Int J Comput Math Learning (2009) 14:61–65
DOI 10.1007/s10758-009-9145-y

COMPUTATIONAL DIVERSIONS

Computational Diversions: Web Fame, Web Games

Michael Eisenberg

Published online: 24 March 2009

The cognitive scientist Gerd Gigerenzer—an expert in the subject of people's judgment and decision-making abilities—once did a fascinating experiment demonstrating the "positive value of ignorance." A full description can be found in (Gigerenzer et al. 1999); here, I'll just summarize the basic idea.

Gigerenzer began with two classrooms of students—one from Germany, one from the United States. He then gave each group two tests relating to cities from the two nations. (That is to say, American students took a test on German cities and a test on American cities; and the German students likewise took the same two tests.) An interesting result occurred: the German students did better on the *American city* test than did the students from the US. Now, before you jump to the conclusion (as I did) that this is just one more proof of the dismal state of American education, it should also be noted that the American students did better on the *German city* test than did the German students. In other words, the students on average outperformed their counterparts on the *other* group's geography.

How could such an unexpected result occur? After all, don't German (or American) kids know more about their own country than kids living on a different continent? The reason for Gigerenzer's result lies in the format of his test. All of the questions had the same basic structure, asking students to judge between two cities as to which had greater population. Here are a couple of sample questions, one from each test:

Which is bigger: DORTMUND or MUNICH?
Which is bigger: SAN JOSE or SAN ANTONIO?

Now, suppose (for example) that you are a German student looking at that second question, and further assume that (realistically) you've never paid much attention to the subject of city populations. You can now be in one of three situations: (a) you've never heard of either San Jose or San Antonio, (b) you've heard of only one of those two towns, or (c) you've heard of both of them. If you're in situation (a), you're out of luck—you may as well flip a coin in answering the question. If you're in situation (c), you might bring to bear whatever meager information you have about the two cities in

M. Eisenberg (✉)
University of Colorado, Boulder, CO, USA
e-mail: ijcml-diversions@ccl.northwestern.edu

answering the question. But if you're in situation (b), you might reason that the town that you've heard of is larger, simply because… well, simply because. It makes sense, after all, that if you've *heard of* a place, it's probably larger than a place you've never heard of. The correlation between "fame" and "population," in other words, is likely to be positive; so if you answer the question with the city that you've heard of, you're pretty likely to get the right answer. (In point of fact, by the way, the correct answer is "San Antonio"—so if you're a German student who has heard of San Antonio but not San Jose, you're in luck on this question.)

Now we're in a position to understand Gigerenzer's result. Suppose (just to take plausible values) that you're a student answering questions about city populations—whether those cities are in your own country or somewhere else. You will get the correct answer 50% of the time in situation (a), when you've never heard of either town; after all, you're just flipping a coin in total ignorance. In situation (c), let's suppose you get the answer right 60% of the time; in this case, you can use whatever you know about the cities to guess at an answer, and you'll do somewhat better than chance based on your knowledge. In situation (b), though, you're using your environment to answer the question for you—the environment that has *familiarized you* with one city but not another. In this case, let's say that 75% of the time the environment has suggested the right answer; or, to go back to our previous example, of those German students who have heard of only one of the two towns in question, three quarters will have heard of San Antonio and one quarter will have heard of San Jose.

With these (admittedly hypothetical but not implausible) parameters, how would a student from country X do on a test of her own cities? The answer is about 60%—she's *heard of* every city, but in many cases has little else to go on in answering the question of which is larger. How will she do on the test of country Y? Well, let's say that she knows half the cities on the test, and that every pair of N distinct cities is tested. In that case, she'll answer about one quarter of the questions in situation (a), when she's just guessing; one quarter in situation (c), where she can use her knowledge about both cities; and about one half in situation (b), when she's heard of one city but not the other. (To be exact, if there are 100 cities, she'll see 50 * 49/2 questions in which both cities are known, 50 * 49/2 questions in which both cities are unknown, and 50 * 50 questions in which only one city is known.) Her score in this case will be about 0.25 *(0.5 + 0.6) + 0.5 *(0.75), or 0.65. In other words, the student will do about 5% better on the unfamiliar country's test than on her own.

Is this a weird, unrealistic scenario? Gigerenzer would argue that in fact we're in situations of this sort fairly often—situations where we make a judgment (between brands, between politicians, between colleges) based on what we hope is a reasonably good correlation between familiarity and quality. I was once placed in a very obvious situation of this type while at a conference in Europe: the World Cup was going on, and the audience was asked to choose "which of these four players would you want to take a penalty shot for your team?" Now, I know absolutely nothing about soccer (strictly a baseball fan), and three of the four names were totally unfamiliar to me. One, David Beckham, was a name that I recognized as "a world-famous soccer player." So I chose him as my hypothetical penalty-kicker, on the assumption that if I'd heard of Beckham he has to be on that basis alone a pretty prominent—and therefore, I'd assume, pretty remarkable—player. (People like Babe Ruth and Michael Jordan don't achieve world notoriety by being merely good at their respective sports.)

Gigerenzer's discussion of his experiment is eminently plausible, but it is missing one rather important corroborative element—namely, an independent metric for "fame." Okay, we might say: San Antonio is larger in population than San Jose. Does that mean that San Antonio is more *famous* than San Jose? How would we go about measuring that?

Conceivably, we could (say) do an exhaustive survey of a selection of newspapers and television programs for a month and see which of the two cities gets mentioned more; but in the age of the search engine, we have a much easier metric.

Suppose we use the number of occurrences of a city's name on the Web as a measure of its fame. In this case, we'd expect San Antonio to be mentioned more often on the Web than San Jose. (In fact, we might even use the gap between the number of mentions of the two cities as a metric of how much more San Antonio is famous than San Jose, but for our present purposes we needn't go that far.) With this in mind, we can try an experiment in which the Web, effectively, "takes Gigerenzer's test" by using fame as a stand-in for population. That is, we'll have the Web respond to each city-comparison question by answering that the larger city is the one with a higher number of mentions in our favorite search engine. Notice that we could, of course, use the Web to find a city's estimated population directly; but instead, we're just having the Web guess at the population in the same way that a partially ignorant student might.

Using Google as my search engine, I came up with the results in Table 1. I actually ran the experiment a couple of ways. Column 1 shows, in order from top to bottom, the 20 most populous cities in the United States (as returned by the website infoplease.com (City Population website: http://www.infoplease.com/ipa/A0763098.html)); and Column 2 shows each city's estimated population in thousands as of 2007. Column 3 shows the result of doing a Google search of each city paired with its state name—for instance, we find that there are 11.9 million references on the Web to "Phoenix, Arizona" and 8.8 million references to "Philadelphia, Pennsylvania"; Column 4 shows the city ranking that we get

Table 1 American cities, ranked by population and by "Web Fame" (as described in the text)

City	Population (in thousands)	References (with state name; in thousands)	Rank (using Column 3)	References (name alone; in millions)	Rank (using Column 5)
New York	8,275	26,500	1	1,070	1
Los Angeles	3,834	20,500	4	312	3
Chicago	2,837	22,000	3	359	2
Houston	2,208	22,900	2	180	7
Phoenix	1,552	11,900	8	191	6
Philadelphia	1,450	8,810	13	173	8
San Antonio	1,329	9,410	12	252	5
San Diego	1,267	11,000	10	170	9
Dallas	1,241	16,300	6	166	10
San Jose	940	5,780	16	107	16
Detroit	917	5,720	17	139	13
Jacksonville	806	6,300	15	60	20
Indianapolis	795	6,790	14	66	19
San Francisco	765	11,900	9	267	4
Columbus	748	15,400	7	99	17
Austin	743	18,200	5	166	11
Fort Worth	682	4,810	19	117	14
Memphis	674	4,310	20	69	18
Charlotte	671	4,850	18	142	12
Baltimore	637	10,400	11	108	15

from these references. Column 5 repeats the experiment, but this time without requiring the state names; obviously, since we're now looking at a superset of the Column 3 references, we'll obtain larger numbers. The word "phoenix" now gives us 191 million references, and "philadelphia" 173 million. The rankings obtained this way are shown in Column 6.

In point of fact, it's hard to argue that any one way of measuring fame in this fashion is inarguably the correct way. To take the difference between my two experiments: should we look for "Phoenix, Arizona" or just "phoenix?" Maybe some references to "phoenix" (I didn't bother with capitals in the second experiment) are alluding to the mythical bird? So perhaps we should always employ the more formal version of the city's name…. But then again, don't most typical city references dispense with the state name? (Who says "I'm going to Chicago, Illinois" as opposed to "I'm going to Chicago?") In any event, the two versions of the city experiment come up with similar, though not identical results. Some cities—San Francisco, for instance—appear to be disproportionately famous in relation to their size; and San Francisco's geographic neighbor, San Jose, appears to be a little bit more obscure than its population would suggest. But by and large, the Web does pretty well on Gigerenzer's test. On a full test of the 20 largest American cities, the Web gets 72% of its answers correct using Columns 3 and 4, and 78% of its answers correct using Columns 5 and 6.

These results suggest that the "environment," as represented by the Web, is in fact doing a pretty good job of indicating quantities such as population size through the stand-in concept of fame. Gigerenzer's argument is lent some support by this observation: in the absence of more specific information, we can often make reasonable judgments about measures like size (and perhaps more controversially, other measures like "quality") simply based on name recognition.

With this in mind, we can try some other informal experiments—or party games—using the Web as an arbiter for fame. For instance, suppose we try the following questions:

Order from most to least famous:

Silent Movie Stars

(a) Clara Bow
(b) Charlie Chaplin
(c) Harold Lloyd
(d) Mary Pickford
(e) Rudolph Valentino

Greek Deities

(a) Aphrodite
(b) Apollo
(c) Athena
(d) Poseidon
(e) Zeus

Mathematicians

(a) Leonhard Euler
(b) Evariste Galois
(c) David Hilbert
(d) Alan Turing
(e) John von Neumann

Alternatively, you might imagine a different sort of Web game based on search engine counts. Suppose we type in a plausible but rather idiosyncratic series of phrases—for example, which of these would you expect to appear most often on the Web?:

(a) "Hey, put her down"
(b) "Hey, put him down"
(c) "Hey, put me down"
(d) "Hey, put that down"

Or one might try the following as a challenge at a party: given a specific phrase, ask everyone to guess the number of times that phrase occurs on the Web. The person who guesses the closest value—or maybe the closest underestimate—wins. Or—just to take one more challenge—ask everyone to suggest a phrase that will appear more than 10 times on the Web, but fewer than 100 times. (This is a lot harder than it sounds at first. It took me about ten tries before I hit on "evil rat terrier," with 28 occurrences on Google.)

If you have a suggestion for a Web pastime in a similar spirit—or if you have interesting examples of experiments similar to those described here—please email them to the column at ijcml-diversions@ccl.northwestern.edu.

Reference

Gigerenzer, G., et al. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.