1

What Do Schoolteachers and Sumo Wrestlers Have in Common?

Imagine for a moment that you are the manager of a day-care center. You have a clearly stated policy that children are supposed to be picked up by 4 p.m. But very often parents are late. The result: at day's end, you have some anxious children and at least one teacher who must wait around for the parents to arrive. What to do?

A pair of economists who heard of this dilemma—it turned out to be a rather common one—offered a solution: fine the tardy parents. Why, after all, should the day-care center take care of these kids for free?

The economists decided to test their solution by conducting a study of ten day-care centers in Haifa, Israel. The study lasted twenty weeks, but the fine was not introduced immediately. For the first four weeks, the economists simply kept track of the number of parents who came late; there were, on average, eight late pickups per week per day-care center. In the fifth week, the fine was enacted. It was announced that any parent arriving more than ten minutes late would pay $3 per child for each incident. The fee would be added to the parents' monthly bill, which was roughly $380.

After the fine was enacted, the number of late pickups promptly went…up. Before long there were twenty late pickups per week, more than double the original average. The incentive had plainly backfired.

Economics is, at root, the study of incentives: how people get what they want, or need, especially when other people want or need the same thing. Economists love incentives. They love to dream them up and enact them, study them and tinker with them. The typical economist believes the world has not yet invented a problem that he cannot fix if given a free hand to design the proper incentive scheme. His solution may not always be pretty—it may involve coercion or exorbitant penalties or the violation of civil liberties—but the original problem, rest assured, will be fixed. An incentive is a bullet, a lever, a

key: an often tiny object with astonishing power to change a situation.

We all learn to respond to incentives, negative and positive, from the outset of life. If you toddle over to the hot stove and touch it, you burn a finger. But if you bring home straight A's from school, you get a new bike. If you are spotted picking your nose in class, you get ridiculed. But if you make the basketball team, you move up the social ladder. If you break curfew, you get grounded. But if you ace your SATs, you get to go to a good college. If you flunk out of law school, you have to go to work at your father's insurance company. But if you perform so well that a rival company comes calling, you become a vice president and no longer have to work for your father. If you become so excited about your new vice president job that you drive home at eighty mph, you get pulled over by the police and fined $100. But if you hit your sales projections and collect a year-end bonus, you not only aren't worried about the $100 ticket but can also afford to buy that Viking range you've always wanted—and on which your toddler can now burn her own finger.

An incentive is simply a means of urging people to do more of a good thing and less of a bad thing. But most incentives don't come about organically. Someone—an economist or a politician or a parent—has to invent them. Your three-year-old eats all her vegetables for a week? She wins a trip to the toy store. A big steelmaker belches too much smoke into the air? The company is fined for each cubic foot of pollutants over the legal limit. Too many Americans aren't paying their share of income tax? It was the economist Milton Friedman who helped come up with a solution to this one: automatic tax withholding from employees' paychecks.

There are three basic flavors of incentive: economic, social, and moral. Very often a single incentive scheme will include all three varieties. Think about the anti-smoking campaign of recent years. The addition of a $3-per-pack "sin tax" is a strong economic incentive against buying cigarettes. The banning of cigarettes in restaurants and bars is a powerful social incentive. And when the U.S. government asserts that terrorists raise money by selling black-market cigarettes, that acts as a rather jarring moral incentive.

Some of the most compelling incentives yet invented have been put in place to deter crime. Considering this fact, it might be worthwhile to take a familiar question—why is there so much crime in modern society?—and stand it on its head: why isn't there a lot *more* crime?

After all, every one of us regularly passes up opportunities to maim, steal, and defraud. The chance of going to jail—thereby losing your job, your house, and your freedom, all of which are essentially economic penalties—is certainly a strong incentive. But when it comes to crime, people also respond to moral

incentives (they don't want to do something they consider wrong) and social incentives (they don't want to be seen by others as doing something wrong). For certain types of misbehavior, social incentives are terribly powerful. In an echo of Hester Prynne's scarlet letter, many American cities now fight prostitution with a "shaming" offensive, posting pictures of convicted johns (and prostitutes) on websites or on local-access television. Which is a more horrifying deterrent: a $500 fine for soliciting a prostitute or the thought of your friends and family ogling you on www.HookersAndJohns.com?

So through a complicated, haphazard, and constantly readjusted web of economic, social, and moral incentives, modern society does its best to militate against crime. Some people would argue that we don't do a very good job. But taking the long view, that is clearly not true. Consider the historical trend in homicide (not including wars), which is both the most reliably measured crime and the best barometer of a society's overall crime rate. These statistics, compiled by the criminologist Manuel Eisner, track the historical homicide levels in five European regions.

HOMICIDES
(per 100,000 People)

	England	Netherlands and Belgium	Scandinavia	Germany and Switzerland	Italy
13th and 14th c.	23.0	47.0	n.a.	37.0	56.0
15th c.	n.a.	45.0	46.0	16.0	73.0
16th c.	7.0	25.0	21.0	11.0	47.0
17th c.	5.0	7.5	18.0	7.0	32.0
18th c.	1.5	5.5	1.9	7.5	10.5
19th c.	1.7	1.6	1.1	2.8	12.6
1900–1949	0.8	1.5	0.7	1.7	3.2
1950–1994	0.9	0.9	0.9	1.0	1.5

The steep decline of these numbers over the centuries suggests that, for one of the gravest human concerns—getting murdered—the incentives that we collectively cook up are working better and better.

So what was wrong with the incentive at the Israeli day-care centers?

You have probably already guessed that the $3 fine was simply too small. For that price, a parent with one child could afford to be late every day and only

pay an extra $60 each month—just one-sixth of the base fee. As babysitting goes, that's pretty cheap. What if the fine had been set at $100 instead of $3? That would have likely put an end to the late pickups, though it would have also engendered plenty of ill will. (Any incentive is inherently a trade-off; the trick is to balance the extremes.)

But there was another problem with the day-care center fine. It substituted an economic incentive (the $3 penalty) for a moral incentive (the guilt that parents were supposed to feel when they came late). For just a few dollars each day, parents could buy off their guilt. Furthermore, the small size of the fine sent a signal to the parents that late pickups weren't such a big problem. If the day-care center suffers only $3 worth of pain for each late pickup, why bother to cut short your tennis game? Indeed, when the economists eliminated the $3 fine in the seventeenth week of their study, the number of late-arriving parents didn't change. Now they could arrive late, pay no fine, *and* feel no guilt.

Such is the strange and powerful nature of incentives. A slight tweak can produce drastic and often unforeseen results. Thomas Jefferson noted this while reflecting on the tiny incentive that led to the Boston Tea Party and, in turn, the American Revolution: "So inscrutable is the arrangement of causes and consequences in this world that a two-penny duty on tea, unjustly imposed in a sequestered part of it, changes the condition of all its inhabitants."

In the 1970s, researchers conducted a study that, like the Israeli day-care study, pitted a moral incentive against an economic incentive. In this case, they wanted to learn about the motivation behind blood donations. Their discovery: when people are given a small stipend for donating blood rather than simply being praised for their altruism, they tend to donate *less* blood. The stipend turned a noble act of charity into a painful way to make a few dollars, and it wasn't worth it.

What if the blood donors had been offered an incentive of $50, or $500, or $5,000? Surely the number of donors would have changed dramatically.

But something else would have changed dramatically as well, for every incentive has its dark side. If a pint of blood were suddenly worth $5,000, you can be sure that plenty of people would take note. They might literally steal blood at knifepoint. They might pass off pig blood as their own. They might circumvent donation limits by using fake IDs. Whatever the incentive, whatever the situation, dishonest people will try to gain an advantage by whatever means necessary.

Or, as W. C. Fields once said: a thing worth having is a thing worth cheating for.

Who cheats?

Well, just about anyone, if the stakes are right. You might say to yourself, *I* don't cheat, regardless of the stakes. And then you might remember the time you cheated on, say, a board game. Last week. Or the golf ball you nudged out of its bad lie. Or the time you really wanted a bagel in the office break room but couldn't come up with the dollar you were supposed to drop in the coffee can. And then took the bagel anyway. And told yourself you'd pay double the next time. And didn't.

For every clever person who goes to the trouble of creating an incentive scheme, there is an army of people, clever and otherwise, who will inevitably spend even more time trying to beat it. Cheating may or may not be human nature, but it is certainly a prominent feature in just about every human endeavor. Cheating is a primordial economic act: getting more for less. So it isn't just the boldface names—inside-trading CEOs and pill-popping ballplayers and perk-abusing politicians—who cheat. It is the waitress who pockets her tips instead of pooling them. It is the Wal-Mart payroll manager who goes into the computer and shaves his employees' hours to make his own performance look better. It is the third grader who, worried about not making it to the fourth grade, copies test answers from the kid sitting next to him.

Some cheating leaves barely a shadow of evidence. In other cases, the evidence is massive. Consider what happened one spring evening at midnight in 1987: seven million American children suddenly disappeared. The worst kidnapping wave in history? Hardly. It was the night of April 15, and the Internal Revenue Service had just changed a rule. Instead of merely listing the name of each dependent child, tax filers were now required to provide a Social Security number. Suddenly, seven million children—children who had existed only as phantom exemptions on the previous year's 1040 forms—vanished, representing about one in ten of all dependent children in the United States.

The incentive for those cheating taxpayers was quite clear. The same for the waitress, the payroll manager, and the third grader. But what about that third grader's *teacher*? Might she have an incentive to cheat? And if so, how would she do it?

Imagine now that instead of running a day-care center in Haifa, you are running the Chicago Public Schools, a system that educates 400,000 students each year.

The most volatile current debate among American school administrators, teachers, parents, and students concerns "high-stakes" testing. The stakes are considered high because instead of simply testing students to measure their

progress, schools are increasingly held accountable for the results.

The federal government mandated high-stakes testing as part of the No Child Left Behind law, signed by President Bush in 2002. But even before that law, most states gave annual standardized tests to students in elementary and secondary school. Twenty states rewarded individual schools for good test scores or dramatic improvement; thirty-two states sanctioned the schools that didn't do well.

The Chicago Public School system embraced high-stakes testing in 1996. Under the new policy, a school with low reading scores would be placed on probation and face the threat of being shut down, its staff to be dismissed or reassigned. The CPS also did away with what is known as social promotion. In the past, only a dramatically inept or difficult student was held back a grade. Now, in order to be promoted, every student in third, sixth, and eighth grade had to manage a minimum score on the standardized, multiple-choice exam known as the Iowa Test of Basic Skills.

Advocates of high-stakes testing argue that it raises the standards of learning and gives students more incentive to study. Also, if the test prevents poor students from advancing without merit, they won't clog up the higher grades and slow down good students. Opponents, meanwhile, worry that certain students will be unfairly penalized if they don't happen to test well, and that teachers may concentrate on the test topics at the exclusion of more important lessons.

Schoolchildren, of course, have had incentive to cheat for as long as there have been tests. But high-stakes testing has so radically changed the incentives for teachers that they too now have added reason to cheat. With high-stakes testing, a teacher whose students test poorly can be censured or passed over for a raise or promotion. If the entire school does poorly, federal funding can be withheld; if the school is put on probation, the teacher stands to be fired. High-stakes testing also presents teachers with some positive incentives. If her students do well enough, she might find herself praised, promoted, and even richer: the state of California at one point introduced bonuses of $25,000 for teachers who produced big test-score gains.

And if a teacher were to survey this newly incentivized landscape and consider somehow inflating her students' scores, she just might be persuaded by one final incentive: teacher cheating is rarely looked for, hardly ever detected, and just about never punished.

How might a teacher go about cheating? There are any number of possibilities, from brazen to subtle. A fifth-grade student in Oakland recently came home from school and gaily told her mother that her super-nice teacher

had written the answers to the state exam right there on the chalkboard. Such instances are certainly rare, for placing your fate in the hands of thirty prepubescent witnesses doesn't seem like a risk that even the worst teacher would take. (The Oakland teacher was duly fired.) There are more nuanced ways to inflate students' scores. A teacher can simply give students extra time to complete the test. If she obtains a copy of the exam early—that is, illegitimately—she can prepare them for specific questions. More broadly, she can "teach to the test," basing her lesson plans on questions from past years' exams, which isn't considered cheating but may well violate the spirit of the test. Since these tests all have multiple-choice answers, with no penalty for wrong guesses, a teacher might instruct her students to randomly fill in every blank as the clock is winding down, perhaps inserting a long string of Bs or an alternating pattern of Bs and Cs. She might even fill in the blanks for them after they've left the room.

But if a teacher *really* wanted to cheat—and make it worth her while—she might collect her students' answer sheets and, in the hour or so before turning them in to be read by an electronic scanner, erase the wrong answers and fill in correct ones. (And you always thought that no. 2 pencil was for the *children* to change their answers.) If this kind of teacher cheating is truly going on, how might it be detected?

To catch a cheater, it helps to think like one. If you were willing to erase your students' wrong answers and fill in correct ones, you probably wouldn't want to change too many wrong answers. That would clearly be a tip-off. You probably wouldn't even want to change answers on every student's test—another tip-off. Nor, in all likelihood, would you have enough time, because the answer sheets have to be turned in soon after the test is over. So what you might do is select a string of eight or ten consecutive questions and fill in the correct answers for, say, one-half or two-thirds of your students. You could easily memorize a short pattern of correct answers, and it would be a lot faster to erase and change that pattern than to go through each student's answer sheet individually. You might even think to focus your activity toward the end of the test, where the questions tend to be harder than the earlier questions. In that way, you'd be most likely to substitute correct answers for wrong ones.

If economics is a science primarily concerned with incentives, it is also—fortunately—a science with statistical tools to measure how people respond to those incentives. All you need are some data.

In this case, the Chicago Public School system obliged. It made available a database of the test answers for every CPS student from third grade through seventh grade from 1993 to 2000. This amounts to roughly 30,000 students per grade per year, more than 700,000 sets of test answers, and nearly 100 million

individual answers. The data, organized by classroom, included each student's question-by-question answer strings for reading and math tests. (The actual paper answer sheets were not included; they were habitually shredded soon after a test.) The data also included some information about each teacher and demographic information for every student, as well as his or her past and future test scores—which would prove a key element in detecting the teacher cheating.

Now it was time to construct an algorithm that could tease some conclusions from this mass of data. What might a cheating teacher's classroom look like?

The first thing to search for would be unusual answer patterns in a given classroom: blocks of identical answers, for instance, especially among the harder questions. If ten very bright students (as indicated by past and future test scores) gave correct answers to the exam's first five questions (typically the easiest ones), such an identical block shouldn't be considered suspicious. But if ten poor students gave correct answers to the *last* five questions on the exam (the hardest ones), that's worth looking into. Another red flag would be a strange pattern within any one student's exam—such as getting the hard questions right while missing the easy ones—especially when measured against the thousands of students in other classrooms who scored similarly on the same test. Furthermore, the algorithm would seek out a classroom full of students who performed far better than their past scores would have predicted and who then went on to score significantly lower the following year. A dramatic one-year spike in test scores might initially be attributed to a *good* teacher; but with a dramatic fall to follow, there's a strong likelihood that the spike was brought about by artificial means.

Consider now the answer strings from the students in two sixth-grade Chicago classrooms who took the identical math test. Each horizontal row represents one student's answers. The letter a, b, c, or d indicates a correct answer; a number indicates a wrong answer, with 1 corresponding to a, 2 corresponding to b, and so on. A zero represents an answer that was left blank. One of these classrooms almost certainly had a cheating teacher and the other did not. Try to tell the difference—although be forewarned that it's not easy with the naked eye.

Classroom A

112a4a342cb214d0001acd24a3a12dadbcb4a0000000
d4a2341cacbddad3142a2344a2ac23421c00adb4b3cb

1b2a34d4ac42d23b141acd24a3a12dadbcb4a2134141
dbaab3dcacb1dadbc42ac2cc31012dadbcb4adb40000
d12443d43232d32323c213c22d2c23234c332db4b300
db2abad1acbdda212b1acd24a3a12dadbcb400000000
d4aab2124cbddadbcb1a42cca3412dadbcb423134bc1
1b33b4d4a2b1dadbc3ca22c000000000000000000000
d43a3a24acb1d32b412acd24a3a12dadbcb422143bc0
313a3ad1ac3d2a23431223c000012dadbcb400000000
db2a33dcacbd32d313c21142323cc300000000000000
d43ab4d1ac3dd43421240d24a3a12dadbcb400000000
db223a24acb11a3b24cacd12a241cdadbcb4adb4b300
db4abadcacb1dad3141ac212a3a1c3a144ba2db41b43
1142340c2cbddadb4b1acd24a3a12dadbcb43d133bc4
214ab4dc4cbdd31b1b2213c4ad412dadbcb4adb00000
1423b4d4a23d24131413234123a243a2413a21441343
3b3ab4d14c3d2ad4cbcac1c003a12dadbcb4adb40000
dba2ba21ac3d2ad3c4c4cd40a3a12dadbcb400000000
d122ba2cacbd1a13211a2d02a2412d0dbcb4adb4b3c0
144a3adc4cbddadbcbc2c2cc43a12dadbcb4211ab343
d43aba3cacbddadbcbca42c2a3212dadbcb42344b3cb

Classroom B

db3a431422bd131b4413cd422a1acda332342d3ab4c4
d1aa1a11acb2d3dbc1ca22c23242c3a142b3adb243c1
d42a12d2a4b1d32b21ca2312a3411d00000000000000
3b2a34344c32d21b1123cdc0000000000000000000000
34aabad12cbdd3d4c1ca112cad2ccd00000000000000
d33a3431a2b2d2d44b2acd2cad2c2223b40000000000
23aa32d2a1bd2431141342c13d212d233c34a3b3b000
d32234d4a1bdd23b242a22c2a1a1cda2b1baa33a0000
d3aab23c4cbddadb23c322c2a222223232b443b24bc3
d13a14313c31d42b14c421c42332cd2242b3433a3343
d13a3ad122b1da2b11242dc1a3a12100000000000000
d12a3ad1a13d23d3cb2a21ccada24d2131b440000000
314a133c4cbd142141ca424cad34c122413223ba4b40
d42a3adcacbddadbc42ac2c2ada2cda341baa3b24321

db1134dc2cb2dadb24c412c1ada2c3a341ba20000000
d1341431acbddad3c4c213412da22d3d1132a1344b1b
1ba41a21a1b2dadb24ca22c1ada2cd32413200000000
dbaa33d2a2bddadbcbca11c2a2accda1b2ba20000000

If you guessed that classroom A was the cheating classroom, congratulations. Here again are the answer strings from classroom A, now reordered by a computer that has been asked to apply the cheating algorithm and seek out suspicious patterns.

Classroom A
(With cheating algorithm applied)

1. 112a4a342cb214d0001**acd24a3a12dadbcb4**a0000000
2. 1b2a34d4ac42d23b141**acd24a3a12dadbcb4**a2134141
3. db2abad1acbdda212b1**acd24a3a12dadbcb4**00000000
4. d43a3a24acb1d32b412**acd24a3a12dadbcb4**22143bc0
5. 1142340c2cbddadb4b1**acd24a3a12dadbcb4**3d133bc4
6. d43ab4d1ac3dd43421240d24**a3a12dadbcb4**00000000
7. dba2ba21ac3d2ad3c4c4cd40**a3a12dadbcb4**00000000
8. 144a3adc4cbddadbcbc2c2cc4**3a12dadbcb4**211ab343
9. 3b3ab4d14c3d2ad4cbcac1c00**3a12dadbcb4**adb40000
10. d43aba3cacbddadbcbca42c2a32**12dadbcb4**2344b3cb
11. 214ab4dc4cbdd31b1b2213c4ad4**12dadbcb4**adb00000
12. 313a3ad1ac3d2a23431223c0000**12dadbcb4**00000000
13. d4aab2124cbddadbcb1a42cca34**12dadbcb4**23134bc1
14. dbaab3dcacb1dadbc42ac2cc310**12dadbcb4**adb40000
15. db223a24acb11a3b24cacd12a241c**dadbcb4**adb4b300
16. d122ba2cacbd1a13211a2d02a2412d0dbcb4adb4b3c0
17. 1423b4d4a23d24131413234123a243a2413a21441343
18. db4abadcacb1dad3141ac212a3a1c3a144ba2db41b43
19. db2a33dcacbd32d313c21142323cc300000000000000
20. 1b33b4d4a2b1dadbc3ca22c0000000000000000000000
21. d12443d43232d32323c213c22d2c23234c332db4b300
22. d4a2341cacbddad3142a2344a2ac23421c00adb4b3cb

Take a look at the answers in bold. Did fifteen out of twenty-two students somehow manage to reel off the same six consecutive correct answers (the d-a-d-b-c-b string) all by themselves?

There are at least four reasons this is unlikely. One: those questions, coming near the end of the test, were harder than the earlier questions. Two: these were mainly subpar students to begin with, few of whom got six consecutive right answers elsewhere on the test, making it all the more unlikely they would get right the same six hard questions. Three: up to this point in the test, the fifteen students' answers were virtually uncorrelated. Four: three of the students (numbers 1, 9, and 12) left more than one answer blank *before* the suspicious string and then ended the test with another string of blanks. This suggests that a long, unbroken string of blank answers was broken not by the student but by the teacher.

There is another oddity about the suspicious answer string. On nine of the fifteen tests, the six correct answers are preceded by another identical string, 3-a-1-2, which includes three of four *incorrect* answers. And on all fifteen tests, the six correct answers are followed by the same incorrect answer, a 4. Why on earth would a cheating teacher go to the trouble of erasing a student's test sheet and then fill in the *wrong* answer?

Perhaps she is merely being strategic. In case she is caught and hauled into the principal's office, she could point to the wrong answers as proof that she didn't cheat. Or perhaps—and this is a less charitable but just as likely answer—she doesn't know the right answers herself. (With standardized tests, the teacher is typically not given an answer key.) If this is the case, then we have a pretty good clue as to why her students are in need of inflated grades in the first place: they have a bad teacher.

Another indication of teacher cheating in classroom A is the class's overall performance. As sixth graders who were taking the test in the eighth month of the academic year, these students needed to achieve an average score of 6.8 to be considered up to national standards. (Fifth graders taking the test in the eighth month of the year needed to score 5.8, seventh graders 7.8, and so on.) The students in classroom A averaged 5.8 on their sixth-grade tests, which is a full grade level below where they should be. So plainly these are poor students. A year earlier, however, these students did even worse, averaging just 4.1 on their fifth-grade tests. Instead of improving by one full point between fifth and sixth grade, as would be expected, they improved by 1.7 points, nearly two grades' worth. But this miraculous improvement was short-lived. When these sixth-grade students reached seventh grade, they averaged 5.5—more than two grade levels below standard and even *worse* than they did in sixth grade. Consider the

erratic year-to-year scores of three particular students from classroom A:

	5TH GRADE SCORE	6TH GRADE SCORE	7TH GRADE SCORE
Student 3	3.0	6.5	5.1
Student 5	3.6	6.3	4.9
Student 14	3.8	7.1	5.6

The three-year scores from classroom B, meanwhile, are also poor but at least indicate an honest effort: 4.2, 5.1, and 6.0. So an entire roomful of children in classroom A suddenly got very smart one year and very dim the next, or more likely, their sixth-grade teacher worked some magic with her pencil.

There are two noteworthy points to be made about the children in classroom A, tangential to the cheating itself. The first is that they are obviously in poor academic shape, which makes them the very children whom high-stakes testing is promoted as helping the most. The second point is that these students (and their parents) would be in for a terrible shock once they reached the seventh grade. All they knew was that they had been successfully promoted due to their test scores. (No child left behind, indeed.) *They* weren't the ones who artificially jacked up their scores; they probably expected to do great in the seventh grade—and then they failed miserably. This may be the cruelest twist yet in high-stakes testing. A cheating teacher may tell herself that she is helping her students, but the fact is that she would appear far more concerned with helping herself.

An analysis of the entire Chicago data reveals evidence of teacher cheating in more than two hundred classrooms per year, roughly 5 percent of the total. This is a conservative estimate, since the algorithm was able to identify only the most egregious form of cheating—in which teachers systematically changed students' answers—and not the many subtler ways a teacher might cheat. In a recent study among North Carolina schoolteachers, some 35 percent of the respondents said they had witnessed their colleagues cheating in some fashion, whether by giving students extra time, suggesting answers, or manually changing students' answers.

What are the characteristics of a cheating teacher? The Chicago data shows that male and female teachers are equally prone to cheating. A cheating teacher tends to be younger and less qualified than average. She is also more likely to cheat after her incentives change. Because the Chicago data ran from 1993 to 2000, it bracketed the introduction of high-stakes testing in 1996. Sure enough,

there was a pronounced spike in cheating in 1996. Nor was the cheating random. It was the teachers in the lowest-scoring classrooms who were most likely to cheat. It should also be noted that the $25,000 bonus for California teachers was eventually revoked, in part because of suspicions that too much of the money was going to cheaters.

Not every result of the Chicago cheating analysis was so dour. In addition to detecting cheaters, the algorithm could also identify the best teachers in the school system. A good teacher's impact was nearly as distinctive as a cheater's. Instead of getting random answers correct, her students would show real improvement on the easier types of questions they had previously missed, an indication of actual learning. And a good teacher's students carried over all their gains into the next grade.

Most academic analyses of this sort tend to languish, unread, on a dusty library shelf. But in early 2002, the new CEO of the Chicago Public Schools, Arne Duncan, contacted the study's authors. He didn't want to protest or hush up their findings. Rather, he wanted to make sure that the teachers identified by the algorithm as cheaters were truly cheating—and then do something about it.

Duncan was an unlikely candidate to hold such a powerful job. He was only thirty-six when appointed, a onetime academic all-American at Harvard who later played pro basketball in Australia. He had spent just three years with the CPS—and never in a job important enough to have his own secretary—before becoming its CEO. It didn't hurt that Duncan had grown up in Chicago. His father taught psychology at the University of Chicago; his mother ran an afterschool program for forty years, without pay, in a poor neighborhood. When Duncan was a boy, his afterschool playmates were the underprivileged kids his mother cared for. So when he took over the public schools, his allegiance lay more with schoolchildren and their families than with teachers and their union.

The best way to get rid of cheating teachers, Duncan had decided, was to readminister the standardized exam. He only had the resources to retest 120 classrooms, however, so he asked the creators of the cheating algorithm to help choose which classrooms to test.

How could those 120 retests be used most effectively? It might have seemed sensible to retest only the classrooms that likely had a cheating teacher. But even if their retest scores were lower, the teachers could argue that the students did worse merely because they were told that the scores wouldn't count in their official record—which, in fact, all retested students would be told. To make the retest results convincing, some non-cheaters were needed as a control group. The best control group? The classrooms shown by the algorithm to have the best teachers, in which big gains were thought to have been legitimately

attained. If those classrooms held their gains while the classrooms with a suspected cheater lost ground, the cheating teachers could hardly argue that their students did worse only because the scores wouldn't count.

So a blend was settled upon. More than half of the 120 retested classrooms were those suspected of having a cheating teacher. The remainder were divided between the supposedly excellent teachers (high scores but no suspicious answer patterns) and, as a further control, classrooms with mediocre scores and no suspicious answers.

The retest was given a few weeks after the original exam. The children were not told the reason for the retest. Neither were the teachers. But they may have gotten the idea when it was announced that CPS officials, not the teachers, would administer the test. The teachers were asked to stay in the classroom with their students, but they would not be allowed to even touch the answer sheets.

The results were as compelling as the cheating algorithm had predicted. In the classrooms chosen as controls, where no cheating was suspected, scores stayed about the same or even rose. In contrast, the students with the teachers identified as cheaters scored far worse, by an average of more than a full grade level.

As a result, the Chicago Public School system began to fire its cheating teachers. The evidence was only strong enough to get rid of a dozen of them, but the many other cheaters had been duly warned. The final outcome of the Chicago study is further testament to the power of incentives: the following year, cheating by teachers fell more than 30 percent.

You might think that the sophistication of teachers who cheat would increase along with the level of schooling. But an exam given at the University of Georgia in the fall of 2001 disputes that idea. The course was called Coaching Principles and Strategies of Basketball, and the final grade was based on a single exam that had twenty questions. Among the questions:

How many halves are in a college basketball game?

- a. 1 b. 2 c. 3 d. 4

How many points does a 3-pt. field goal account for in a basketball game?

- a. 1 b. 2 c. 3 d. 4

What is the name of the exam which all high school seniors in the state of Georgia must pass?

- a. Eye Exam
- b. How Do the Grits Taste Exam
- c. Bug Control Exam
- d. Georgia Exit Exam

In your opinion, who is the best Division I assistant coach in the country?

- a. Ron Jirsa
- b. John Pelphrey
- c. Jim Harrick Jr.
- d. Steve Wojciechowski

If you are stumped by the final question, it might help to know that Coaching Principles was taught by Jim Harrick Jr., an assistant coach with the university's basketball team. It might also help to know that his father, Jim Harrick Sr., was the head basketball coach. Not surprisingly, Coaching Principles was a favorite course among players on the Harricks' team. Every student in the class received an A. Not long afterward, both Harricks were relieved of their coaching duties.

If it strikes you as disgraceful that Chicago schoolteachers and University of Georgia professors will cheat—a teacher, after all, is meant to instill values along with the facts—then the thought of cheating among sumo wrestlers may also be deeply disturbing. In Japan, sumo is not only the national sport but also a repository of the country's religious, military, and historical emotion. With its purification rituals and its imperial roots, sumo is sacrosanct in a way that American sports will never be. Indeed, sumo is said to be less about competition than about honor itself.

It is true that sports and cheating go hand in hand. That's because cheating is more common in the face of a bright-line incentive (the line between winning

and losing, for instance) than with a murky incentive. Olympic sprinters and weightlifters, cyclists in the Tour de France, football linemen and baseball sluggers: they have all been shown to swallow whatever pill or powder may give them an edge. It is not only the participants who cheat. Cagey baseball managers try to steal an opponent's signs. In the 2002 Winter Olympic figure-skating competition, a French judge and a Russian judge were caught trying to swap votes to make sure their skaters medaled. (The man accused of orchestrating the vote swap, a reputed Russian mob boss named Alimzhan Tokhtakhounov, was also suspected of rigging beauty pageants in Moscow.)

An athlete who gets caught cheating is generally condemned, but most fans at least appreciate his motive: he wanted so badly to win that he bent the rules. (As the baseball player Mark Grace once said, "If you're not cheating, you're not trying.") An athlete who cheats to *lose*, meanwhile, is consigned to a deep circle of sporting hell. The 1919 Chicago White Sox, who conspired with gamblers to throw the World Series (and are therefore known forever as the Black Sox), retain a stench of iniquity among even casual baseball fans. The City College of New York's championship basketball team, once beloved for its smart and scrappy play, was instantly reviled when it was discovered in 1951 that several players had taken mob money to shave points—intentionally missing baskets to help gamblers beat the point spread. Remember Terry Malloy, the tormented former boxer played by Marlon Brando in *On the Waterfront*? As Malloy saw it, all his troubles stemmed from the one fight in which he took a dive. Otherwise, he could have had class; he could have been a contender.

If cheating to lose is sport's premier sin, and if sumo wrestling is the premier sport of a great nation, cheating to lose couldn't possibly exist in sumo. Could it?

Once again, the data can tell the story. As with the Chicago school tests, the data set under consideration here is surpassingly large: the results from nearly every official match among the top rank of Japanese sumo wrestlers between January 1989 and January 2000, a total of 32,000 bouts fought by 281 different wrestlers.

The incentive scheme that rules sumo is intricate and extraordinarily powerful. Each wrestler maintains a ranking that affects every slice of his life: how much money he makes, how large an entourage he carries, how much he gets to eat, sleep, and otherwise take advantage of his success. The sixty-six highest-ranked wrestlers in Japan, comprising the *makuuchi* and *juryo* divisions, make up the sumo elite. A wrestler near the top of this elite pyramid may earn millions and is treated like royalty. Any wrestler in the top forty earns at least $170,000 a year. The seventieth-ranked wrestler in Japan, meanwhile, earns only

$15,000 a year. Life isn't very sweet outside the elite. Low-ranked wrestlers must tend to their superiors, preparing their meals, cleaning their quarters, and even soaping up their hardest-to-reach body parts. So ranking is everything.

A wrestler's ranking is based on his performance in the elite tournaments that are held six times a year. Each wrestler has fifteen bouts per tournament, one per day over fifteen consecutive days. If he finishes the tournament with a winning record (eight victories or better), his ranking will rise. If he has a losing record, his ranking falls. If it falls far enough, he is booted from the elite rank entirely. The eighth victory in any tournament is therefore critical, the difference between promotion and demotion; it is roughly four times as valuable in the rankings as the typical victory.

So a wrestler entering the final day of a tournament on the bubble, with a 7–7 record, has far more to gain from a victory than an opponent with a record of 8–6 has to lose.

Is it possible, then, that an 8–6 wrestler might allow a 7–7 wrestler to beat him? A sumo bout is a concentrated flurry of force and speed and leverage, often lasting only a few seconds. It wouldn't be very hard to let yourself be tossed. Let's imagine for a moment that sumo wrestling *is* rigged. How might we measure the data to prove it?

The first step would be to isolate the bouts in question: those fought on a tournament's final day between a wrestler on the bubble and a wrestler who has already secured his eighth win. (Because more than half of all wrestlers end a tournament with either seven, eight, or nine victories, hundreds of bouts fit these criteria.) A final-day match between two 7–7 wrestlers isn't likely to be fixed, since both fighters badly need the victory. A wrestler with ten or more victories probably wouldn't throw a match either, since he has his own strong incentive to win: the $100,000 prize for overall tournament champion and a series of $20,000 prizes for the "outstanding technique" award, "fighting spirit" award, and others.

Let's now consider the following statistic, which represents the hundreds of matches in which a 7–7 wrestler faced an 8–6 wrestler on a tournament's final day. The left column tallies the probability, based on all past meetings between the two wrestlers fighting that day, that the 7–7 wrestler will win. The right column shows how often the 7–7 wrestler actually did win.

7–7 Wrestler's Predicted Win Percentage Against8–6 Opponent	7–7 Wrestler's Actual Win Percentage Against 8–6 Opponent
48.7	79.6

So the 7–7 wrestler, based on past outcomes, was expected to win just less than half the time. This makes sense; their records in this tournament indicate that the 8–6 wrestler is slightly better. But in actuality, the wrestler on the bubble won *almost eight out of ten* matches against his 8–6 opponent. Wrestlers on the bubble also do astonishingly well against 9–5 opponents:

7–7 Wrestler's Predicted Win Percentage Against 9–5 Opponent	7–7 Wrestler's Actual Win Percentage Against 9–5 Opponent
47.2	73.4

As suspicious as this looks, a high winning percentage alone isn't enough to prove that a match is rigged. Since so much depends on a wrestler's eighth win, he should be expected to fight harder in a crucial bout. But perhaps there are further clues in the data that prove collusion.

It's worth thinking about the incentive a wrestler might have to throw a match. Maybe he accepts a bribe (which would obviously not be recorded in the data). Or perhaps some other arrangement is made between the two wrestlers. Keep in mind that the pool of elite sumo wrestlers is extraordinarily tight-knit. Each of the sixty-six elite wrestlers fights fifteen of the others in a tournament every two months. Furthermore, each wrestler belongs to a stable that is typically managed by a former sumo champion, so even the rival stables have close ties. (Wrestlers from the same stable do not wrestle one another.)

Now let's look at the win-loss percentage between the 7–7 wrestlers and the 8–6 wrestlers the *next* time they meet, when neither one is on the bubble. In this case, there is no great pressure on the individual match. So you might expect the wrestlers who won their 7–7 matches in the previous tournament to do about as well as they had in earlier matches against these same opponents—that is, winning roughly 50 percent of the time. You certainly wouldn't expect them to uphold their 80 percent clip.

As it turns out, the data show that the 7–7 wrestlers win only 40 percent of the rematches. Eighty percent in one match and 40 percent in the next? How do you make sense of that?

The most logical explanation is that the wrestlers made a quid pro quo agreement: you let me win today, when I really need the victory, and I'll let you win the next time. (Such an arrangement wouldn't preclude a cash bribe.) It's especially interesting to note that by the two wrestlers' *second* subsequent meeting, the win percentages revert to the expected level of about 50 percent,

suggesting that the collusion spans only two matches.

And it isn't only the individual wrestlers whose records are suspect. The collective records of the various sumo stables are similarly aberrational. When one stable's wrestlers fare well on the bubble against wrestlers from a second stable, they tend to do especially *poorly* when the second stable's wrestlers are on the bubble. This indicates that some match rigging may be choreographed at the highest level of the sport—much like the Olympic skating judges' vote swapping.

No formal disciplinary action has ever been taken against a Japanese sumo wrestler for match rigging. Officials from the Japanese Sumo Association typically dismiss any such charges as fabrications by disgruntled former wrestlers. In fact, the mere utterance of the words "sumo" and "rigged" in the same sentence can cause a national furor. People tend to get defensive when the integrity of their national sport is impugned.

Still, allegations of match rigging do occasionally find their way into the Japanese media. These occasional media storms offer one more chance to measure possible corruption in sumo. Media scrutiny, after all, creates a powerful incentive: if two sumo wrestlers or their stables *have* been rigging matches, they might be leery to continue when a swarm of journalists and TV cameras descend upon them.

So what happens in such cases? The data show that in the sumo tournaments held immediately after allegations of match rigging, 7–7 wrestlers win only 50 percent of their final-day matches against 8–6 opponents instead of the typical 80 percent. No matter how the data are sliced, they inevitably suggest one thing: it is hard to argue that sumo wrestling isn't rigged.

Several years ago, two former sumo wrestlers came forward with extensive allegations of match rigging—and more. Aside from the crooked matches, they said, sumo was rife with drug use and sexcapades, bribes and tax evasion, and close ties to the *yakuza*, the Japanese mafia. The two men began to receive threatening phone calls; one of them told friends he was afraid he would be killed by the *yakuza*. Still, they went forward with plans to hold a press conference at the Foreign Correspondents' Club in Tokyo. But shortly beforehand, the two men died—hours apart, in the same hospital, of a similar respiratory ailment. The police declared there had been no foul play but did not conduct an investigation. "It seems very strange for these two people to die on the same day at the same hospital," said Mitsuru Miyake, the editor of a sumo magazine. "But no one has seen them poisoned, so you can't prove the skepticism."

Whether or not their deaths were intentional, these two men had done what

no other sumo insider had previously done: named names. Of the 281 wrestlers covered in the data cited above, they identified 29 crooked wrestlers and 11 who were said to be incorruptible.

What happens when the whistle-blowers' corroborating evidence is factored into the analysis of the match data? In matches between two supposedly corrupt wrestlers, the wrestler who was on the bubble won about 80 percent of the time. In bubble matches against a supposedly clean opponent, meanwhile, the bubble wrestler was no more likely to win than his record would predict. Furthermore, when a supposedly corrupt wrestler faced an opponent whom the whistle-blowers did not name as either corrupt or clean, the results were nearly as skewed as when two corrupt wrestlers met—suggesting that most wrestlers who *weren't* specifically named were also corrupt.

So if sumo wrestlers, schoolteachers, and day-care parents all cheat, are we to assume that mankind is innately and universally corrupt? And if so, how corrupt?

The answer may lie in…bagels. Consider this story about a man named Paul Feldman.

Once upon a time, Feldman dreamed big dreams. With early training in agricultural economics, he wanted to tackle world hunger. Instead, he took a job in Washington, analyzing weapons expenditures for the U.S. Navy. This was in 1962. For the next twenty-odd years, he did further analytic work in Washington. He held senior-level jobs and earned good money, but he wasn't always recognized for his best work. At the office Christmas party, colleagues would introduce him to their wives not as "the head of the public research group" (which he was) but as "the guy who brings in the bagels."

The bagels had begun as a casual gesture: a boss treating his employees whenever they won a research contract. Then he made it a habit. Every Friday, he would bring in some bagels, a serrated knife, and cream cheese. When employees from neighboring floors heard about the bagels, they wanted some too. Eventually he was bringing in fifteen dozen bagels a week. In order to recoup his costs, he set out a cash basket and a sign with the suggested price. His collection rate was about 95 percent; he attributed the underpayment to oversight, not fraud.

In 1984, when his research institute fell under new management, Feldman took a look at his future and grimaced. He decided to quit his job and sell bagels. His economist friends thought he had lost his mind, but his wife supported him. The last of their three children was finishing college, and they had retired their

mortgage.

Driving around the office parks that encircle Washington, he solicited customers with a simple pitch: early in the morning, he would deliver some bagels and a cash basket to a company's snack room; he would return before lunch to pick up the money and the leftovers. It was an honor-system commerce scheme, and it worked. Within a few years, Feldman was delivering 8,400 bagels a week to 140 companies and earning as much as he had ever made as a research analyst. He had thrown off the shackles of cubicle life and made himself happy.

He had also—quite without meaning to—designed a beautiful economic experiment. From the beginning, Feldman kept rigorous data on his bagel business. So by measuring the money collected against the bagels taken, he found it possible to tell, down to the penny, just how honest his customers were. Did they steal from him? If so, what were the characteristics of a company that stole versus a company that did not? Under what circumstances did people tend to steal more, or less?

As it happens, Feldman's accidental study provides a window onto a form of cheating that has long stymied academics: white-collar crime. (Yes, shorting the bagel man is white-collar crime, writ however small.) It might seem ludicrous to address as large and intractable a problem as white-collar crime through the life of a bagel man. But often a small and simple question can help chisel away at the biggest problems.

Despite all the attention paid to rogue companies like Enron, academics know very little about the practicalities of white-collar crime. The reason? There are no good data. A key fact of white-collar crime is that we hear about only the very slim fraction of people who are *caught* cheating. Most embezzlers lead quiet and theoretically happy lives; employees who steal company property are rarely detected.

With street crime, meanwhile, that is not the case. A mugging or a burglary or a murder is usually tallied whether or not the criminal is caught. A street crime has a victim, who typically reports the crime to the police, who generate data, which in turn generate thousands of academic papers by criminologists, sociologists, and economists. But white-collar crime presents no obvious victim. From whom, exactly, did the masters of Enron steal? And how can you measure something if you don't know to whom it happened, or with what frequency, or in what magnitude?

Paul Feldman's bagel business was different. It did present a victim. The victim was Paul Feldman.

When he started his business, he expected a 95 percent payment rate, based on the experience at his own office. But just as crime tends to be low on a street where a police car is parked, the 95 percent rate was artificially high: Feldman's presence had deterred theft. Not only that, but those bagel eaters knew the provider and had feelings (presumably good ones) about him. A broad swath of psychological and economic research has shown that people will pay different amounts for the same item depending on who is providing it. The economist Richard Thaler, in his 1985 "Beer on the Beach" study, showed that a thirsty sunbather would pay $2.65 for a beer delivered from a resort hotel but only $1.50 for the same beer if it came from a shabby grocery store.

In the real world, Feldman learned to settle for less than 95 percent. He came to consider a company "honest" if its payment rate was above 90 percent. He considered a rate between 80 and 90 percent "annoying but tolerable." If a company habitually paid below 80 percent, Feldman might post a hectoring note, like this one:

> The cost of bagels has gone up dramatically since the beginning of the year. Unfortunately, the number of bagels that disappear without being paid for has also gone up. Don't let that continue. I don't imagine that you would teach your children to cheat, so why do it yourselves?

In the beginning, Feldman left behind an open basket for the cash, but too often the money vanished. Then he tried a coffee can with a money slot in its plastic lid, which also proved too tempting. In the end, he resorted to making small plywood boxes with a slot cut into the top. The wooden box has worked well. Each year he drops off about seven thousand boxes and loses, on average, just one to theft. This is an intriguing statistic: the same people who routinely steal more than 10 percent of his bagels almost never stoop to stealing his money box—a tribute to the nuanced social calculus of theft. From Feldman's perspective, an office worker who eats a bagel without paying is committing a crime; the office worker probably doesn't think so. This distinction probably has less to do with the admittedly small amount of money involved (Feldman's bagels cost one dollar each, cream cheese included) than with the context of the "crime." The same office worker who fails to pay for his bagel might also help himself to a long slurp of soda while filling a glass in a self-serve restaurant, but he is very unlikely to leave the restaurant without paying.

So what do the bagel data have to say? In recent years, there have been two

noteworthy trends in the overall payment rate. The first was a long, slow decline that began in 1992. By the summer of 2001, the overall rate had slipped to about 87 percent. But immediately after September 11 of that year, the rate spiked a full 2 percent and hasn't slipped much since. (If a 2 percent gain in payment doesn't sound like much, think of it this way: the nonpayment rate fell from 13 to 11 percent, which amounts to a 15 percent decline in theft.) Because many of Feldman's customers are affiliated with national security, there may have been a patriotic element to this 9/11 Effect. Or it may have represented a more general surge in empathy.

The data also show that smaller offices are more honest than big ones. An office with a few dozen employees generally outpays by 3 to 5 percent an office with a few hundred employees. This may seem counterintuitive. In a bigger office, a bigger crowd is bound to convene around the bagel table, providing more witnesses to make sure you drop your money in the box. But in the big-office/small-office comparison, bagel crime seems to mirror street crime. There is far less street crime per capita in rural areas than in cities, in large part because a rural criminal is more likely to be known (and therefore caught). Also, a smaller community tends to exert greater social incentives against crime, the main one being shame.

The bagel data also reflect how much personal mood seems to affect honesty. Weather, for instance, is a major factor. Unseasonably pleasant weather inspires people to pay at a higher rate. Unseasonably cold weather, meanwhile, makes people cheat prolifically; so do heavy rain and wind. Worst are the holidays. The week of Christmas produces a 2 percent drop in payment rates—again, a 15 percent increase in theft, an effect on the same magnitude, in reverse, as that of 9/11. Thanksgiving is nearly as bad; the week of Valentine's Day is also lousy, as is the week straddling April 15. There are, however, several good holidays: the weeks that include the Fourth of July, Labor Day, and Columbus Day. The difference in the two sets of holidays? The low-cheating holidays represent little more than an extra day off from work. The high-cheating holidays are fraught with miscellaneous anxieties and the high expectations of loved ones.

Feldman has also reached some of his own conclusions about honesty, based more on his experience than the data. He has come to believe that morale is a big factor—that an office is more honest when the employees like their boss and their work. He also believes that employees further up the corporate ladder cheat more than those down below. He got this idea after delivering for years to one company spread out over three floors—an executive floor on top and two lower floors with sales, service, and administrative employees. (Feldman

wondered if perhaps the executives cheated out of an overdeveloped sense of entitlement. What he didn't consider is that perhaps cheating was how they got to *be* executives.)

If morality represents the way we would like the world to work and economics represents how it actually does work, then the story of Feldman's bagel business lies at the very intersection of morality and economics. Yes, a lot of people steal from him, but the vast majority, even though no one is watching over them, do not. This outcome may surprise some people—including Feldman's economist friends, who counseled him twenty years ago that his honor-system scheme would never work. But it would not have surprised Adam Smith. In fact, the theme of Smith's first book, *The Theory of Moral Sentiments*, was the innate honesty of mankind. "How selfish soever man may be supposed," Smith wrote, "there are evidently some principles in his nature, which interest him in the fortune of others, and render their happiness necessary to him, though he derives nothing from it, except the pleasure of seeing it."

There is a tale, "The Ring of Gyges," that Feldman sometimes tells his economist friends. It comes from Plato's *Republic*. A student named Glaucon offered the story in response to a lesson by Socrates—who, like Adam Smith, argued that people are generally good even without enforcement. Glaucon, like Feldman's economist friends, disagreed. He told of a shepherd named Gyges who stumbled upon a secret cavern with a corpse inside that wore a ring. When Gyges put on the ring, he found that it made him invisible. With no one able to monitor his behavior, Gyges proceeded to do woeful things—seduce the queen, murder the king, and so on. Glaucon's story posed a moral question: could any man resist the temptation of evil if he knew his acts could not be witnessed? Glaucon seemed to think the answer was no. But Paul Feldman sides with Socrates and Adam Smith—for he knows that the answer, at least 87 percent of the time, is yes.

2

How Is the Ku Klux Klan Like a Group of Real-Estate Agents?

As institutions go, the Ku Klux Klan has had a markedly up-and-down history. It was founded in the immediate aftermath of the Civil War by six former Confederate soldiers in Pulaski, Tennessee. The six young men, four of whom were budding lawyers, saw themselves as merely a circle of like-minded friends. Thus the name they chose, "kuklux," a slight mangling of *kuklos*, the Greek word for "circle." In the beginning, their activities were said to be harmless midnight pranks—for instance, riding horses through the countryside while draped in white sheets and pillowcase hoods. But soon the Klan evolved into a multistate terrorist organization designed to frighten and kill emancipated slaves. Among its regional leaders were five former Confederate generals; its staunchest supporters were the plantation owners for whom Reconstruction posed an economic and political nightmare. In 1872, President Ulysses S. Grant spelled out for the House of Representatives the true aims of the Ku Klux Klan: "By force and terror, to prevent all political action not in accord with the views of the members, to deprive colored citizens of the right to bear arms and of the right of a free ballot, to suppress the schools in which colored children were taught, and to reduce the colored people to a condition closely allied to that of slavery."

The early Klan did its work through pamphleteering, lynching, shooting, burning, castrating, pistol-whipping, and a thousand forms of intimidation. They targeted former slaves and any whites who supported the blacks' rights to vote, acquire land, or gain an education. But within barely a decade, the Klan had been extinguished, largely by legal and military interventions out of Washington, D.C.

If the Klan itself was defeated, however, its aims had largely been achieved through the establishment of Jim Crow laws. Congress, which during Reconstruction had been quick to enact measures of legal, social, and economic freedom for blacks, just as quickly began to roll them back. The federal

government agreed to withdraw its occupation troops from the South, allowing the restoration of white rule. In *Plessy v. Ferguson*, the U.S. Supreme Court gave the go-ahead to full-scale racial segregation.

The Ku Klux Klan lay largely dormant until 1915, when D. W. Griffith's film *The Birth of a Nation* (originally titled *The Clansman*) helped spark its rebirth. Griffith presented the Klan as crusaders for white civilization itself, and as one of the noblest forces in American history. The film quoted a line from *A History of the American People*, written by a renowned historian: "At last there had sprung into existence a great Ku Klux Klan, a veritable empire of the South, to protect the Southern country." The historian in question was U.S. president Woodrow Wilson, onetime scholar and president of Princeton University.

By the 1920s, a revived Klan claimed eight million members. This time around, the Klan was not confined to the South but ranged throughout the country; this time, it concerned itself not only with blacks but also with Catholics, Jews, communists, unionists, immigrants, agitators, and other disrupters of the status quo. In 1933, with Hitler ascendant in Germany, Will Rogers was the first to draw a line between the new Klan and the new threat in Europe: "Papers all state Hitler is trying to copy Mussolini," he wrote. "Looks to me like it's the Ku Klux that he is copying."

The onset of World War II and a number of internal scandals once again laid the Klan low. Public sentiment turned against the Klan as the unity of a country at war trumped its message of separatism.

But within a few years, there were already signs of a massive revival. As wartime anxiety gave way to postwar uncertainty, Klan membership flourished. Barely two months after V-J Day, the Klan in Atlanta burned a 300-foot cross on the face of Stone Mountain, site of a storied rock carving of Robert E. Lee. The extravagant cross burning, one Klansman later said, was intended "just to let the niggers know the war is over and that the Klan is back on the market."

Atlanta had by now become Klan headquarters. The Klan was thought to hold great sway with key Georgia politicians, and its Georgia chapters were said to include many policemen and sheriff's deputies. Yes, the Klan was a secret society, reveling in passwords and cloak-and-dagger ploys, but its real power lay in the very public fear that it fostered, exemplified by the open secret that the Ku Klux Klan and the law-enforcement establishment were brothers in arms.

Atlanta—the Imperial City of the KKK's Invisible Empire, in Klan jargon—was also home to Stetson Kennedy, a thirty-year-old man with the bloodlines of a Klansman but a temperament that ran opposite. He came from a good southern family which claimed ancestors including two signers of the Declaration of Independence, an officer in the Confederate Army, and John B.

Stetson, founder of the famed hat company and the man for whom Stetson University was named.

Stetson Kennedy grew up in a fourteen-room house in Jacksonville, Florida, the youngest of five children. His uncle Brady was a Klansman. But Kennedy would go on to become a self-described "dissident at large," writing numberless articles and several books that railed against bigotry. He first worked as a folklorist, traveling around Florida to collect old native tales and songs. Years later, when he served as a rare white correspondent for the *Pittsburgh Courier*, the country's largest black newspaper, he wrote under the pseudonym Daddy Mention—after a black folk hero who, as myth told it, could outrun the blast of a sheriff's shotgun.

What drove Kennedy was a hatred of small-mindedness, ignorance, obstructionism, and intimidation—which, in his view, were displayed by no organization more proudly than the Ku Klux Klan. Kennedy saw the Klan as the terrorist arm of the white establishment itself. This struck him as an intractable problem, for a variety of reasons. The Klan was in cahoots with political, business, and law-enforcement leaders. The public was frightened and felt powerless to act against the Klan. And the few anti-hate groups that existed at the time had little leverage or even information about the Klan. "Almost all of the things written on the subject were editorials, not exposés," Kennedy would later explain. "The writers were *against* the Klan, all right, but they had precious few inside facts *about* it."

So Kennedy set out to gather those facts. He would spend years interviewing Klan leaders and sympathizers, sometimes taking advantage of his own background and lineage to pretend that he was on their side of the issues. He also attended public Klan events and, as he would later write, he even set about to infiltrate the Klan in Atlanta.

The Klan Unmasked, Kennedy's memoir of his exploits "inside" the Klan, is in fact more of a novelization than a straight nonfiction account. Kennedy, a folklorist at heart, apparently wanted to put across the most dramatic story possible, and therefore included not only his own anti-Klan activities but those of another man, code-named John Brown. Brown was a union worker and a former Klan official who had changed his ways and offered to infiltrate the Klan. It was John Brown who apparently performed many of the most dramatic and dangerous episodes portrayed in *The Klan Unmasked*—physically attending Klan meetings and other functions in Atlanta—but since Stetson Kennedy was the man who later wrote the book, he rendered Brown's actions as his own.

Regardless, there was a great deal of information to be gleaned from this Brown/Kennedy collaboration. Brown divulged what he was learning at the

weekly Klan meetings: the identities of the Klan's local and regional leaders; their upcoming plans; the Klan's current rituals, passwords, and language. It was Klan custom, for instance, to append a *Kl* to many words. (Thus would two Klansmen hold a Klonversation in the local Klavern.) The secret Klan handshake was a left-handed, limp-wristed fish wiggle. When a traveling Klansman wanted to locate brethren in a strange town, he would ask for a "Mr. Ayak"—"Ayak" being code for "Are You a Klansman?" He would hope to hear this response: "Yes, and I also know a Mr. Akai"—code for "A Klansman Am I."

Before long, John Brown was invited to join the Klavaliers, the Klan's secret police and "flog squad." For an infiltrator, this posed a particularly sticky problem: What would happen if he were called upon to inflict violence?

But as it happened, a central tenet of life in the Klan—and of terrorism in general—is that most of the threatened violence never goes beyond the threat stage.

Consider lynching, the Klan's hallmark sign of violence. Here, compiled by the Tuskegee Institute, are the decade-by-decade statistics on the lynching of blacks in the United States:

YEARS	LYNCHINGS OF BLACKS
1890–1899	1,111
1900–1909	791
1910–1919	569
1920–1929	281
1930–1939	119
1940–1949	31
1950–1959	6
1960–1969	3

Bear in mind that these figures represent not only lynchings attributed to the Ku Klux Klan but the total number of reported lynchings. The statistics reveal at least three noteworthy facts. The first is the obvious decrease in lynchings over time. The second is the absence of a correlation between lynchings and Klan membership: there were actually *more* lynchings of blacks between 1900 and 1909, when the Klan was dormant, than during the 1920s, when the Klan had millions of members—which suggests that the Ku Klux Klan carried out far fewer lynchings than is generally thought.

Third, relative to the size of the black population, lynchings were exceedingly rare. To be sure, one lynching is one too many. But by the turn of the century, lynchings were hardly the everyday occurrence that they are often considered in the public recollection. Compare the 281 victims of lynchings in the 1920s to the number of black infants who were dying at that time as a result of malnutrition, pneumonia, diarrhea, and the like. As of 1920, about 13 out of every 100 black children died in infancy, or roughly 20,000 children each year—compared to 28 people who were lynched in a year. As late as 1940, about 10,000 black infants died each year.

What larger truths do these lynching figures suggest? What does it mean that lynchings were relatively rare and that they fell precipitously over time, even in the face of a boom in Klan membership?

The most compelling explanation is that all those early lynchings *worked*. White racists—whether or not they belonged to the Ku Klux Klan—had through their actions and their rhetoric developed a strong incentive scheme that was terribly clear and terribly frightening. If a black person violated the accepted code of behavior, whether by talking back to a bus driver or daring to try to vote, he knew he might well be punished, perhaps by death.

So it may be that by the mid-1940s, when Stetson Kennedy was trying to bust up the Klan, it didn't really *need* to use as much violence. Many blacks, having long been told to behave like second-class citizens—or else—simply obliged. One or two lynchings went a long way toward inducing docility among even a large group of people, for people respond strongly to strong incentives. And there are few incentives more powerful than the fear of random violence—which, in essence, is why terrorism is so effective.

But if the Ku Klux Klan of the 1940s wasn't uniformly violent, what was it? The Klan that Stetson Kennedy wrote about was in fact a sorry fraternity of men, most of them poorly educated and with poor prospects, who needed a place to vent—and an excuse for occasionally staying out all night. That their fraternity engaged in quasi-religious chanting and oath taking and hosanna hailing, all of it top secret, made it that much more appealing.

Kennedy also found the Klan to be a slick money-making operation, at least for those near the top of the organization. Klan leaders had any number of revenue sources: thousands of dues-paying rank-and-file members; business owners who hired the Klan to scare off the unions or who paid the Klan protection money; Klan rallies that generated huge cash donations; even the occasional gunrunning or moonshine operation. Then there were rackets like the Klan's Death Benefit Association, which sold insurance policies to Klan members and accepted only cash or personal checks made out to the Grand

Dragon himself.

And, even though the Klan may not have been as deadly as generally thought, it was plenty violent and, perhaps worse, had ever greater designs on political influence. Kennedy was therefore eager to damage the Klan in any way he could. When he heard about Klan plans for a union-busting rally, he fed the information to a union friend. He passed along Klan information to the assistant attorney general of Georgia, an established Klan buster. After researching the Klan's corporate charter, Kennedy wrote to the governor of Georgia suggesting the grounds upon which the charter should be revoked: the Klan had been designated a non-profit, non-political organization, but Kennedy had proof that it was clearly devoted to both profits and politics.

The problem was that most of Kennedy's efforts weren't producing the desired effect. The Klan was so entrenched and broad-based that Kennedy felt as if he were tossing pebbles at a giant. And even if he could somehow damage the Klan in Atlanta, the hundreds of other chapters around the country would go untouched.

Kennedy was supremely frustrated, and out of this frustration was born a new strategy. He had noticed one day a group of young boys playing some kind of spy game in which they exchanged silly secret passwords. It reminded him of the Klan. Wouldn't it be nice, he thought, to get the Klan's passwords and the rest of its secrets into the hands of kids all across the country—and their parents too? What better way to defang a secret society than to make public its most secret information? Instead of futilely attacking the Klan from the outside, what if he could somehow unleash all the secret inside information that John Brown was gathering from the Klan's weekly meetings? Between Brown's inside dope and everything that Kennedy had learned via his own investigations, he probably knew more Klan secrets than the average Klansman.

Kennedy turned to the most powerful mass medium of his day: radio. He began feeding Klan reports to the journalist Drew Pearson, whose *Washington Merry-Go-Round* program was heard by millions of adults every day, and to the producers of the *Adventures of Superman* show, which reached millions of children each night. He told them about Mr. Ayak and Mr. Akai, and he passed along overheated passages from the Klan's bible, which was called the Kloran. (Kennedy never did learn why a white Christian supremacist group would give its bible essentially the same name as the most holy book of Islam.) He explained the role of Klan officers in any local Klavern: the Klaliff (vice president), Klokard (lecturer), Kludd (chaplain), Kligrapp (secretary), Klabee (treasurer), Kladd (conductor), Klarogo (inner guard), Klexter (outer guard), the Klokann (a five-man investigative committee), and the Klavaliers (whose leader

was called Chief Ass Tearer). He spelled out the Klan hierarchy as it proceeded from the local to the national level: an Exalted Cyclops and his twelve Terrors; a Great Titan and his twelve Furies; a Grand Dragon and his nine Hydras; and the Imperial Wizard and his fifteen Genii. And Kennedy passed along all the information and gossip that John Brown gleaned by infiltrating the main Klan chapter, Nathan Bedford Forrest Klavern No. 1, Atlanta, Realm of Georgia.

During the war, the *Adventures of Superman* program had portrayed its hero fighting Hitler and Mussolini and Hirohito. But now he was in need of fresh villains. The Klan was a perfect target, and Superman turned his powers against them. Drew Pearson, an avowed Klan hater, now began giving regular Klan updates on his radio show, and then gave further updates, based on John Brown's inside reports, to show how the original updates were infuriating Klan officials. Pearson's work created an echo chamber that seemed to be driving Grand Dragon Samuel Green crazy. Here is Pearson's radio report from November 17, 1948:

> Speaking at Klavern No. 1, Atlanta, Ga., the week after elections, the Grand Dragon wrung his hands and once again cautioned Klansmen to be careful about leaks.
>
> "I have to talk frankly at these meetings," he said, "but I might as well call Drew Pearson before I come to the meeting and give him the information, for [the] next day he gives it out to everybody from coast to coast. The A.P. and U.P. are both calling me about it next morning while I am eating breakfast."…
>
> The Grand Dragon spoke about plans for a big cross-burning to be held in Macon, Ga., on Dec. 10. It would be the biggest in Klan history, he said, and he expected 10,000 Klansmen to be there—in their robes….
>
> He added that the Klavalier Klub—the Klan's whipping and flogging department—was now on the job and had plenty of friends on the Atlanta police force.

As the Pearson and *Superman* radio shows played on, and as Stetson Kennedy continued to relay the Klan secrets obtained by John Brown to other broadcast and print outlets, a funny thing happened: attendance at Klan meetings began to fall, as did applications for new membership. Of all the ideas that Kennedy had thought up to fight bigotry, this campaign was easily the cleverest. He turned the Klan's secrecy against itself by making its private information

public; he converted heretofore precious knowledge into ammunition for mockery.

Americans who might have been philosophically inclined to oppose the Klan had now been given enough specific information to oppose them more actively, and public sentiment began to shift. Americans who might have been philosophically inclined to *embrace* the Klan had now been given all sorts of caution against doing so. Although the Klan would never quite die, especially down south—David Duke, a smooth-talking Klan leader from Louisiana, mounted substantive bids for the U.S. Senate and other offices—it was certainly handicapped, at least in the short term, by Kennedy's brazen dissemination of inside information. While it is impossible to tease out the exact impact that his work had on the Klan, many people have given him a great deal of credit for damaging an institution that was in grave need of being damaged.

This did not come about because Stetson Kennedy was courageous or resolute or unflappable, even though he was all of these. It happened because he understood the raw power of information. The Ku Klux Klan—much like politicians or real-estate agents or stockbrokers—was a group whose power was derived in large part from the fact that it hoarded information. Once that information falls into the wrong hands (or, depending on your point of view, the *right* hands), much of the group's advantage disappears.

In the late 1990s, the price of term life insurance fell dramatically. This posed something of a mystery, for the decline had no obvious cause. Other types of insurance, including health and automobile and homeowners' coverage, were certainly not falling in price. Nor had there been any radical changes among insurance companies, insurance brokers, or the people who buy term life insurance. So what happened?

The Internet happened. In the spring of 1996, Quotesmith.com became the first of several websites that enabled a customer to compare, within seconds, the price of term life insurance sold by dozens of different companies. For such websites, term life insurance was a perfect product. Unlike other forms of insurance—especially whole life insurance, which is a far more complicated financial instrument—term life policies are fairly homogeneous: any given thirty-year, guaranteed policy for $1 million is essentially identical to the next. So what really matters is the price. Shopping around for the cheapest policy, a process that had been convoluted and time-consuming, was suddenly made simple. With customers able to instantaneously find the cheapest policy, the more expensive companies had no choice but to lower their prices. Suddenly

customers were paying $1 billion less a year for term life insurance.

It is worth noting that these websites only listed prices; they didn't even sell the policies. So it wasn't really insurance they were peddling. Like Stetson Kennedy, they were dealing in information. (Had the Internet been around when Kennedy was attacking the Klan, he probably would have been blogging his brains out.) To be sure, there are differences between exposing the Ku Klux Klan and exposing insurance companies' high premiums. The Klan trafficked in secret information whose secrecy engendered fear, while insurance prices were less a secret than a set of facts dispensed in a way that made comparisons difficult. But in both instances, the dissemination of the information diluted its power. As Supreme Court Justice Louis D. Brandeis once wrote, "Sunlight is said to be the best of disinfectants."

Information is a beacon, a cudgel, an olive branch, a deterrent—all depending on who wields it and how. Information is so powerful that the *assumption* of information, even if the information does not actually exist, can have a sobering effect. Consider the case of a one-day-old car.

The day that a car is driven off the lot is the worst day in its life, for it instantly loses as much as a quarter of its value. This might seem absurd, but we know it to be true. A new car that was bought for $20,000 cannot be resold for more than perhaps $15,000. Why? Because the only person who might logically want to resell a brand-new car is someone who found the car to be a lemon. So even if the car isn't a lemon, a potential buyer assumes that it is. He assumes that the seller has some information about the car that he, the buyer, does not have—and the seller is punished for this assumed information.

And if the car *is* a lemon? The seller would do well to wait a year to sell it. By then, the suspicion of lemonness will have faded; by then, some people will be selling their perfectly good year-old cars, and the lemon can blend in with them, likely selling for more than it is truly worth.

It is common for one party to a transaction to have better information than another party. In the parlance of economists, such a case is known as an information asymmetry. We accept as a verity of capitalism that someone (usually an expert) knows more than someone else (usually a consumer). But information asymmetries everywhere have in fact been gravely wounded by the Internet.

Information is the currency of the Internet. As a medium, the Internet is brilliantly efficient at shifting information from the hands of those who have it into the hands of those who do not. Often, as in the case of term life insurance prices, the information existed but in a woefully scattered way. (In such instances, the Internet acts like a gigantic horseshoe magnet waved over an

endless sea of haystacks, plucking the needle out of each one.) The Internet has accomplished what even the most fervent consumer advocates usually cannot: it has vastly shrunk the gap between the experts and the public.

The Internet has proven particularly fruitful for situations in which a face-to-face encounter with an expert might actually *exacerbate* the problem of asymmetrical information—situations in which an expert uses his informational advantage to make us feel stupid or rushed or cheap or ignoble. Consider a scenario in which your loved one has just died and now the funeral director (who knows that you know next to nothing about his business and are under emotional duress to boot) steers you to the $8,000 mahogany casket. Or consider the automobile dealership: a salesman does his best to obscure the car's base price under a mountain of add-ons and incentives. Later, however, in the cool-headed calm of your home, you can use the Internet to find out exactly how much the dealer paid the manufacturer for that car. Or you might just log on to www.TributeDirect.com and buy that mahogany casket yourself for only $3,595, delivered overnight. Unless you decide to spend $2,300 for "The Last Hole" (a casket with golf scenes) or "Memories of the Hunt" (featuring big-racked bucks and other prey) or one of the much cheaper models that the funeral director somehow failed even to mention.

The Internet, powerful as it is, has hardly slain the beast that is information asymmetry. Consider the so-called corporate scandals of the early 2000s. The crimes committed by Enron included hidden partnerships, disguised debt, and the manipulation of energy markets. Henry Blodget of Merrill Lynch and Jack Grubman of Salomon Smith Barney wrote glowing research reports of companies they knew to be junk. Sam Waksal dumped his ImClone stock when he got early word of a damaging report from the Food and Drug Administration; his friend Martha Stewart also dumped her shares, then lied about the reason. WorldCom and Global Crossing fabricated billions of dollars in revenues to pump up their stock prices. One group of mutual fund companies let preferred customers trade at preferred prices, and another group was charged with hiding management fees.

Though extraordinarily diverse, these crimes all have a common trait: they were sins of information. Most of them involved an expert, or a gang of experts, promoting false information or hiding true information; in each case the experts were trying to keep the information asymmetry as asymmetrical as possible.

The practitioners of such acts, especially in the realm of high finance, inevitably offer this defense: "Everybody else was doing it." Which may be

largely true. One characteristic of information crimes is that very few of them are detected. Unlike street crimes, they do not leave behind a corpse or a broken window. Unlike a bagel criminal—that is, someone who eats one of Paul Feldman's bagels but doesn't pay—an information criminal typically doesn't have someone like Feldman tallying every nickel. For an information crime to reach the surface, something drastic must happen. When it does, the results tend to be pretty revealing. The perpetrators, after all, weren't thinking about their private actions being made public. Consider the "Enron tapes," the secretly recorded conversations of Enron employees that surfaced after the company imploded. During a phone conversation on August 5, 2000, two traders chatted about how a wildfire in California would allow Enron to jack up its electricity prices. "The magical word of the day," one trader said, "is 'Burn, Baby, Burn.'" A few months later, a pair of Enron traders named Kevin and Bob talked about how California officials wanted to make Enron refund the profits of its price gouging.

> KEVIN: They're fucking taking all the money back from you guys? All the money you guys stole from those poor grandmas in California?
> BOB: Yeah, Grandma Millie, man.
> KEVIN: Yeah, now she wants her fucking money back for all the power you jammed right up her ass for fucking $250 a megawatt hour.

If you were to assume that many experts use their information to your detriment, you'd be right. Experts depend on the fact that you don't have the information they do. Or that you are so befuddled by the complexity of their operation that you wouldn't know what to do with the information if you had it. Or that you are so in awe of their expertise that you wouldn't dare challenge them. If your doctor suggests that you have angioplasty—even though some current research suggests that angioplasty often does little to prevent heart attacks—you aren't likely to think that the doctor is using his informational advantage to make a few thousand dollars for himself or his buddy. But as David Hillis, an interventional cardiologist at the University of Texas Southwestern Medical Center in Dallas, explained to the *New York Times*, a doctor may have the same economic incentives as a car salesman or a funeral director or a mutual fund manager: "If you're an invasive cardiologist and Joe Smith, the local internist, is sending you patients, and if you tell them they don't need the procedure, pretty soon Joe Smith doesn't send patients anymore."

Armed with information, experts can exert a gigantic, if unspoken, leverage: fear. Fear that your children will find you dead on the bathroom floor of a heart attack if you do not have angioplasty surgery. Fear that a cheap casket will expose your grandmother to a terrible underground fate. Fear that a $25,000 car will crumple like a toy in an accident, whereas a $50,000 car will wrap your loved ones in a cocoon of impregnable steel. The fear created by commercial experts may not quite rival the fear created by terrorists like the Ku Klux Klan, but the principle is the same.

Consider a transaction that wouldn't seem, on the surface, to create much fear: selling your house. What's so scary about that? Aside from the fact that selling a house is typically the largest financial transaction in your life, and that you probably have scant experience in real estate, and that you may have an enormous emotional attachment to your house, there are at least two pressing fears: that you will sell the house for far less than it is worth and that you will not be able to sell it at all.

In the first case, you fear setting the price too low; in the second, you fear setting it too high. It is the job of your real-estate agent, of course, to find the golden mean. She is the one with all the information: the inventory of similar houses, the recent sales trends, the tremors of the mortgage market, perhaps even a lead on an interested buyer. You feel fortunate to have such a knowledgeable expert as an ally in this most confounding enterprise.

Too bad she sees things differently. A real-estate agent may see you not so much as an ally but as a mark. Think back to the study cited at the beginning of this book, which measured the difference between the sale prices of homes that belonged to real-estate agents themselves and the houses they sold for their clients. The study found that an agent keeps her own house on the market an average ten extra days, waiting for a better offer, and sells it for over 3 percent more than your house—or $10,000 on the sale of a $300,000 house. That's $10,000 going into her pocket that does not go into yours, a nifty profit produced by the abuse of information and a keen understanding of incentives. The problem is that the agent only stands to personally gain an additional $150 by selling your house for $10,000 more, which isn't much reward for a lot of extra work. So her job is to convince you that a $300,000 offer is in fact a very good offer, even a generous one, and that only a fool would refuse it.

This can be tricky. The agent does not want to come right out and call you a fool. So she merely implies it—perhaps by telling you about the much bigger, nicer, newer house down the block that has sat unsold for six months. Here is the agent's main weapon: the conversion of information into fear. Consider this true story, related by John Donohue, a law professor who in 2001 was teaching at

Stanford University: “I was just about to buy a house on the Stanford campus,” he recalls, “and the seller’s agent kept telling me what a good deal I was getting because the market was about to zoom. As soon as I signed the purchase contract, he asked me if I would need an agent to sell my previous Stanford house. I told him that I would probably try to sell without an agent, and he replied, ‘John, that might work under normal conditions, but with the market tanking now, you really need the help of a broker.’”

Within five minutes, a zooming market had tanked. Such are the marvels that can be conjured by an agent in search of the next deal.

Consider now another true story of a real-estate agent’s information abuse. The tale involves K., a close friend of one of this book’s authors. K. wanted to buy a house that was listed at $469,000. He was prepared to offer $450,000 but he first called the seller’s agent and asked her to name the lowest price that she thought the homeowner might accept. The agent promptly scolded K. “You ought to be ashamed of yourself,” she said. “That is clearly a violation of real-estate ethics.”

K. apologized. The conversation turned to other, more mundane issues. After ten minutes, as the conversation was ending, the agent told K., “Let me say one last thing. My client is willing to sell this house for a lot less than you might think.”

Based on this conversation, K. then offered $425,000 for the house instead of the $450,000 he had planned to offer. In the end, the seller accepted $430,000. Thanks to *his own agent’s* intervention, the seller lost at least $20,000. The agent, meanwhile, only lost $300—a small price to pay to ensure that she would quickly and easily lock up the sale, which netted her a commission of $6,450.

So a big part of a real-estate agent’s job, it would seem, is to persuade the homeowner to sell for less than he would like while at the same time letting potential buyers know that a house can be bought for less than its listing price. To be sure, there are more subtle means of doing so than coming right out and telling the buyer to bid low. The study of real-estate agents cited above also includes data that reveals how agents convey information through the for-sale ads they write. A phrase like “well maintained,” for instance, is as full of meaning to an agent as “Mr. Ayak” was to a Klansman; it means that a house is old but not quite falling down. A savvy buyer will know this (or find out for himself once he sees the house), but to the sixty-five-year-old retiree who is selling his house, “well maintained” might sound like a compliment, which is just what the agent intends.

An analysis of the language used in real-estate ads shows that certain words are powerfully correlated with the final sale price of a house. This doesn’t

necessarily mean that labeling a house “well maintained” *causes* it to sell for less than an equivalent house. It does, however, indicate that when a real-estate agent labels a house “well maintained,” she may be subtly encouraging a buyer to bid low.

Listed below are ten terms commonly used in real-estate ads. Five of them have a strong positive correlation to the ultimate sale price, and five have a strong negative correlation. Guess which are which.

Ten Common Real-Estate Ad Terms

Fantastic
Granite
Spacious
State-of-the-Art
!
Corian
Charming
Maple
Great Neighborhood
Gourmet

A “fantastic” house is surely fantastic enough to warrant a high price, isn’t it? What about a “charming” and “spacious” house in a “great neighborhood!”? No, no, no, and no. Here’s the breakdown:

Five Terms Correlated to a Higher Sale Price

Granite
State-of-the-Art
Corian
Maple
Gourmet

Five Terms Correlated to a Lower Sale Price

Fantastic
Spacious
!
Charming
Great Neighborhood

Three of the five terms correlated with a higher sale price are physical descriptions of the house itself: granite, Corian, and maple. As information goes, such terms are specific and straightforward—and therefore pretty useful. If you like granite, you might like the house; but even if you don't, "granite" certainly doesn't connote a fixer-upper. Nor does "gourmet" or "state-of-the-art," both of which seem to tell a buyer that a house is, on some level, truly fantastic.

"Fantastic," meanwhile, is a dangerously ambiguous adjective, as is "charming." Both these words seem to be real-estate agent code for a house that doesn't have many specific attributes worth describing. "Spacious" homes, meanwhile, are often decrepit or impractical. "Great neighborhood" signals a buyer that, well, *this* house isn't very nice but others nearby may be. And an exclamation point in a real-estate ad is bad news for sure, a bid to paper over real shortcomings with false enthusiasm.

If you study the words in ads for a real-estate agent's *own* home, meanwhile, you see that she indeed emphasizes descriptive terms (especially "new," "granite," "maple," and "move-in condition") and avoids empty adjectives (including "wonderful," "immaculate," and the telltale "!"). Then she patiently waits for the best buyer to come along. She might tell this buyer about a house nearby that just sold for $25,000 *above* the asking price, or another house that is currently the subject of a bidding war. She is careful to exercise every advantage of the information asymmetry she enjoys.

Does this make her a bad person? That's hard to say, at least hard for *us* to say. The point here is not that real-estate agents are bad people, but that they simply *are* people—and people inevitably respond to incentives. The incentives of the real-estate business, as currently configured, plainly encourage some agents to act against the best interests of their customers.

But like the funeral director and the car salesman and the life-insurance company, the real-estate agent has also seen her advantage eroded by the Internet. After all, anyone selling a home can now get online and gather her own information about sales trends and housing inventory and mortgage rates. The information has been set loose. And recent sales data show the results. Real-estate agents still get a higher price for their own homes than comparable homes

owned by their clients, but since the proliferation of real-estate websites, the gap between the two prices has shrunk by a third.

It would be naïve to suppose that people abuse information only when they are acting as experts or as agents of commerce. After all, agents and experts are people too—which suggests that we are likely to abuse information in our personal lives as well, whether by withholding true information or editing the information we choose to put forth. A real-estate agent may wink and nod when she lists a "well-maintained" house, but we each have our equivalent hedges.

Think about how you describe yourself during a job interview versus how you might describe yourself on a first date. (For even more fun, compare that first-date conversation to a conversation with the same person during your tenth year of marriage.) Or think about how you might present yourself if you were going on national television for the first time. What sort of image would you want to project? Perhaps you want to seem clever or kind or good-looking; presumably you *don't* want to come off as cruel or bigoted. During the heyday of the Ku Klux Klan, its members took pride in publicly disparaging anybody who wasn't a conservative white Christian. But public bigotry has since been vastly curtailed. Even subtle displays of bigotry, if they become public, are now costly. Trent Lott, the majority leader of the U.S. Senate, learned this in 2002 after making a toast at a one hundredth birthday party for Strom Thurmond, his fellow senator and fellow southerner. Lott made a reference in his toast to Thurmond's 1948 campaign for president, which was built on a platform of segregation; Mississippi—Lott's home state—was one of just four states that Thurmond carried. "We're proud of it," Lott told the partygoers. "And if the rest of the country had followed our lead, we wouldn't have had all these problems over all these years either." The implication that Lott was a fan of segregation raised enough of a fury that he was forced to quit his Senate leadership post.

Even if you are a private citizen, you surely wouldn't want to seem bigoted while appearing in public. Might there be a way to test for discrimination in a public setting?

Unlikely as it may seem, the television game show *The Weakest Link* provides a unique laboratory to study discrimination. An import from the United Kingdom, *The Weakest Link* for a short time became wildly popular in the United States. The game includes eight contestants (or, in a later daytime version, six) who each answer trivia questions and compete for a single cash jackpot. But the player who answers the most questions correctly isn't

necessarily the player who advances. After each round, every contestant votes to eliminate one other contestant. A player's trivia-answering ability is presumably the only worthwhile factor to consider; race, gender, and age wouldn't seem to matter. But do they? By measuring a contestant's actual votes against the votes that would truly best serve his self-interest, it's possible to tell if discrimination is at play.

The voting strategy changes as the game progresses. In the first several rounds, it makes sense to eliminate bad players since the jackpot grows only when correct answers are given. In later rounds, the strategic incentives are flipped. The value of building the jackpot is now outweighed by each contestant's desire to win the jackpot. It's easier to do that if you eliminate the other good players. So, roughly speaking, the typical contestant will vote to eliminate the worse players in the early rounds and the better players in the later rounds.

The key to measuring the *Weakest Link* voting data is to tease out a contestant's playing ability from his race, gender, and age. If a young black man answers a lot of questions correctly but is voted off early, discrimination would seem to be a factor. Meanwhile, if an elderly white woman doesn't answer a single question correctly and is still not voted off, some sort of discriminatory favoritism would seem to be at play.

Again, keep in mind that all of this is happening on camera. A contestant knows that his friends, family, and co-workers—along with a few million strangers—are watching. So who, if anyone, is discriminated against on *The Weakest Link*?

Not, as it turns out, blacks. An analysis of more than 160 episodes reveals that black contestants, in both the early and late rounds of the game, are eliminated at a rate commensurate with their trivia-answering abilities. The same is true for female contestants. In a way, neither of these findings is so surprising. Two of the most potent social campaigns of the past half-century were the civil rights movement and the feminist movement, which demonized discrimination against blacks and women, respectively.

So perhaps, you say hopefully, discrimination was practically eradicated during the twentieth century, like polio.

Or more likely, it has become so unfashionable to discriminate against certain groups that all but the most insensitive people take pains to at least *appear* fair-minded, at least in public. This hardly means that discrimination itself has ended—only that people are embarrassed to show it. How might you determine whether the lack of discrimination against blacks and women represents a true absence or just a charade? The answer can be found by looking

at other groups that society doesn't protect as well. Indeed, the *Weakest Link* voting data do indicate two kinds of contestants who *are* consistently discriminated against: the elderly and Latinos.

Among economists, there are two leading theories of discrimination. Interestingly, elderly *Weakest Link* contestants seem to suffer from one type, while Latinos suffer the other. The first type is called taste-based discrimination, which means that one person discriminates simply because he prefers to not interact with a particular type of other person. In the second type, known as information-based discrimination, one person believes that another type of person has poor skills, and acts accordingly.

On *The Weakest Link*, Latinos suffer information-based discrimination. Other contestants seem to view the Latinos as poor players, even when they are not. This perception translates into Latinos' being eliminated in the early rounds even if they are doing well and *not* being eliminated in the later rounds, when other contestants want to keep the Latinos around to weaken the field.

Elderly players, meanwhile, are victims of taste-based discrimination: in the early rounds *and* late rounds, they are eliminated far out of proportion to their skills. It seems as if the other contestants—this is a show on which the average age is thirty-four—simply don't want the older players around.

It's quite possible that a typical *Weakest Link* contestant isn't even cognizant of his discrimination toward Latinos and the elderly (or, in the case of blacks and women, his *lack* of discrimination). He is bound to be nervous, after all, and excited, playing a fast-moving game under the glare of television lights. Which naturally suggests another question: how might that same person express his preferences—and reveal information about himself—in the privacy of his home?

In a given year, some forty million Americans swap intimate truths about themselves with complete strangers. It all happens on Internet dating sites. Some of them, like Match.com, eHarmony.com, and Yahoo! Personals, appeal to a broad audience. Others cater to more specific tastes: ChristianSingles.com, JDate.com, LatinMatcher .com, BlackSinglesConnection.com, CountryWesternSingles.com, USMilitarySingles.com, OverweightDate.com, and Gay.com. Dating websites are the most successful subscription-based business on the Internet.

Each site operates a bit differently, but the gist is this: You compose a personal ad about yourself that typically includes a photo, vital statistics, your income range, level of education, likes and dislikes, and so on. If the ad catches

someone’s fancy, that someone will e-mail you and perhaps arrange a date. On many sites, you also specify your dating aims: “long-term relationship,” “a casual lover,” or “just looking.”

So there are two massive layers of data to be mined here: the information that people include in their ads and the level of response gleaned by any particular ad. Each layer of the data can be asked its own question. In the case of the ads, how forthright (and honest) are people when it comes to sharing their personal information? And in the case of the responses, what kind of information in personal ads is considered the most (and least) desirable?

Two economists and a psychologist recently banded together to address these questions. Günter J. Hitsch, Ali Hortaçsu, and Dan Ariely analyzed the data from one of the mainstream dating sites, focusing on more than 20,000 active users, half in Boston and half in San Diego. Fifty-six percent of the users were men, and the median age range for all users was twenty-one to thirty-five. Although they represented an adequate racial mix to reach some conclusions about race, they were predominantly white.

They were also a lot richer, taller, skinnier, and better-looking than average. That, at least, is what they wrote about themselves. More than 4 percent of the online daters claimed to earn more than $200,000 a year, whereas fewer than 1 percent of typical Internet users actually earn that much, suggesting that three of the four big earners were exaggerating. Male and female users typically reported that they are about an inch taller than the national average. As for weight, the men were in line with the national average, but the women typically said they weighed about twenty pounds less than the national average.

Most impressively, fully 72 percent of the women claimed “above average” looks, including 24 percent claiming “very good looks.” The online men too were gorgeous: 68 percent called themselves “above average,” including 19 percent with “very good looks.” This leaves only about 30 percent of the users with “average” looks, including a paltry 1 percent with “less than average” looks—which suggests that the typical online dater is either a fabulist, a narcissist, or simply resistant to the meaning of “average.” (Or perhaps they are all just pragmatists: as any real-estate agent knows, the typical house isn’t “charming” or “fantastic,” but unless you say it is, no one will even bother to take a look.) Twenty-eight percent of the women on the site said they were blond, a number far beyond the national average, which indicates a lot of dyeing, or lying, or both.

Some users, meanwhile, were bracingly honest. Seven percent of the men conceded that they were married, with a significant minority of these men reporting that they were “happily married.” But the fact that they were honest

doesn't mean they were rash. Of the 243 "happily married" men in the sample, only 12 chose to post a picture of themselves. The reward of gaining a mistress was evidently outweighed by the risk of having your wife discover your personal ad. ("And what were *you* doing on that website?" the husband might bluster, undoubtedly to little avail.)

Of the many ways to fail on a dating website, not posting a photo of yourself is perhaps the most certain. (Not that the photo necessarily *is* a photo of yourself; it may well be some better-looking stranger, but such deception would obviously backfire in time.) A man who does not include his photo gets only 60 percent of the volume of e-mail response of a man who does; a woman who doesn't include her photo gets only 24 percent as much. A low-income, poorly educated, unhappily employed, not very attractive, slightly overweight, and balding man who posts his photo stands a better chance of gleaning some e-mails than a man who says he makes $200,000 and is deadly handsome but doesn't post a photo. There are plenty of reasons someone might not post a photo—he's technically challenged or is ashamed of being spotted by friends or is just plain unattractive—but as in the case of a brand-new car with a For Sale sign, prospective customers will assume he's got something seriously wrong under the hood.

Getting a date is hard enough as it is. Fifty-six percent of the men who post ads don't receive even one e-mail; 21 percent of the women don't get a single response. The traits that do draw a big response, meanwhile, will not be a big surprise to anyone with even a passing knowledge of the sexes. In fact, the preferences expressed by online daters fit snugly with the most common stereotypes about men and women.

For instance, men who say they want a long-term relationship do much better than men looking for an occasional lover. But women looking for an occasional lover do great. For men, a woman's looks are of paramount importance. For women, a man's income is terribly important. The richer a man is, the more e-mails he receives. But a woman's income appeal is a bell-shaped curve: men do not want to date *low*-earning women, but once a woman starts earning too much, they seem to be scared off. Women are eager to date military men, policemen, and firemen (possibly the result of a 9/11 Effect, like the higher payments to Paul Feldman's bagel business), along with lawyers and doctors; they generally avoid men with manufacturing jobs. For men, being short is a big disadvantage (which is probably why so many lie about it), but weight doesn't much matter. For women, being overweight is deadly (which is probably why *they* lie). For a man, having red hair or curly hair is a downer, as is "bald with a fringe"—but a shaved head is okay. For a woman, salt-and-pepper hair is bad,

while blond hair is, not surprisingly, very good.

In addition to all the information about income, education, and looks, men and women on the dating site listed their race. They were also asked to indicate a preference regarding the race of their potential dates. The two preferences were "the same as mine" or "it doesn't matter." Like the *Weakest Link* contestants, the website users were now publicly declaring how they felt about people who didn't look like them. They would reveal their *actual* preferences later, in confidential e-mails to the people they wanted to date.

Roughly half of the white women on the site and 80 percent of the white men declared that race didn't matter to them. But the response data tell a different story. The white men who said that race didn't matter sent 90 percent of their e-mail queries to white women. The white women who said race didn't matter sent about 97 percent of their e-mail queries to white men. This means that an Asian man who is good-looking, rich, and well educated will receive fewer than 25 percent as many e-mails from white women as a white man with the same qualifications would receive; similarly, black and Latino men receive about half as many e-mails from white women as they would if they were white.

Is it possible that race really didn't matter for these white women and men and that they simply never happened to browse a nonwhite date that interested them? Or, more likely, did they say that race didn't matter because they wanted to come across—especially to potential mates of their own race—as open-minded?

The gulf between the information we publicly proclaim and the information we know to be true is often vast. (Or, put a more familiar way: we say one thing and do another.) This can be seen in personal relationships, in commercial transactions, and of course in politics.

By now we are fully accustomed to the false public proclamations of politicians themselves. But voters lie too. Consider an election between a black candidate and a white candidate. Might white voters lie to pollsters, claiming they will vote for the black candidate in order to appear more color-blind than they actually are? Apparently so. In New York City's 1989 mayoral race between David Dinkins (a black candidate) and Rudolph Giuliani (who is white), Dinkins won by only a few points. Although Dinkins became the city's first black mayor, his slender margin of victory came as a surprise, for preelection polls showed Dinkins winning by nearly 15 points. When the white supremacist David Duke ran for the U.S. Senate in 1990, he garnered nearly 20 percent more of the vote than pre-election polls had projected, an indication that thousands of

Louisiana voters did not want to admit their preference for a candidate with racist views.

Duke, though he never won the high political office he often sought, proved himself a master of information abuse. As Grand Wizard of the Knights of the Ku Klux Klan, he was able to compile a mailing list of thousands of rank-and-file Klansmen and other supporters who would eventually become his political base. Not content to use the list only for himself, he sold it for $150,000 to the governor of Louisiana. Years later, Duke would once again use the list himself, letting his supporters know that he'd fallen on hard times and needed their donations. In this way Duke was able to raise hundreds of thousands of dollars for his continuing work in the field of white supremacy. He had explained to his supporters in a letter that he was so broke that the bank was trying to repossess his house.

In truth, Duke had already sold his house for a solid profit. (It isn't known whether he used a real-estate agent.) And most of the money he raised from his supporters was being used not to promote any white supremacist cause but rather to satisfy Duke's gambling habit. It was a sweet little scam he was running—until he was arrested and sent to federal prison in Big Spring, Texas.