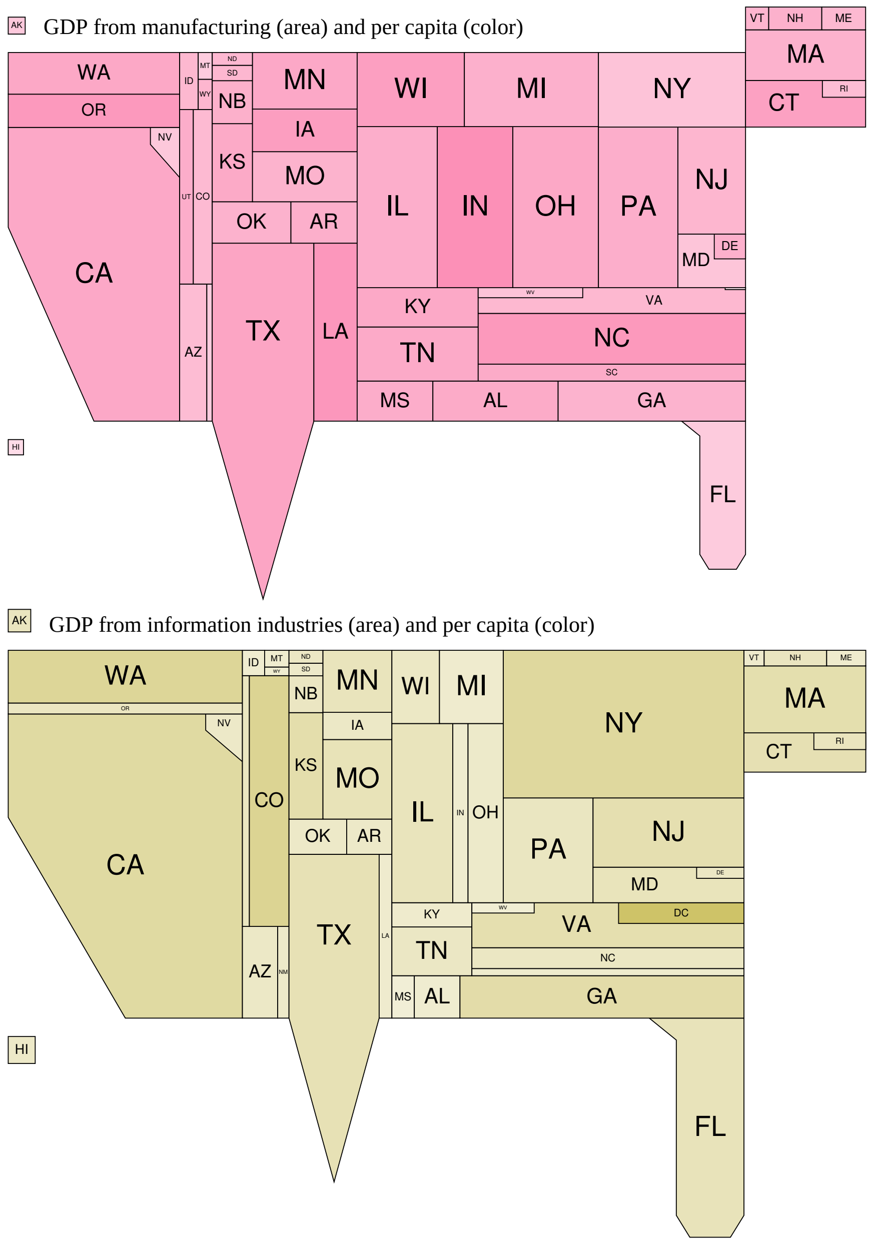

Figure 18: Top: GDP from manufacturing industries (area) and GDP per capita (shading). Bottom: Same, for information industries, including broadcasting, publishing, and motion pictures. (data source: www.statemaster.com)

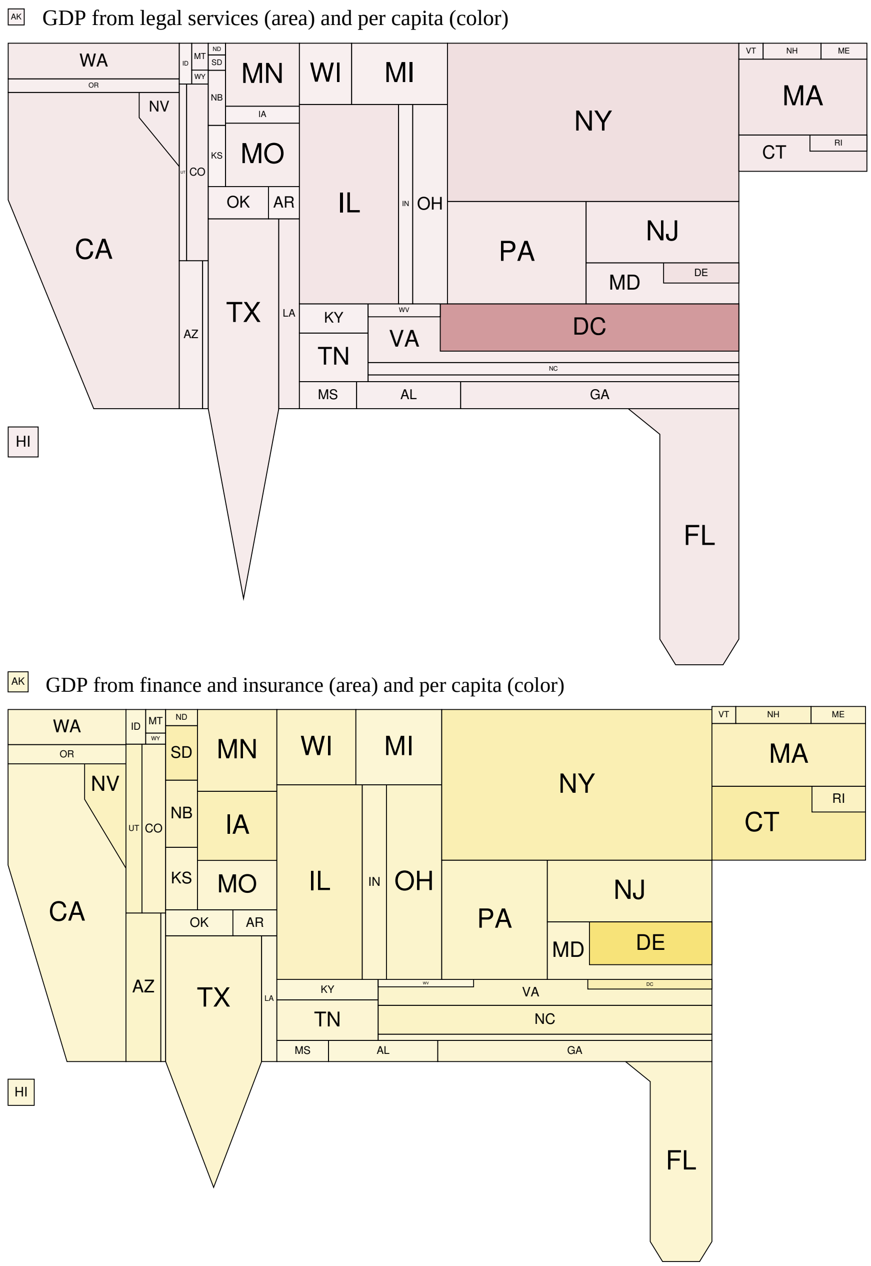

Figure 19: Top: GDP from legal services (area) and GDP per capita (shading). Bottom: Same, for the finance and insurances industries (data source: www.statemaster.com)

5.5 GDP from Agriculture and Extractive Industries

The economic sectors already shown are sized by their respective blue-collar (manufacturing) or white-collar (information industries, finance, etc.) workforces. Thus, they tend to distributed roughly as population, with some added weight, for blue-collar, to the Rust Belt; or, for white-collar, to the two coasts plus two or three Midwestern hubs.

A very different geographical picture emerges if we look at sectors of the economy that depend on land, or on natural resources. Figure 20, top, shows GDP from crops and animals, including farming, ranching, hunting and fishing. Intensity of shading is used to show the land-use intensity of the economic activity, as measured in GDP per acre. The farm belt of the Midwest dominates this map, along with the Pacific Coast and Florida. Unsurprisingly, arid state of the Intermountain West have a low intensity of land use (light shading). Perhaps more surprising is how small is the aggregate GDP that their ranching and farming actually produce.

The picture changes completely in Figure 20, bottom, where we show GDP from mining, oil and gas. Here again shading is used to show the land intensity of the economic activity, as measured in GDP per acre. Texas is now depicted at a scale that can satisfy most Texans. Texas, Louisiana, and Oklahoma are dominant in oil and gas production. Pennsylvania, West Virginia, and Kentucky are coal states. Wyoming, Colorado, New Mexico, and Arizona dominate in mining and other extractive industries. California has both oil production and mining.

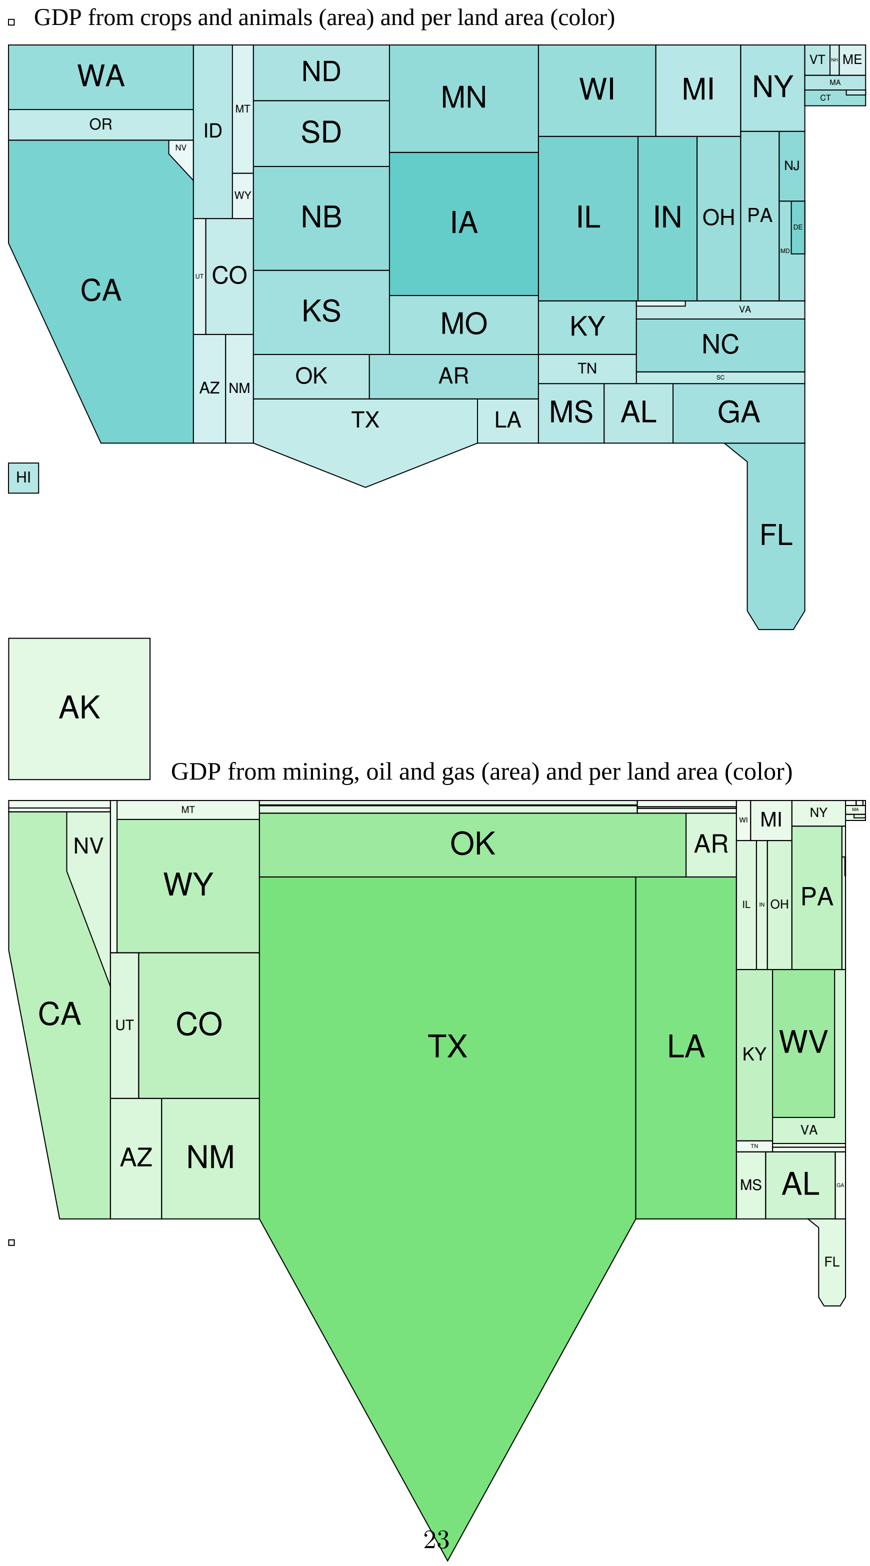

Figure 20: Top: GDP from crops and animals (area) and GDP per land area (shading). Bottom: Same, for mining, oil and gas industries (data source: www.statemaster.com)

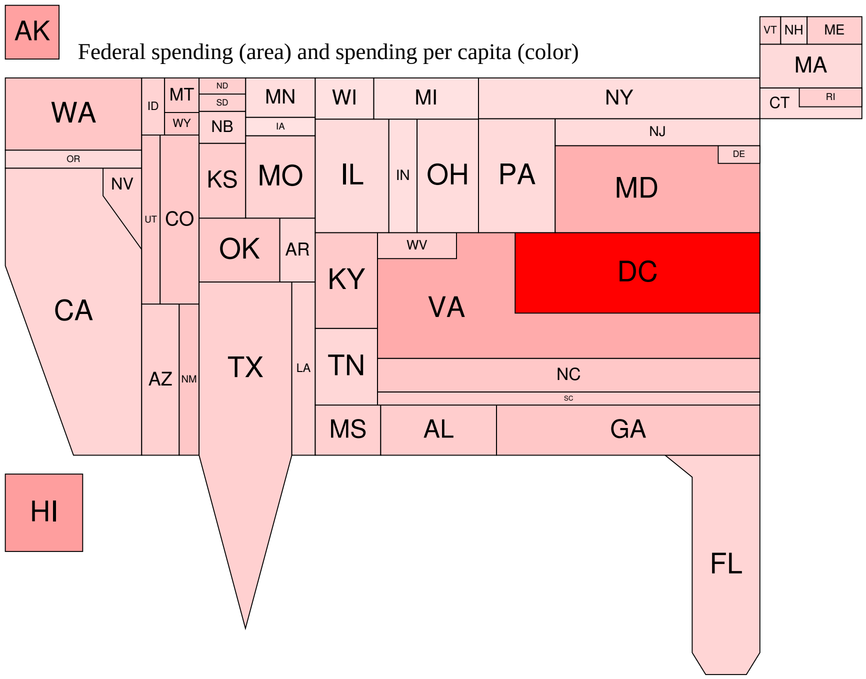

Figure 21: Federal spending by state is shown by area. Spending per resident is shown by intensity of shading. (data source: www.taxfoundation.org/research)

5.6 Federal Expenditures and Net Funds Flows

Where does the Federal government spend its money? A nominal answer is shown in Figure 21. The District of Columbia, along with neighboring Virginia and Maryland, show a surprising dominance. However, without digging deeper into the data, it is hard to know the extent to which Federal funds attributed to DC and surroundings are actually spent there, as opposed to representing Federal contracts with companies in other states that are booked at the location of the contracting Federal agency.

With this caveat about the data, an interesting way to view Federal taxes and expenditures is by their net flows across state lines. To a rough approximation most states get back from the Federal government about the same amount as their residents pay in Federal taxes. But this is only a rough equality, and many states are exceptions to the rule. In 2005, Federal spending was about \$2.5 Trillion, of which about \$300 Billion (net) crossed state lines. The United States can thus be viewed as two countries, a Donor States of America, where more taxes are collected than Federal funds are spent; and a Needy States of America, where more Federal funds are spent than taxes collected.

What do the maps of these two countries look like, scaling each state by dollars transferred? The answer is shown in Figure 22. As before, the dominance

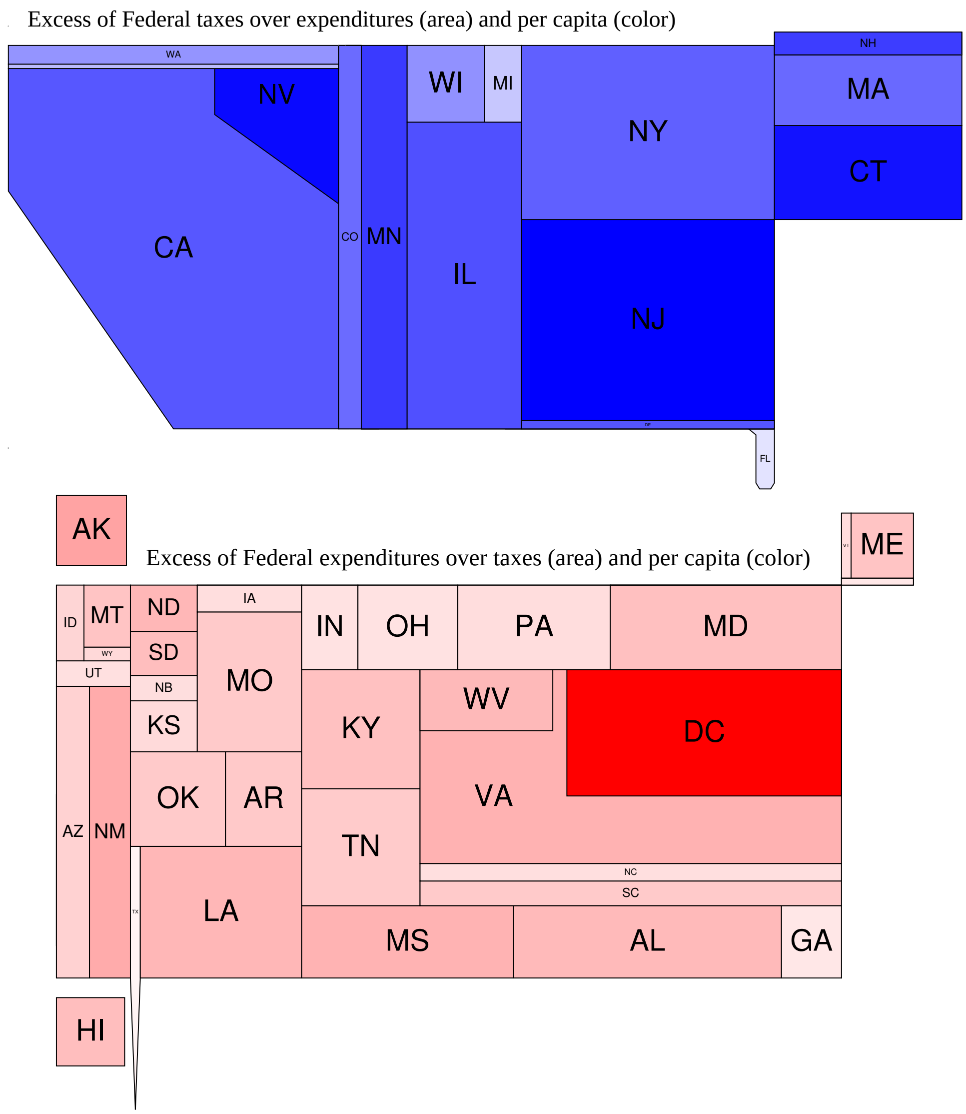

Figure 22: Top: the "Donor States of America" shown with size proportional to their net excess of Federal taxes over expenditures. Bottom: the "Needy States of America", with size proportional to their net excess of Federal expenditures over taxes. In both figures, shading intensity represents dollar amounts per capita. (data source: www.taxfoundation.org/research)

of DC, Maryland, and Virgina among the Needy States is likely exaggerated to an unknown extent by procurement bookkeeping practices; however, it seems likely that they are correctly shown as being “in the red”.

It is striking in Figure 22 that both the Donor and Needy States of America are (mostly) contiguous geographical areas. You could drive from Los Angeles to New York with only two quick hops through Needy states (through Utah, Nebraska, and Iowa, which are themselves only marginally in the red). Or, you could drive from Tucson to Atlanta never leaving the red at all. Another interesting fact is how closely in balance is Texas, a populous state. It is only barely in the red; Texas, in other words, is basically a country unto itself.

Also striking is to compare Figure 22 to Figure 12, which mapped U.S. politics on a red-blue scale. One can view the 2008 election results as being somewhat Democrat-skewed, and the 2010 results as being correspondingly Republican-skewed. So, a state’s long-term red-blueness is some compromise between the two maps. But that almost exactly describes the red-blue assignments in Figure 22. In other words, the Donor states are the long-term Democratic states, while the Needy states are the long-term Republican states. Put differently, it is the typical Democratic voter whose Federal taxes are subsidizing the typical Republican voter, not the other way around! That is certainly not what one expects from contemporary populist (e.g., Tea Party) rhetoric, and is a paradox that is known to political scientists. [10, 11, 12]

6 Conclusion

One should not be too quick to use the visual appeal of this paper’s cartographic treemaps (or of other cartograms elsewhere) in support of U.S. regional stereotypes. Nevertheless, regional distinctions do emerge clearly from the data shown here. We should celebrate the fact that now, 235 years after its founding and despite its historical fractures and present, sometimes challenging, diversity, the U.S. remains in fact a United States – and its outline remains highly recognizable even in quite stylized graphical representations of diverse data.

Acknowledgments

I thank Douglas Eardley for mathematical discussions, and Richard Rapp for pointing me to Lacy’s work on tax imbalances.

References

[1] Schulten, S. “Visualizing Slavery”, *New York Times*, December 9, 2010, at http://opinionator.blogs.nytimes.com/2010/12/09/visualizing-slavery/.

[2] Wikipedia, “Transformation Matrix,” at http://en.wikipedia.org/wiki/Transformation_matrix.

[3] Weisstein, E.W. "Conformal Projection," in MathWorld at http://mathworld.wolfram.com/ConformalProjection.html.

[4] Synder, J.P. *Flattening the earth: two thousand years of map projections* (University of Chicago Press, 1993) p. 65.

[5] Gastner, M.T. and Newman, M.E.J. "Diffusion-based method for producing density-equalizing maps," *PNAS*, **101**, 7499-7504 (2004).

[6] Wikipedia, "Treemapping," at http://en.wikipedia.org/wiki/Treemapping.

[7] Map of the Market, at http://www.smartmoney.com/map-of-the-market/. For its history, see http://www.bewitched.com/marketmap.html.

[8] Newman, M.E.J. "Maps of the 2008 US presidential election results," (2008) at http://www-personal.umich.edu/ mejn/election/2008/

[9] Gastner, M.T. "Cartograms of the 2010 US midterm election" (2010) at http://www2.imperial.ac.uk/ mgastner/cartogram/election10/election10.html.

[10] Lacy, D. "A Curious Paradox of the Red States and Blue States: Federal Spending and Electoral Votes in the 2000 Election," (2002) at http://psweb.sbs.ohio-state.edu/faculty/hweisberg/conference/Lacy-OSUConf.PDF

[11] Lacy, D. "Why Do Red States Vote Republican While Blue States Pay the Bills?" Paper presented at the Annual Meeting of the American Political Science Association (Toronto, September 3-6, 2009).

[12] Morris, D. "Do States Vote Against Their Economic Self-Interest?" blog at http://static.newrules.org/drdave/9-spending.html (seen 12/25/2010).