Cooperation and Coordination

FOR WHOM THE BELL CURVE TOLLS

In the 1950s the Ivy League colleges were faced with a problem. Each school wanted to produce a winning football team. The colleges found themselves overemphasizing athletics and compromising their academic standards in order to build a championship team. Yet no matter how often they practiced or how much money they spent, at the end of the season the standings were much as they had been before. The average win-loss record was still 50:50. The inescapable mathematical fact is that for every winner there had to be a loser. All the extra work canceled itself out.

The excitement of college sports depends as much on the closeness and intensity of the competition as on the level of skill. Many fans prefer college basketball and football to the professional versions; while the level of skill is lower, there is often more excitement and intensity to the competition. With this idea in mind, the colleges got smart. They joined together and agreed to limit spring training to one day. Although there were more fumbles, the games were no less exciting. Athletes had more time to concentrate on their studies. Everyone was better off, except some alumni who wanted their alma maters to excel at football and forget about academic work.

Many students would like to have a similar agreement with their fellow students before examinations. When grades are based on a traditional bell curve, one's relative standing in the class matters more than the absolute level of one's knowledge. It matters not how much you know, only that others know less than you. The way to gain an advantage over the other students is to study more. If they all do so, they all have more knowledge, but the relative standings and therefore the bottom line—the grades—are largely unchanged. If only everyone in the class could agree to limit spring studying to one (preferably rainy) day, they would get the same grades with less effort.

The feature common to these situations is that success is determined by *relative* rather than *absolute* performance. When one participant improves his own ranking, he necessarily worsens everyone else's ranking. But the fact that one's victory requires someone else's defeat does not make the game zero-sum. In a zero-sum game it is not possible to make everyone better off. Here, it is. The scope for gain comes from reducing the inputs. While there might always be the same number of winners and losers, it can be less costly for everyone to play the game.

The source of the problem of why (some) students study too much is that they do not have to pay a price or compensation to the others. Each student's studying is akin to a factory's polluting: it makes it more difficult for all the other students to breathe. Because there is no market for buying and selling studying time, the result is a rat race: each participant strives too hard, with too little to show for his efforts. But no one team or student is willing to be the only one, or the leader, in reducing the effort. This is just like a prisoners' dilemma with more than two prisoners. An escape from the horns of this dilemma requires an enforceable collective agreement.

As with the Ivy League or OPEC, the trick is to form a cartel to limit competition. The problem for high-school students is that the cartel cannot easily detect cheating. For the collectivity of students, a cheater is one who studies more to sneak an advantage over the others. It is hard to tell if some are secretly studying until after they have aced the test. By then it is too late.

In some small towns, high-school students do have a way to enforce "no-studying" cartels. Everyone gets together and cruises Main Street at night. The absence of those home studying is noticed. Punishment can be social ostracism or worse.

To arrange a self-enforcing cartel is difficult. It is all the better if an outsider enforces the collective agreement limiting competition. This is just what happened for cigarette advertising, although not intentionally. In the old days, cigarette companies used to spend money to convince consumers to "walk a mile" for their product or to "fight rather than switch." The different campaigns made advertising agencies rich, but their main purpose was defensive—each company advertised because the others did, too. Then, in 1968, cigarette advertisements were banned from TV by law. The companies thought this restriction would hurt them and fought against it. But, when the smoke cleared, they saw that the ban helped them all avoid costly advertising campaigns and thus improved all their profits.

THE ROUTE LESS TRAVELED

There are two main ways to commute from Berkeley to San Francisco. One is driving over the Bay Bridge, and the other is taking public transportation, the Bay Area Rapid Transit train (BART). Crossing the bridge is the shortest route, and with no traffic, a car can make the trip in 20 minutes. But that is rarely the case. The bridge has only four lanes and is easily congested.* We suppose that each additional 2,000 cars (per hour) causes a 10-minute delay for everyone on the road. For example, with 2,000 cars the travel time rises to 30 minutes; at 4,000 cars, to 40 minutes.

The BART train makes a number of stops, and one has to walk to the station and wait for the train. It is fair to say that the trip takes closer to 40 minutes along this route, but the train never fights traffic. When train usage rises, they put on more cars, and the commuting time stays roughly constant.

If, during rush hour, 10,000 commuters want to go from Berkeley to San Francisco, how will the commuters be distributed over the two routes? Each commuter will act selfishly, choosing the route that minimizes his own transportation time. Left to their own devices, 40 percent will drive and 60 percent will take the train. The commuting time will be 40 minutes for everyone. This outcome is the equilibrium of a game.

We can see this result by asking what would happen if the split were different. Suppose only 2,000 drivers took the Bay Bridge. With less congestion, the trip would take less time (30 minutes) along this route. Then some of the 8,000 BART commuters would find out that they could save time by switching, and would do so. Conversely, if there were, say, 8,000 drivers using the Bay Bridge, each spending 60 minutes, some of them would switch to the train for the faster trip it provides. But when there are 4,000 drivers on the Bay Bridge and 6,000 on the train, no one can gain by switching: the commuters have reached an equilibrium.

We can show the equilibrium using a simple chart, which is quite similar in spirit to the one in chapter 4 describing the classroom experiment of the prisoners' dilemma. In this chart, we are holding the total number of commuters constant at 10,000, so that when there are 2,000 cars using the bridge, it implies that 8,000 commuters are using BART. The rising line shows how the trip time on the Bay Bridge increases as the number of drivers on it increases. The flat line shows the constant time of 40 minutes for the train. The lines intersect at *E*, showing that the trip times on the two routes are equal when the number of drivers on the Bay Bridge is 4,000. This graphic depiction is a useful tool to describe the equilibrium, and we will use it often in this chapter.

Is this equilibrium good for the commuters as a whole? Not really. It is easy to find a better pattern. Suppose only 2,000 take the Bay Bridge. Each of them saves 10 minutes. The 2,000 who switch to the train are still spending the same time as they did before, namely 40 minutes. So are the 6,000 who were already taking the train. We have just saved 20,000 person-minutes (or almost two weeks) from the total travel time.

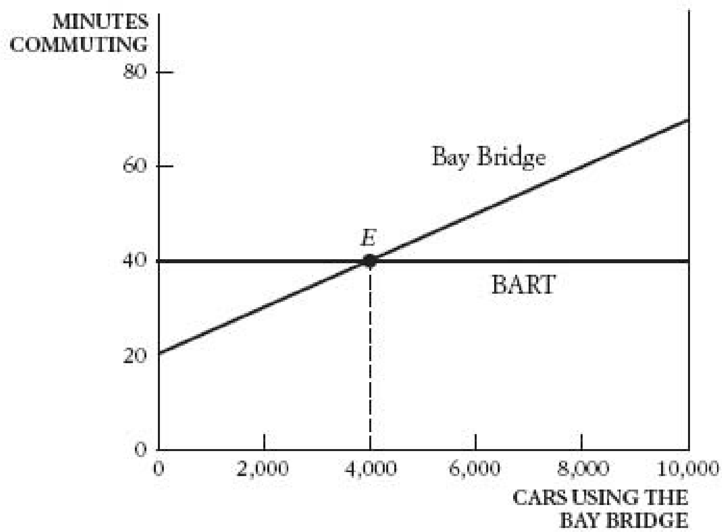

Why is this saving possible? Or, in other words, why were the drivers left to themselves not guided by an invisible hand to the best mix of routes? The answer again lies in the cost that each user of the Bay Bridge inflicts on the others. When an extra driver takes this road, the travel time of all the other users goes up by a little bit. But the newcomer is not required to pay a price that reflects this cost. He takes into account only his own travel time.

What traffic pattern is best for the group of drivers as a whole? In fact, the one we constructed, with 2,000 cars on the Bay Bridge and a total time saving of 20,000 minutes, is best. To see this, try a couple of others. If there are 3,000 cars on the Bay Bridge, the travel time is 35 minutes, with a saving of 5 minutes each, or 15,000 minutes in all. With only 1,000 cars, the travel time is 25 minutes, and each saves 15 minutes, but the total saving is again only 15,000 minutes. The intermediate point with 2,000 drivers, each saving 10 minutes, is best.

How can the best pattern be achieved? Devotees of central planning will think of issuing 2,000 licenses to use the Bay Bridge. If they are worried about the inequity of allowing those with licenses to travel in 30 minutes while the other 8,000 must take the train and spend 40 minutes, they will devise an ingenious system of rotating the licenses among the population every month.

A market-based solution charges people for the harm they cause to others. Suppose each person values an hour of time at $12, that is, each would be willing to pay $12 to save an hour. Then charge a toll for driving on the Bay Bridge; set the toll $2 above the BART fare. By our supposition, people regard an extra $2 cost as equivalent to 10 minutes of time. Now the equilibrium commuting pattern will have 2,000 cars on the Bay Bridge and 8,000 riders on BART. Each user of the Bay Bridge spends 30 minutes plus an extra $2 in commuting costs; each BART rider spends 40 minutes. The total effective costs are the same, and no one wants to switch to the other route. In the process we have collected $4,000 of toll revenue (plus an additional 2,000 BART fares), which can then go into the county's budget, thus benefiting everyone because taxes can be lower than they would otherwise be.

A solution even closer to the spirit of free enterprise would be to allow private ownership of the Bay Bridge. The owner realizes that people are willing to pay for the advantage of a faster trip on a less congested road. He charges a price, therefore, for the privilege. How can he maximize his revenue? By maximizing the total value of the time saved, of course.

The invisible hand guides people to an optimal commuting pattern only when the good "commuting time" is priced. With the profit-maximizing toll on the bridge, time really is money. Those commuters who ride BART are selling time to those who use the bridge.

Finally, we recognize that the cost of collecting the toll sometimes exceeds the resulting benefit of saving people's time. Creating a marketplace is not a free lunch. The toll booths may be a primary cause of the congestion. If so, it may be best to tolerate the original inefficient route choices.

CATCH-22?

Chapter 4 offered the first examples of games with many equilibria. Where should two strangers meet in New York City: Times Square or the Empire State Building? Who should return a disconnected phone call? In those examples it was not important which of the conventions was chosen, so long as everyone agreed on the same convention. But sometimes one convention is much better than another. Even so, that doesn't mean it will always get adopted. If one convention has become established and then some change in circumstances makes another one more desirable, it can be especially hard to bring about the change.

The keyboard design on most typewriters is a case in point. In the late 1800s, there was no standard pattern for the arrangement of letters on the typewriter keyboard. Then in 1873 Christopher Scholes helped design a "new, improved" layout. The layout became known as QWERTY, after the letter arrangement of the first six letters in the top row. QWERTY was chosen to *maximize* the distance between the most frequently used letters. This was a good solution in its day; it deliberately slowed down the typist, and reduced the jamming of keys on manual typewriters. By 1904, the Remington Sewing Machine Company of New York was mass-producing typewriters with this layout, and it became the de facto industry standard. But with the advent of electric typewriters and, later, computers, this jamming problem became irrelevant. Engineers developed new keyboard layouts, such as DSK (Dvorak's Simplified Keyboard), which reduced the distance typists' fingers traveled by over 50 percent. The same material can be typed in 5 to 10 percent less time using DSK than QWERTY.[1] But QWERTY is the established system. Almost all keyboards use it, so we all learn it and are reluctant to learn a second layout. Keyboard manufacturers continue, therefore, with QWERTY. The vicious circle is complete.[2]

If history had worked differently, and if the DSK standard had been adopted from the outset, that would have been better for today's technology. However, given where we are, the question of whether or not we should switch standards involves further considerations. There is a lot of inertia, in the form of machines, keyboards, and trained typists, behind QWERTY. Is it worthwhile to retool?

From the point of view of society as a whole, the answer would seem to be yes. During the Second World War, the U.S. Navy used DSK typewriters on a large scale, and retrained typists to use them. It found that the cost of retraining could be fully recouped in only ten days of use.

However, this study and the overall advantage of DSK has been called into question by Professors Stan Liebowitz and Stephen Margolis.[3] It appears that an interested party, one Lieutenant Commander August Dvorak, was involved in conducting the original study. A 1956 General Services Administration study found that it took a month of four-hour-a-day training for typists to catch up to their old QWERTY speed. At that point, further training on the Dvorak keyboard was less effective than providing training to QWERTY typists. To the extent that DSK is superior, the biggest gain is when typists learn this system from the start.

If the typist becomes so good that he or she almost never has to look at the keyboard, then learning DSK makes sense. With today's software, it is a relatively simple matter to reassign the keys from one layout to another. (On a Mac, it is a simple switch on the keyboard menu.) Thus the keyboard layout almost doesn't matter. Almost. The problem is: how does one learn to touch type on a mislabeled keyboard? Anyone who wants to reassign the layout from QWERTY to DSK but cannot yet touch type must look at the keyboard and mentally convert each key to its DSK value. This is not practical. Therefore beginners have to learn QWERTY anyway and that greatly reduces the gains from also learning DSK.

No individual user can change the social convention. The uncoordinated decisions of individuals keep us tied to QWERTY. The problem is called a bandwagon effect and can be illustrated using the following chart. On the horizontal axis we show the fraction of typists using QWERTY. The vertical axis details the chance that a new typist will learn QWERTY as opposed to DSK. As drawn, if 85 percent of typists are using QWERTY, then the chances are 95 percent that a new typist will choose to learn QWERTY and only 5 percent that the new typist will learn DSK. The way the curve is drawn is meant to emphasize the superiority of the DSK layout. A majority of new typists will learn DSK rather than QWERTY provided that QWERTY has anything less than a 70 percent market share. In spite of this handicap, it is possible for QWERTY to dominate in equilibrium. (Indeed, this possibility is just what has happened in the prevailing equilibrium.)

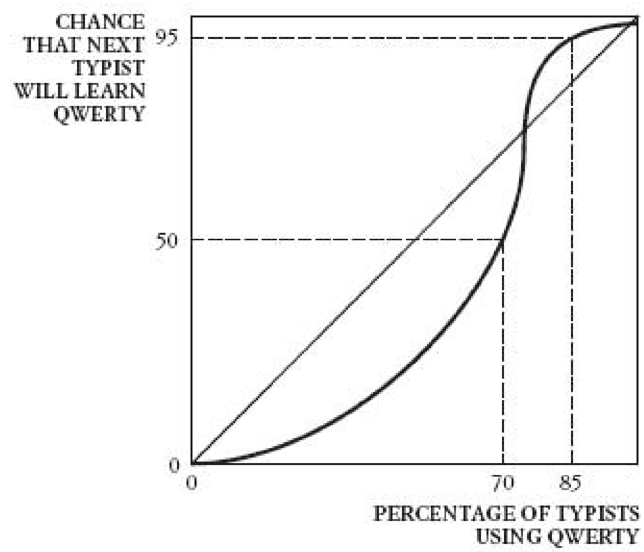

The choice of which keyboard to use is a strategy. When the fraction using each technology is constant over time, we are at an equilibrium of the game. Showing that this game converges to an equilibrium is not easy. The random choice of each new typist is constantly disrupting the system. Recent high-powered mathematical tools in the field of stochastic approximation theory have allowed economists and statisticians to prove that this dynamic game does converge to an equilibrium.[4] We now describe the possible outcomes.

If the fraction of typists using QWERTY exceeds 72 percent, there is the expectation that an even greater fraction of people will learn QWERTY. The prevalence of QWERTY expands until it reaches 98 percent. At that point, the fraction of new typists learning QWERTY just equals its predominance in the population, 98 percent, and so there is no more upward pressure.[*]

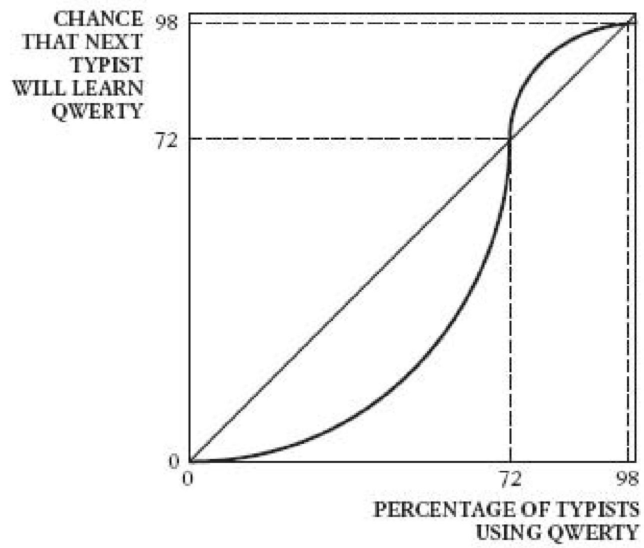

Conversely, if the fraction of typists using QWERTY falls below 72 percent, then there is the expectation that DSK will take over. Fewer than 72 percent of the new typists learn QWERTY, and the subsequent fall in its usage gives new typists an even greater incentive to learn the superior layout of DSK. Once all typists are using DSK there is no reason for a new typist to learn QWERTY, and QWERTY will die out.

The mathematics says only that we will end up at one of these two possible outcomes: everyone using DSK or 98 percent using QWERTY. It does not say which will occur. If we were starting from scratch, the odds are in favor of DSK being the predominant keyboard layout. But we are not. History matters. The historical accident that led to QWERTY capturing nearly 100 percent of typists ends up being self-perpetuating, even though the original motivation for QWERTY is long since obsolete.

Since bad luck or the convergence to an inferior equilibrium is self-perpetuating, there is the possibility of making everyone better off. But it requires coordinated action. If the major computer manufacturers coordinate on a new keyboard layout or a major employer, such as the federal government, trains its employees on a new keyboard, this could switch the equilibrium all the way from one extreme to the other. The essential point is that it is not necessary to convert everyone, just a critical mass. Given enough of a toehold, the better technology can take it from there.

The QWERTY problem is but one minor example of a more widespread problem. Our preference for gasoline engines over steam and light-water nuclear reactors over gas-cooled is better explained by historical accidents than by the superiority of the adopted technologies. Brian Arthur, an economist at Stanford and one of the developers of the mathematical tools used to study bandwagon effects, tells the story of how we ended up with gasoline-powered cars.

> In 1890 there were three ways to power automobiles—steam, gasoline, and electricity—and of these one was patently *inferior* to the other two: gasoline…. [A turning point for gasoline was]an 1895 horseless carriage competition sponsored by the *Chicago Times-Herald*. This was won by a gasoline-powered Duryea—one of only two cars to finish out of six starters—and has been cited as the possible inspiration for R. E. Olds to patent in 1896 a gasoline power source, which he subsequently mass-produced in the "Curved-Dash Olds." Gasoline thus overcame its slow start. Steam continued to be viable as an automotive power source until 1914, when there was an outbreak of hoof-and-mouth disease in North America. This led to the withdrawal of horse troughs—which is where steam cars could fill with water. It took the Stanley brothers about three years to develop a condenser and boiler system that did not need to be filled every thirty or forty miles. But by then it was too late. The steam engine never recovered.[5]

While there is little doubt that today's gasoline technology is better than steam, that's not the right comparison. How good would steam have been if it had had the benefit of seventy-five years of research and development? While we may never know, some engineers believe that steam was the better bet.[6]

In the United States, almost all nuclear power is generated by light-water reactors. Yet there are reasons to believe that the alternative technologies of heavy-water or gas-cooled reactors would have been superior, especially given the same amount of learning and experience. Canada's experience with heavy-water reactors allows them to generate power for 25 percent less cost than light-water reactors of equivalent size in the United States. Heavy-water reactors can operate without the need to reprocess fuel. Perhaps most important is the safety comparison. Both heavy-water and gas-cooled reactors have a significantly lower risk of a meltdown—the former because the high pressure is distributed over many tubes rather than a single core vessel, and the latter because of the much slower temperature rise in the event of coolant loss.[7]

The question of how light-water reactors came to dominate has been studied by Robin Cowen in a 1987 Stanford University Ph.D. thesis. The first consumer for nuclear power was the U.S. Navy. In 1949, then Captain Rickover made the pragmatic choice in favor of light-water reactors. He had two good reasons. It was the most compact technology at the time, an important consideration for submarines, and it was the furthest advanced, suggesting that it would have the quickest route to implementation. In 1954, the first nuclear-powered submarine, *Nautilus*, was launched. The results looked positive.

At the same time civilian nuclear power became a high priority. The Soviets had exploded their first nuclear bomb in 1949. In response, Atomic Energy Commissioner T. Murray warned, "Once we become fully conscious of the possibility that [energy-poor] nations will gravitate towards the USSR if it wins the nuclear power race, it will be quite clear that this race is no Everest-climbing, kudos-providing contest."[8] General Electric and Westinghouse, with their experience producing light-water reactors for the nuclear-powered submarines, were the natural choice to develop civilian power stations. Considerations of proven reliability and speed of implementation took precedence over finding the most cost-effective and safe technology. Although light-water was first chosen as an interim technology, this gave it enough of a head start down the learning curve that the other options have never had the chance to catch up.

The adoption of QWERTY, gasoline engines, and light-water reactors are but three demonstrations of how history matters in determining today's technology choices, though the historical reasons may be irrelevant considerations in the present. Typewriter-key jamming, hoof-and-mouth disease, and submarine space constraints are not relevant to today's trade-offs between the competing technologies. The important insight from game theory is to recognize early on the potential for future lock-in—once one option has enough of a head start, superior technological alternatives may never get the chance to develop. Thus there is a potentially great payoff in the early stages from spending more time figuring out not only what technology meets today's constraints but also what options will be the best for the future.

FASTER THAN A SPEEDING TICKET

Just how fast should you drive? In particular, should you abide by the speed limit? Again the answer is found by looking at the game where your decision interacts with those of all the other drivers.

If nobody is abiding by the law, then you have two reasons to break it too. First, some experts argue that it is actually safer to drive at the same speed as the flow of traffic.[9] On most highways, anyone who tries to drive at fifty-five miles per hour creates a dangerous obstacle that everyone else must go around.

Second, when you tag along with the other speeders, your chances of getting caught are almost zero. The police simply cannot pull over more than a small percentage of the speeding cars. As long as you go with the flow of traffic, there is safety in numbers.*

As more people become law-abiding, both reasons to speed vanish. It becomes more dangerous to speed, since this requires weaving in and out of traffic. And your chances of getting caught increase dramatically.

We show this in a chart similar to the one for commuters from Berkeley to San Francisco. The horizontal axis measures the percentage of drivers who abide by the speed limit. The lines *A* and *B* show each driver's calculation of his benefit from (*A*) abiding by and (*B*) breaking the law. Our argument says that if no one else is keeping under the limit (the far left end), neither should you (line *B* is higher than line *A*); if everyone else is law-abiding (the far right end), you should be too (line *A* is higher than line *B*). Once again there are three equilibria, of which only the extreme ones can arise from the process of social dynamics as drivers adjust to one another's behavior.

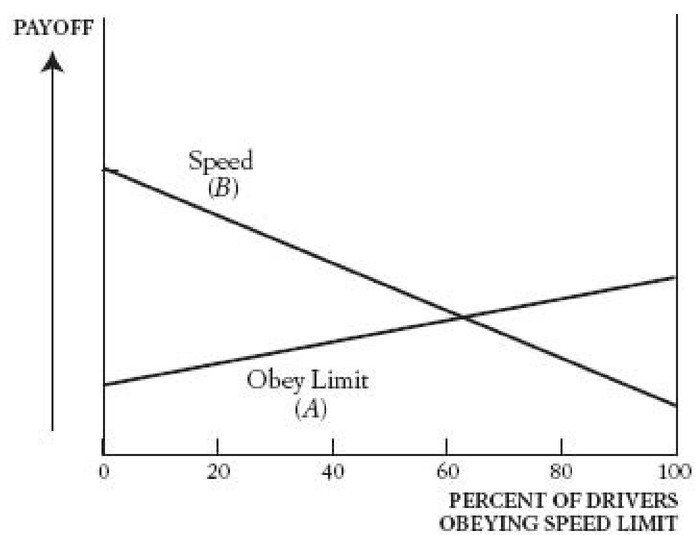

In the case of the commuters choosing between the Bay Bridge and BART, the dynamics converged on the equilibrium in the middle. Here the tendency is toward one of the extremes. The difference arises because of the way interactions work. With commuting, either choice becomes *less* attractive when more of the others follow you, whereas with speeding, additional company makes it *more* attractive.

The general theme of one person's decision affecting the others applies here, too. If one driver speeds up, he makes it a little safer for the others to speed. If no one is speeding, no one is willing to be the first to do so and provide this “benefit” to the others without being “rewarded” for doing so. But there is a new twist: if everyone is speeding, then no one wants to be the only one to slow down.

Can this situation be affected by changing the speed limit? The chart is drawn for a specific speed limit, say 55 m.p.h. Suppose the limit is raised to 65. The value of breaking the limit falls, since beyond a point, higher speeds do become dangerous, and the extra advantage of going 75 instead of 65 is less than the gain of going 65 over 55. Furthermore, above 55 miles an hour, gasoline consumption goes up exponentially with speed. It may be 20 percent more expensive to drive at 65 than at 55, but it could easily be 40 percent more expensive to drive at 75 rather than at 65.

What can lawmakers learn from this if they want to encourage people to drive at the speed limit? It is not necessary to set the speed limit so high that everyone is happy to obey it. The key is to get a critical mass of drivers obeying the speed limit. Thus a short phase of extremely strict enforcement and harsh penalties can change the behavior of enough drivers to generate the momentum toward full compliance. The equilibrium moves from one extreme (where everyone speeds) to the other (where everyone complies). With the new equilibrium, the police can cut back on enforcement, and the compliance behavior is self-sustaining. More generally, what this suggests is that short but intense enforcement can be significantly more effective than the same total effort applied at a more moderate level for a longer time.[10]

A similar logic applies to fuel-economy standards. For many years, the vast majority of Americans supported a large increase in the Corporate Average Fuel Economy (CAFE) standards. Finally, in 2007 President Bush signed legislation mandating an increase from 27.5 mpg to 35 mpg for cars (and similar increases for trucks) to be phased in gradually starting in 2011 before taking full effect in 2020. But if most people want higher fuel economy, nothing prevents them from buying a fuel-efficient car. Why is it that folks who want higher fuel standards keep on driving gas-guzzling SUVs?

One reason is that people are concerned that fuel-efficient cars are lighter and thus less safe in the event of an accident. Light cars are especially unsafe when hit by a Hummer. Folks are more willing to drive a light car when they know that the other cars on the road are light as well. As speeding leads to speeding, the more heavy cars there are out there, the more everyone needs to drive an SUV in order to be safe. Just like people, cars have become 20 percent heavier over the last two decades. The end result is that we end up with low fuel economy and no one is safer. A move to higher CAFE standards is the coordination device that could help shift enough people from heavy to light cars so that (almost) everyone would be happier driving a light car.[11] Perhaps even more important than a technological advance is the coordination change that would shift the mix of cars and thereby allow us to improve fuel economy right away.

Arguments in favor of collective, rather than individual, decisions are not the preserve of liberals, left-wingers, and any remaining socialists. The impeccably conservative economist Milton Friedman made the same logical argument about redistribution of wealth in his classic *Capitalism and Freedom*:

> I am distressed by the sight of poverty; I am benefited by its alleviation; but I am benefited equally whether I or someone else pays for its alleviation; the benefits of other people's charity therefore partly accrue to me. To put it differently, we might all of us be willing to contribute to the relief of poverty, *provided* everyone else did. We might not be willing to contribute the same amount without such assurance. In small communities, public pressure can suffice to realize the proviso even with private charity. In the large impersonal communities that are increasingly coming to dominate our society, it is much more difficult for it to do so. Suppose one accepts, as I do, this line of reasoning as justifying governmental action to relieve poverty…[12]

WHY DID THEY LEAVE?

American cities have few racially integrated neighborhoods. If the proportion of black residents in an area rises above a critical level, it quickly increases further to nearly 100 percent. If it falls below a critical level, the expected course is for the neighborhood to become all white. Preservation of racial balance requires some ingenious public policies.

Is the de facto segregation of most neighborhoods the product of widespread racism? These days, a large majority of urban Americans would regard mixed neighborhoods as desirable.* The more likely difficulty is that segregation can result as the equilibrium of a game in which each household chooses where to live, even when they all have a measure of racial tolerance. This idea is due to

Thomas Schelling.[13] We shall now outline it, and show how it explains the success of the Chicago suburb Oak Park in maintaining an integrated community.

Racial tolerance is not a matter of black or white; there are shades of gray. Different people, black or white, have different views about the best racial mix. For example, very few whites insist on a neighborhood that is 99 or even 95 percent white; yet most will feel out of place in one that is only 1 or 5 percent white. The majority would be happy with a mix somewhere in between.

We can illustrate the evolution of neighborhood dynamics using a chart similar to the one from the QWERTY story. On the vertical axis is the probability that a new person moving into the neighborhood will be white. This is plotted in relationship to the current racial mix, shown on the horizontal axis. The far right end of the curve shows that once a neighborhood becomes completely segregated (all white), the odds are overwhelming that the next person who moves into the neighborhood will also be white. If the current mix falls to 95 percent or 90 percent white, the odds are still very high that the next person to move in will also be white. If the mix changes much further, then there is a sharp drop-off in the probability that the next person to join the community will be white. Finally, as the actual percentage of whites drops to zero, meaning that the neighborhood is now segregated at the other extreme, the probability is very high that the next person to move in will be black.

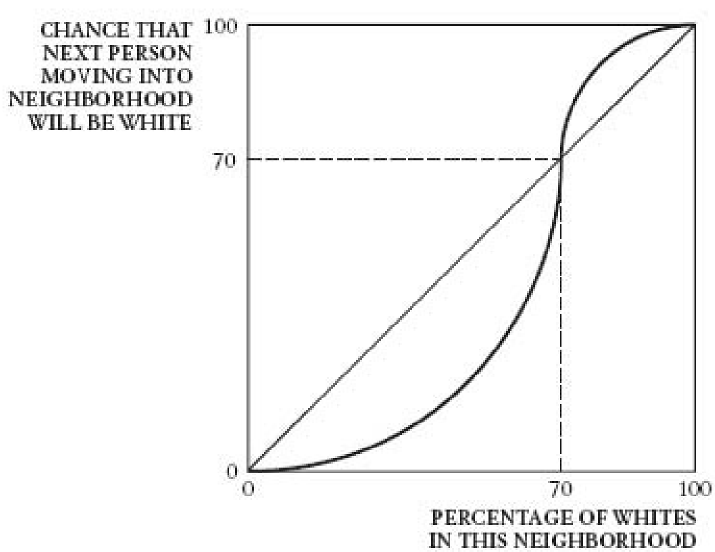

In this situation, the equilibrium will be where the racial mix of the population exactly equals the mix of new entrants to the community. Only in this event are the dynamics stable. There are three such equilibria: two at the extremes where the neighborhood is all white and all black, and one in the middle where there is a mix. The theory so far does not tell us which of the three equilibria is the most likely. In order to answer this question, we need to examine the forces that move the system toward or away from an equilibrium—that is, the social dynamics of the situation.

Social dynamics will always drive the neighborhood to one of the extreme equilibria. Schelling labeled this phenomenon “tipping” (an idea later popularized by Malcolm Gladwell’s book *The Tipping Point*). Let us see why it occurs. Suppose the middle equilibrium has 70 percent whites and 30 percent blacks. By chance, let one black family move out and be replaced by a white family. Then the proportion of whites in this neighborhood becomes slightly above 70 percent. Looking at the chart, the probability that the next entrant will also be white is then above 70 percent. The upward pressure is reinforced by the new entrants. Say the racial mix shifts to 75:25 percent. The tipping pressure continues. The chance that a new entrant will be white is above 75 percent, so the expectation is that the neighborhood will become increasingly segregated. This goes on until the mix of new entrants is the same as the mix in the neighborhood. As drawn, that occurs again only when the neighborhood is all white. If the process had started with one white family moving out and one black family moving in, there would have been a chain reaction in the opposite direction, and the odds are that the neighborhood would have become all black.

The problem is that the 70:30 percent mix is not a stable equilibrium. If this mix is somehow disrupted, as it will surely be, there is a tendency to move toward one of the extremes. Sadly, from the extremes there is no similar tendency to move back toward the middle. Although segregation is the predicted equilibrium, that does not mean that people are better off at this outcome. Everyone might prefer to live in a mixed neighborhood. But they rarely exist and, even when found, tend not to last.

Once again, the source of the problem is the effect of one household’s action on the others. Starting at a 70:30 percent mix, when one white family replaces a black family, this may make the neighborhood a little less attractive for future blacks to move in. But it is not assessed a fine for this. Perhaps there should be a neighborhood departure tax analogous to road tolls. But that would be counter to a more basic principle, namely the freedom to live where one chooses. If society wants to prevent tipping, it must look for some other policy measures.

If we cannot fine a departing family for the damage it causes, both to those who remain and those who now might choose not to come, we must take measures that will reduce the incentives for others to follow suit. If one white family leaves, the neighborhood should not become less attractive to another white family. If one black family leaves, the neighborhood should not become less attractive to another black family. Public policy can help prevent the tipping process from gathering momentum.

The racially integrated Chicago suburb of Oak Park provides an ingenious example of policies that work. It uses two tools: first, the town banned the use of "For Sale" signs in front yards, and, secondly, the town offers insurance that guarantees homeowners that they will not lose the value of their house and property because of a change in the racial mix.

If by chance two houses on the same street are for sale at the same time, "For Sale" signs would spread this news quickly to all neighbors and prospective purchasers. Eliminating such signs makes it possible to conceal the news that would be interpreted as bad; nobody need know until after a house has been sold that it was even up for sale. The result is that panics are avoided (unless they are justified, in which case they are just delayed).

By itself, the first policy is not enough. Homeowners might still worry that they should sell their house while the going is good. If you wait until the neighborhood has tipped, you've waited too long and may find that you've lost most of the value of your home, which is a large part of most people's wealth. Once the town provides insurance, this is no longer an issue. In other words, the insurance removes the *economic* fear that accelerates tipping. In fact, if the guarantee succeeds in preventing tipping, property values will not fall and the policy will not cost the taxpayers anything.

Tipping to an all-black equilibrium has been the more common problem in urban America. But in recent years gentrification, which is just tipping to an all-rich equilibrium, has been on the rise. Left unattended, the free market will often head to these unsatisfactory outcomes. But public policy, combined with an awareness of how tipping works, can help stop the momentum toward tipping and preserve the delicate balances.

IT CAN BE LONELY AT THE TOP

Top law firms generally choose their partners from among their junior associates. Those not chosen must leave the firm, and generally move to a lower-ranked one. At the mythical firm Justin-Case, the standards were so high that for many years no new partners were selected. The junior associates protested about this lack of advancement. The partners responded with a new system that looked very democratic.

Here is what they did. At the time of the annual partnership decision, the abilities of the ten junior associates were rated from 1 to 10, with 10 being the best. The junior associates were told their rating privately. Then they were ushered into a meeting room where they were to decide by majority vote the cutoff level for partnership.

They all agreed that everyone making partner was a good idea and certainly preferable to the old days when nobody made partner. So they began with a cutoff of 1. Then some high-rated junior associate suggested that they raise the cutoff to 2. He argued that this would improve the average quality of the partnership. Eight junior associates agreed with him. The sole dissenting vote came from the least able member, who would no longer make partner.

Next, someone proposed that they raise the standard from 2 to 3. Eight people were still above this standard, and they all voted for this improvement in the quality of the partnership. The person ranked 2 voted against, as this move deprived him of partnership. What was surprising was that the lowest-rated junior associate was in favor of this raising of the standards. In neither case would he make partner. But at least in the latter he would be grouped with someone who had ability 2. Therefore, upon seeing that he was not selected, other law firms would not be able to infer his exact ability. They would guess that he was either a 1 or a 2, a level of uncertainty that would be to his advantage. The proposal to raise the standard to 3 passed 9:1.

With each new cutoff level someone proposed raising it by one. All those strictly above voted in favor so as to raise the quality of the partnership (without sacrificing their own position), while all those strictly below joined in support of raising the standard so as to make their failure less consequential. Each time there was only one dissenter, the associate right at the cutoff level who would no longer make partner. But he was outvoted 9:1.

And so it went, until the standard was raised all the way up to 10. Finally, someone proposed that they raise the standard to 11 so that *nobody* would make partner. Everybody rated 9 and below thought that this was a fine proposal, since once more this improved the average quality of those rejected. Outsiders would not take it as a bad sign that they didn't make partner, as nobody made partner at this law firm. The sole voice against was the most able junior associate, who lost his chance to make partner. But he was outvoted 9:1.

The series of votes brought everybody back to the old system, which they all considered worse than the alternative of promotion for all. Even so, each resolution along the way passed 9:1. There are two morals to this story.

When actions are taken in a piecemeal way, each step of the way can appear attractive to the vast majority of decision makers. But the end is worse than the beginning for everyone. The reason is that voting ignores the intensity of preferences. In our example, all those in favor gain a small amount, while the one person against loses a lot. In the series of ten votes, each junior associate has nine small victories and one major loss that outweighs all the combined gains. Similar problems tend to plague bills involving reforms of taxes or trade tariffs; they get killed by a series of amendments. Each step gets a majority approval, but the end result has enough fatal flaws so that it loses the support of a majority.

Just because an individual recognizes the problem does not mean an individual can stop the process. It is a slippery slope, too dangerous to get onto. The group as a whole must look ahead and reason back in a coordinated way, and set up the rules so as to prevent taking the first steps on the slope. There is safety when individuals agree to consider reforms only as a package rather than as a series of small steps. With a package deal, everyone knows where he will end up. A series of small steps can look attractive at first, but one unfavorable move can more than wipe out the entire series of gains.

In 1989, Congress learned this danger first-hand in its failed attempt to vote itself a 50 percent pay raise. Initially, the pay raise seemed to have wide support in both houses. When the public realized what was about to happen, they protested loudly to their representatives. Consequently, each member of Congress had a private incentive to vote against the pay hike, provided he or she thought that the hike would still pass. The best scenario would be to get the higher salary while having protested against it. Unfortunately (for them), too many members of Congress took this approach, and suddenly passage no longer seemed certain. As each defection moved them further down the slippery slope, there was all the more reason to vote against it. If the pay hike were to fail, the worst possible position would be to go on record supporting the salary hike, pay the political price, and yet not get the raise. At first, there was the potential for a few individuals to selfishly improve their own position. But each defection increased the incentive to follow suit, and soon enough the proposal was dead.

There is a second, quite different moral to the Justin-Case story. If you are going to fail, you might as well fail at a difficult task. Failure causes others to downgrade their expectations of you in the future. The seriousness of this problem depends on what you attempt. Failure to climb Mt. Everest is considerably less damning than failure to finish a 10K race. The point is that when other people’s perception of your ability matters, it might be better for you to do things that *increase* your chance of failing in order to reduce its consequence. People who apply to Harvard instead of the local college or ask the most popular student to prom instead of a more realistic prospect are following such strategies.

Psychologists see this behavior in other contexts. Some individuals are afraid to recognize the limits of their own ability. In these cases they take actions that increase the chance of failure in order to avoid facing their ability. For example, a marginal student may not study for a test so that if he fails, the failure can be blamed on his lack of studying rather than intrinsic ability. Although perverse and counterproductive, there is no invisible hand to protect you in games against yourself.

POLITICIANS AND APPLE CIDER

Two political parties are trying to choose their positions on the liberal-conservative ideological spectrum. First the incumbent takes a stand, then the challenger responds.

Suppose the voters range uniformly over the spectrum. For concreteness, number the political positions from 0 to 100, where 0 represents radical left and 100 represents arch-conservative. If the incumbent chooses a position such as 48, slightly more liberal than the middle of the road, the challenger will take a position between that and the middle—say, 49. Then voters with preferences of 48 and under will vote for the incumbent; all others, making up just over 51 percent of the population, will vote for the challenger. The challenger will win.

If the incumbent takes a position above 50, then the challenger will locate between that and 50. Again this will get him more than half the votes.

By the principle of looking ahead and reasoning backward, the incumbent can figure out that his best bet is to locate right in the middle. (As with highways, the position in the middle of the road is called the median.) When voters' preferences are not necessarily uniform, the incumbent locates at the position where 50 percent of the voters are located to the left and 50 percent are to the right. This median is not necessarily the average position. The median position is determined by where there are an equal number of voices on each side, while the average gives weight to how far the voices are away. At this location, the forces pulling for more conservative or more liberal positions have equal numbers. The best the challenger can do is imitate the incumbent. The two parties take identical stands, so each gets 50 percent of the votes if issues are the only thing that counts. The losers in this process are the voters, who get an echo rather than a choice.

In practice, parties do not take identical hard positions, but each fudges its stand around the middle ground. This phenomenon was first recognized by Columbia University economist Harold Hotelling in 1929. He pointed out similar examples in economic and social affairs: “Our cities become uneconomically large and the business districts within them are too concentrated. Methodist and Presbyterian churches are too much alike; cider is too homogeneous.”[14]

Would the excess homogeneity persist if there were three parties? Suppose they take turns to choose and revise their positions, and have no ideological baggage to tie them down. A party located on the outside will edge closer to its neighbor to chip away some of its support. This will squeeze the party in the middle to such an extent that when its turn comes, it will want to jump to the outside and acquire a whole new and larger base of voters. This process will then continue, and there will be *no* equilibrium. In practice, parties have enough ideological baggage, and voters have enough party loyalty, to prevent such rapid switches.

In other cases, locations won’t be fixed. Consider three people all looking for a taxi in Manhattan. Though they start waiting at the same time, the one at the most uptown position will catch the first taxi going downtown, and the one located farthest downtown will catch the first uptown cab. The one in the middle is squeezed out. If the middle person isn’t willing to be usurped, he will move either uptown or downtown to preempt one of the other two. Until the taxi arrives, there may not be an equilibrium; no individual is content to remain squeezed in the middle. Here we have yet another, and quite different, failure of an uncoordinated decision process; it may not have a determinate outcome at all. In such a situation, society has to find a different and coordinated way of reaching a stable outcome.

A RECAPITULATION

In this chapter we described many instances in which the games people play have more losers than winners. Uncoordinated choices interact to produce a poor outcome for society. Let us summarize the problems briefly, and you can then try out the ideas on the case study.

First we looked at games in which each person had an either-or choice. One problem was the familiar multiperson prisoners’ dilemma: everyone made the same choice, and it was the wrong one. Next we saw examples in which some people made one choice while their colleagues made another, but the proportions were not optimal from the standpoint of the group as a whole. This happened because one of the choices involved spillovers—that is, effects on others—which the choosers failed to take into account. Then we had situations in which either extreme—everyone choosing one thing or everyone choosing the other—was an equilibrium. To choose one, or make sure the right one was chosen, required social conventions, penalties, or restraints on people's behavior. Even then, powerful historical forces might keep the group locked into the wrong equilibrium.

Turning to situations with several alternatives, we saw how the group could voluntarily slide down a slippery path to an outcome it would collectively regret. In other examples, we found a tendency toward excessive homogeneity. Sometimes there might be an equilibrium held together by people's mutually reinforcing expectations about what others think. In still other cases, equilibrium might fail to exist altogether, and another way to reach a stable outcome would have to be found.

The point of these stories is that the free market doesn't always get it right. There are two fundamental problems. One is that history matters. Our greater experience with gasoline engines, QWERTY keyboards, and light-water nuclear reactors may lock us in to continued use of these inferior technologies. Accidents of history cannot necessarily be corrected by today's market. When one looks forward to recognize that lock-in will be a potential problem, this provides a reason for government policy to encourage more diversity before the standard is set. Or if we seem stuck with an inferior standard, public policy can guide a coordinated change from one standard to another. Moving from measurements in inches and feet to the metric system is one example; coordinating the use of daylight saving time is another.

Inferior standards may be behavioral rather than technological. Examples include an equilibrium in which everyone cheats on his taxes, or drives above the speed limit, or even just arrives at parties an hour after the stated time. The move from one equilibrium to a better one can be most effectively accomplished via a short and intense campaign. The trick is to get a critical mass of people to switch, and then the bandwagon effect makes the new equilibrium self-sustaining. In contrast, a little bit of pressure over a long period of time would not have the same effect.

The other general problem with laissez-faire is that so much of what matters in life takes place outside the economic marketplace. Goods ranging from common courtesy to clean air are frequently unpriced, so there is no invisible hand to guide selfish behavior. Sometimes creating a price can solve the problem, as with congestion on the Bay Bridge. Other times, pricing the good changes its nature. For example, donated blood is typically superior to blood that is purchased, because the types of individuals who sell their blood for money are likely to be in a much poorer state of health. The coordination failures illustrated in this chapter are meant to show the role for public policy. But before you get carried away, check the case below.

CASE STUDY: A PRESCRIPTION FOR ALLOCATING DENTISTS

In this case study, we explore the coordination problem of how the invisible hand allocates (or misallocates) the supply of dentists between cities and rural areas. In many ways the problem will seem closely related to our analysis of whether to drive or take the train from Berkeley to San Francisco. Will the invisible hand guide the right numbers to each place?

It is often argued that there is not so much a shortage of dentists as a problem of misallocation. Just as too many drivers, left to their own resources, would take the Bay Bridge, is it the case that too many dentists choose the city over the countryside? And if so, does that mean society should place a toll on those who want to practice city dentistry?

For the purposes of this case study, we greatly simplify the dentists’ decision problem. Living in the city or in the countryside are considered equally attractive. The choice is based solely on financial considerations—they go where they will earn the most money. Like the commuters between Berkeley and San Francisco, the decision is made selfishly; dentists maximize their individual payoffs.

Since there are many rural areas without enough dentists, this suggests that there is room for an increased number of dentists to practice in rural areas without causing any congestion. Thus rural dentistry is like the train route. At its best, being a rural dentist is not quite as lucrative as having a large city practice, but it is a more certain route to an above-average income. Both the incomes and the value to society of rural dentists stay roughly constant as their numbers grow.

Being a city practitioner is akin to driving over the Bay Bridge—it is wonderful when you are alone and not so great when the city gets too crowded. The first dentist in an area can be extremely valuable and maintain a large practice. But with too many dentists around, there is the potential for congestion and price competition. As the number increases, city dentists will be competing for the same patient pool, and their talents will be underutilized. If the population of city dentists grows too much, they may end up earning less than their rural counterparts. In short, as the number of city practices increases, the value of the marginal service that they perform falls, as does their income.

We depict this story in a simple chart, again quite similar to the driving versus train example. Suppose there are 100,000 new dentists choosing between city and rural practices. Thus, if the number of new city dentists is 25,000, then there will be 75,000 new rural dentists.

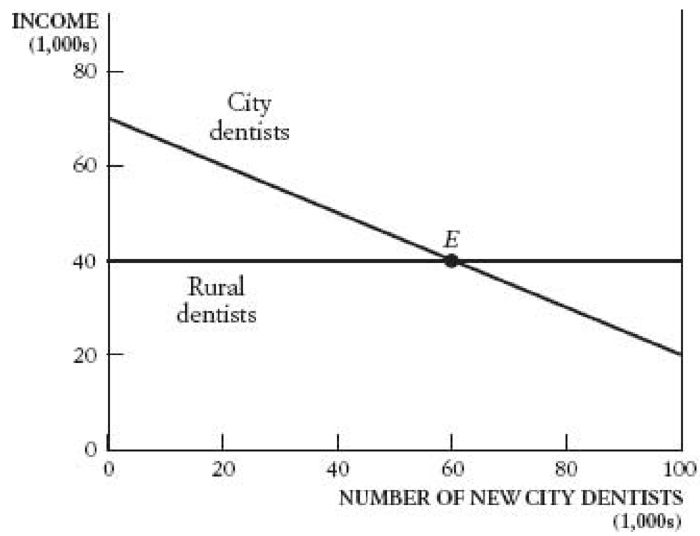

The falling line (city dentists) and the flat line (rural dentists) represent the financial advantages of taking the respective paths. At the far left, where everyone chooses rural practices, city dentists' incomes are above the incomes of those with rural practices. This is reversed at the far right, where everyone chooses city dentistry.

The equilibrium for career choices is at *E*, where the two options provide the same financial rewards. To verify this, suppose that the distribution of career choice results in only 25,000 new city dentists. Since city dentists' incomes are higher there than rural dentists' incomes, we expect that more new dentists will choose city over rural practices. This will move the distribution of city vs. rural to the right. The reverse adjustment would take place if we started to the right of *E*, where city dentists are the lower paid of the two. Only when *E* is reached will next year's career choices broadly replicate those of this year, and the system will settle down to an equilibrium.

Is this outcome the best for society?

Case Discussion

As in the case of the commuters, the equilibrium does not maximize the combined income of dentists. *But society cares about the consumers of dentistry as well as the practitioners*. In fact, left alone, the market solution at *E* is the best for society as a whole. The reason is that there are two side effects created when one more person decides to be a city dentist. The additional city dentist lowers all other dentists' incomes, imposing a cost on the existing city dentists. But this reduction in price is a benefit to consumers. The two side effects exactly cancel each other out. The difference between this story and our commuting example is that no one benefited from the extra commuting time when the Bay Bridge became congested. When the side effect is a change in price (or income), then the purchasers benefit at the producers' cost. There is zero net effect.

From society's viewpoint, a dentist should not worry about lowering colleagues' incomes. Each dentist should pursue the highest-paying practice. As each person makes a selfish choice, we are invisibly led to the right distribution of dentists between city and rural areas. And, the two careers will have equal incomes.*

Of course the American Dental Association may look at this differently. It may place more weight on the loss to city dentists' incomes than on the saving to consumers. From the dental profession's perspective there is indeed a misallocation, with too many dentists practicing in the city. If more dentists took rural practices, then the potential advantages of a city practice would not be "wasted" by competition and congestion. Taken as a whole, the income of dentists would rise if it were possible to keep the number of city dentists below the free-market level. Although dentists cannot place a toll on those who want to practice in the city, it is in the profession's self-interest to create a fund that subsidizes dental students who commit to establish a rural practice.

For some more case studies on cooperation and coordination, see "Here's Mud in Your Eye," "A Burqa for Prices," and "The King Lear Problem" in chapter 14.