

A reprint from
American Scientist
the magazine of Sigma Xi, The Scientific Research Society

This reprint is provided for personal and noncommercial use. For any other use, please send a request Brian Hayes by electronic mail to bhayes@amsci.org.

Playing in Traffic

Can warning drivers of traffic jams make congestion worse? Can closing roads make it better? Mathematically yes, but real-world confirmation is hard to find.

Brian Hayes

I was southbound on Interstate 95, approaching Washington, DC, on a summer afternoon. The Capital Beltway offers two routes around the city—an eastern loop via Greenbelt, MD, and a western arc through Tysons Corner, VA. As I approached the decision point, online traffic maps showed several slow spots on the western branch, whereas the eastern roadway was flowing freely. The choice seemed clear, yet an unsettling thought kept nagging at me: Other drivers had access to the same information I was seeing. If we all followed the recommended route, our strategy would be self-defeating. In collectively avoiding one traffic jam, we'd create a new one.

Traffic patterns present many such puzzles and perplexities. (Pondering them can help pass the time when you're caught in gridlock.) One of the most intriguing ideas in the theory of transport networks is Braess's paradox, which says that building a new road to relieve congestion can sometimes have the opposite effect, causing greater delays for all drivers. Conversely, closing off a road can sometimes speed everyone's journey.

In trying to better understand these counterintuitive cloggings and clearings of roadways, I have been playing with computer simulations in which I can watch individual vehicles wend their way through a network of roads, choosing a path at each intersection. Although the model is a simple one, it does show evidence of Braess's paradox, along with an abundance of other curious instabilities and oscillations.


Brian Hayes is senior writer for American Scientist. Additional material related to the Computing Science column can be found online at <http://bit-player.org>. E-mail: brian@bit-player.org

Whether the results will help drivers—or future driverless vehicles—navigate real highways remains to be seen.

Selfishness on Wheels

Braess's paradox is named for Dietrich Braess of Ruhr University in Bochum, Germany, who described it in 1968. His key assumption in formulating his model is now known as *selfish routing*: Each driver chooses whatever route minimizes his or her own travel time. The system is in equilibrium if no driver can get to the destination quicker by switching to a different route.

The road network is diamond-shaped, with two routes leading from a start node to an end node.




Each route consists of two segments. The thick blue links labeled A and B are wide roads where everyone drives at the speed limit no matter how heavy traffic might become. Each of these segments has a fixed travel time of one hour. The thinner red segments labeled a and b are narrow roads susceptible to congestion. The travel time on one of these roads varies in proportion to the fraction of traffic choosing that road. If there's no one on the road, the travel time goes to zero. But if everyone funnels into a single red segment, the travel time for that road is a full hour (the same as that on the fixed-speed blue links).

In this network a sensible driver chooses the route with less traffic (or chooses randomly if the routes happen to be equal). This is the selfish solution. It is also the optimum solution for the

whole population, attaining the lowest average driving time for everyone. At equilibrium half the cars take the northerly loop and half the southerly, and everyone spends 90 minutes en route.

Now add an extra link to the network, labeled X:



This golden highway has zero traversal time for any number of vehicles. Suppose you are just about to depart the start node when the X link is opened. Up to that moment, traffic has been evenly divided between the Ab and aB routes, with travel times of 90 minutes each. You observe that the aXb route will take only 60 minutes, and so you eagerly select it. The trouble is, everyone else makes the same calculation, and so all the traffic migrates to the aXb pathway. Now both the a and b segments are saturated with 100 percent of the traffic, and everyone has a two-hour trip—30 minutes longer than before the golden link was opened.

When I first read about the Braess paradox, I felt sure there must be some flaw in the argument. Intelligent drivers would realize that they are all suffering needlessly, and so they would return to their original Ab and aB routes, ignoring the new crosslink. Under the rules of selfish routing, however, that happy outcome is unattainable.

Suppose you arrive at the starting node with the network in its dysfunctional condition, with everyone crawling bumper to bumper through route aXb. You might be tempted to strike out in a different direction. All four

pathways— Ab , aB , aXb , and AXB —promise exactly the same travel time of two hours, so the choice appears to be a matter of indifference. But suppose you take route AXB while everyone else continues on aXb . Your defection reduces the load on segments a and b , and so traffic there speeds up slightly. The selfish routing principle requires you to switch back, even though no one will benefit from that decision.

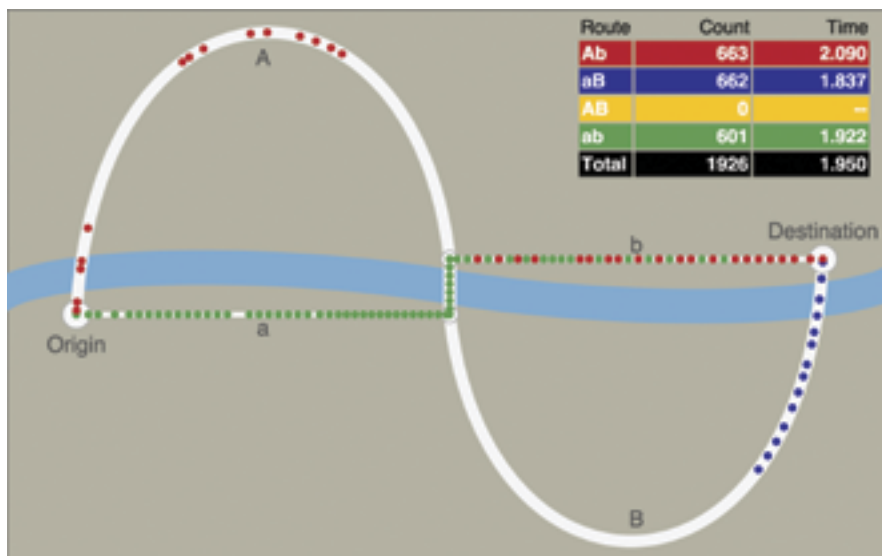
The Price of Anarchy

The Braess paradox might be seen as a tragic quirk of human nature: When everyone strives to get ahead of everyone else, we all fall behind. The phenomenon has been termed “the price of anarchy”—the penalty we pay for failing to coordinate our actions.

But it won’t do to put too much emphasis on the social and psychological aspects of these events. In 1991 Joel E. Cohen of Rockefeller University and Paul Horowitz of Harvard showed that the paradox can affect even inanimate systems, where no human foibles can possibly be to blame. Their most striking example consists of a weight suspended from a certain arrangement of strings and springs. Cutting one of the supporting strings causes the weight to rise rather than fall! The same effect is seen in electrical networks, where opening one branch of a circuit can allow more current to flow.

As a mathematical artifact, the Braess model has the virtues of simplicity, elegance, and determinism. As a model of road traffic, however, the framework is somewhat artificial and unrealistic. For example, some roadways have unlimited capacity, and others allow infinite speed. Moreover, there are really no cars in the model, just rates of flow. Once those rates of flow reach equilibrium, nothing ever changes; it’s a static model.

In recent years a few investigators have addressed the Braess paradox through a more algorithmic and mechanistic style of modeling. The idea is to trace the progress of individual vehicles in time and space as they select paths according to the selfish routing rule. These models reveal the *dynamics* of the system—how vehicles clump together or spread apart, how disturbances propagate through the network, how traffic jams evolve and dissolve. A disadvantage of this approach is that proving theorems becomes more difficult, but gaining intuition may be easier.



A traffic model explores the effects of routing strategies—rules for choosing which pathway to follow from Origin to Destination. Roads A and B are long but wide; they allow high-speed travel regardless of traffic. Roads a and b are shorter but also subject to congestion; in heavy traffic, the shortcuts are no faster than the long way around. The dots are cars moving through the network; their colors indicate their route choices. The table of timings at the upper right shows that in this case the shortest route ab is not the fastest.

Traffic engineers have built highly detailed dynamic models, with realistic roadway geometry and vehicle physics. I have chosen a more schematic design, closer to the abstract formulation of Braess’s paradox, with just a faint whiff of physical plausibility. Roads are lines or curves; cars are colored dots; simple formulas define speed as a function of traffic density. If you would like to play with the model yourself, it’s available at <http://bit-player.org/extras/traffic>.

Bypasses and Shortcuts

The layout of the model is inspired by a version of the Braess paradox introduced in 1997 by Claude M. Penchina of the University of Massachusetts, Amherst. Two cities, Origin and Destination, lie along a river, connected by straight, narrow roads a and b and by looping freeways A and B . The freeways are twice as long as the straight segments, but they allow speed-limit driving for any number of cars; the surface roads, in contrast, are subject to congestion. At the midpoint of the river a bridge connects the two routes, but initially the bridge is closed. Although the geometry of this tableau is quite different from that of the diamond network, the topology is identical. The same four links (or five when the bridge is open) connect the same four nodes.


In the classic Braess model, speed on the a and b links varies from infinity down to 1, and travel time ranges

from 0 up to 1. To bring the dynamic model a little closer to physical realism, maximum speed is limited to 1 everywhere. On the a and b links travel time is proportional to the number of vehicles occupying the link. The constant of proportionality, or congestion coefficient, ranges from 0 (congestion has no effect) to 1 (at maximum occupancy traffic comes to a halt). For most of the experiments described here, the congestion coefficient was equal to $\frac{1}{2}$; with this setting, traversing a congested shortcut takes just as long as driving twice as far on a freeway loop.

The algorithm that controls the model computes motion in discrete time steps. The main loop finds the position of each car, calculates its speed and the distance moved during a step, then updates the position. Before moving the car, however, it’s best to check that there’s nothing in the way. (We want no crumpled fenders.) Only a finite number of cars will fit on a given section of highway, and a car may have to wait its turn to inch ahead in gridlocked traffic.

Queueing Up

The very first run of the simulation vividly demonstrated how modeling with discrete objects differs from the calculation of flow rates. When I pressed the “Go” button, cars poured onto the A loop, rounded it at high speed, and then began spilling onto the b segment.



A space-time diagram captures the trajectories of 500 vehicles moving through the model network with the crosslink bridge open. Time proceeds from left to right, and distance traveled increases from bottom to top. Thus the slope at any point along a line represents the car's speed

at that point; steeper trajectories represent faster motion. Areas where the slope is nearly horizontal result from bottlenecks at intersections. The alternating bands of color reveal a persistent oscillation as successive platoons of cars favor first one route and then another.

But as more cars entered that link, the entire parade slowed down, and soon the queue was backing up onto the *A* loop. Nothing like this can happen in the static model. (I later learned that Carlos F. Daganzo of the University of California at Berkeley had discussed this spillback effect in the 1990s.)

Another surprise awaited me when I began running the model with the crossover bridge closed. In this circumstance the static model tells us that the flow of vehicles should split into two equal streams, coursing through the *Ab* and *aB* paths. That's also the approximate result in the dynamic model if you look at long-term averages. On a moment-by-moment basis, however, the streams are anything but steady and equal. Instead there is a persistent oscillation, with heavy traffic first on one route and then the other. A thought experiment explains why. Suppose it's early morning, and you are the first commuter of the day. The roads are empty, and so the two routes look equally attractive; you flip a coin and drive onto route *aB*. As soon as you enter the *a* segment, speeds there

are slightly reduced because of your presence, with the result that all drivers behind you prefer route *Ab*. However, when the first of those drivers arrive at the *A*-to-*b* junction, their speeds also begin to fall, and the *aB* path returns to favor. In this regime cars tend to form platoons, all selecting the same route until it becomes overcrowded, so that the next platoon goes the other way. (An instability of this kind is exactly what I had imagined on my drive around the Capital Beltway.)

Platooning imposes a penalty for selfish routing that's independent of the Braess paradox; it happens even when the crossover link is closed. In the dynamic model, selfish routing yields longer trips than a simple random choice of routes.

Searching for a Paradox

The next question is whether the model exhibits a Braess effect: When we open the bridge, allowing drivers to slip down the shortcut *ab* path, does travel time increase or decrease? The dynamic model is different in many details from the system Braess stud-

ied, so it's entirely possible that the paradox would not appear. But in fact the model *does* charge a price for anarchy over a wide range of parameter settings. An exception is at very low traffic density, where the effects of congestion are mild, and the short *ab* route is quickest. At densities beyond about 0.25—where vehicles fill one-fourth of the available road space—selfish routers continue to choose *ab* even though it raises their average trip time.

Quite a lot has been written about how the Braess paradox varies as a function of traffic intensity. In 1997 Penchina and independently Eric I. Pas and Shari L. Principio of Duke University found that the paradox exists only within a certain intermediate range of traffic densities; Anna Nagurney of the University of Massachusetts, Amherst, later reached a similar conclusion by another method. Under extreme conditions the added link no longer causes mischief because drivers cease using it. Why do they abandon the shortcut? Not because they have finally learned to cooperate for the common good; it's just that the *ab* path becomes so clogged that other choices are more attractive.

These analyses were based on a static model. Wei-Hua Lin and Hong K. Lo of the University of Arizona have studied the same question in a dynamic model and reached a different conclusion: The paradox persists even at high levels of congestion. My casual experiments yield equivocal results. At the highest traffic loads, usage of the *ab* path and the penalty for selfish routing both decline, but they do not fall to zero. On the other hand, by raising the congestion coefficient from $\frac{1}{2}$ to $\frac{3}{4}$ I can eliminate the paradox, as opening the bridge link becomes strongly beneficial rather than detrimental. In this state most of the traffic on the bridge



Transit time in the model network shows the cost of selfish routing, as well as the paradoxical effect that opening an extra road can slow everyone's trip. Except at very low traffic density, selfishly routed cars reach the destination faster when the midriver bridge is closed. Adopting a random routing strategy has an even greater benefit.

is going the other way—from north to south on the long *AB* route.

Ground Truth

The Braess paradox is not just a mathematician's toy; traffic planners take it quite seriously. However, sightings of the phenomenon in the wild are scanty.

An oft-repeated tale describes the closing of 42nd Street in New York on Earth Day in 1990. Instead of the expected gridlock, traffic flow improved. The incident seems to be documented only in a *New York Times* article by Gina Kolata, which gives no indication that the effects of the street closing were measured or analyzed. Other cases in Stuttgart, Germany, and Seoul, Korea, also seem to be supported more by anecdote than by evidence.

In 2008 Hyejin Youn and Hawoong Jeong of the Korea Advanced Institute of Science and Technology and Michael T. Gastner of the Santa Fe Institute reported finding dozens of streets in Boston, London, and New York that exhibit the paradox; traffic flow between two specific origin and destination nodes would improve if those streets were closed. But this conclusion was reached through theoretical analysis of traffic patterns, not by experiment.

Perhaps the most thorough search for ground truth came in the aftermath of a catastrophe in 2007: the sudden collapse of a major bridge in Minneapolis, carrying Interstate 35W across the Mississippi River. As the replacement bridge was completed a year later, Shanjiang Zhu, David Levinson, and Henry Liu of the University of Minnesota prepared a before-and-after study. Tracking devices were installed in the automobiles of 187 volunteers, whose commuting routes and trip times were recorded over eight weeks. No evidence of a Braess paradox was found: Average travel times improved after the new bridge was opened. A later lane closing on another bridge had a larger effect, but again the change was not in the paradoxical direction.

Urban street grids look nothing like the small and simple networks considered in these models. Cities have hundreds of streets, and drivers navigate between many points of origin and destination. Perhaps this noisy environment protects us from the suboptimal choices of selfish routing—or it may just obscure the signal of the paradox.

It's also possible that we are just now reaching the point where real-world traffic begins to fulfill the underlying assumptions of the models. One of those assumptions is that all drivers have complete, timely, and accurate information about traffic conditions. For many years this was highly unrealistic; our routing decisions were less than perfectly selfish as a result of mere ignorance. GPS units and online mapping services have changed all that, and there is doubtless more to come, such as machine learning to detect and forecast traffic patterns in complex networks. What if everyone adopts the same algorithm?

The self-driving automobile, now under development by companies ranging from Google to Ford, could raise the stakes. With direct access to detailed traffic statistics—past, present, and forecast—such a vehicle would be a formidable player in the game of selfish routing. But an autonomous car could also communicate and cooperate with other vehicles to escape the cycle leading to Braess's paradox. It will be interesting to see which trend wins out.

Bibliography

- Arnott, Richard, and Kenneth Small. 1994. The economics of traffic congestion. *American Scientist* 82:446–455.
- Braess, Dietrich. 1968. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* 12:258–268. English translation by Dietrich Braess, Anna Nagurney, and Tina Wakolbinger. 2005. On a paradox of traffic planning. *Transportation Science* 39:446–450.
- Buscema, Daniele, Matteo Ignaccolo, Giuseppe Inturri, Alessandro Pluchino, Andrea Rapisarda, Corrado Santoro, and Salvatore

Tudisco. 2009. The impact of real time information on transport network routing through intelligent agent-based simulation. *Proceedings of the TIC-STH Symposium on Human Factors and Ergonomics*, p. 72.

Cohen, Joel E. 1988. The counterintuitive in conflict and cooperation. *American Scientist* 76:577–584.

Cohen, Joel E., and Frank P. Kelly. 1990. A paradox of congestion in a queuing network. *Journal of Applied Probability* 27:730–734.

Cohen, Joel E., and Paul Horowitz. 1991. Paradoxical behaviour of mechanical and electrical networks. *Nature* 352:699–701.

Daganzo, Carlos F. 1998. Queue spillovers in transportation networks with a route choice. *Transportation Science* 32(1):3–11.

Frank, Marguerite. 1981. The Braess paradox. *Mathematical Programming* 20:283–302.

Kolata, Gina. 1990. What if they closed 42d street and nobody noticed? *New York Times* 25 December 1990.

Lin, Henry, Tim Roughgarden, Eva Tardos, and Asher Walkover. 2011. Stronger bounds on Braess's paradox and the maximum latency of selfish routing. *SIAM Journal of Discrete Mathematics* 25(4): 1667–1686.

Lin, Wei-Hua, and Hong K. Lo. 2009. Investigating Braess' paradox with time-dependent queues. *Transportation Science* 43(1):117–126.

Nagurney, A. 2010. The negation of the Braess paradox as demand increases: The wisdom of crowds in transportation networks. *Europhysics Letters* 91:48002.

Pas, Eric I., and Shari L. Principio. 1997. Braess' paradox: Some new insights. *Transportation Research B* 31(3):265–276.

Penchina, Claude M. 1997. Braess paradox: maximum penalty in a minimal critical network. *Transportation Research A* 31(5):379–388.

Rapoport, Amnon, Tamar Kugler, Subhasish Dugar, and Eyran J. Gisches. 2009. Choice of routes in congested traffic networks: Experimental tests of the Braess paradox. *Games and Economic Behavior* 65(2):538–571.

Roughgarden, Tim. 2005. *Selfish Routing and the Price of Anarchy*. Cambridge, MA: MIT Press.

Steinberg, Richard, and Willard I. Zangwill. 1983. The prevalence of Braess' paradox. *Transportation Science* 17(3):301–318.

Youn, Hyejin, Michael T. Gastner, and Hawoong Jeong. 2008. The price of anarchy in transportation networks: efficiency and optimality control. *Physical Review Letters* 101:128701.

Zhu, Shanjiang, David Levinson, and Henry Liu. 2009. Measuring winners and losers from the new I-35W Mississippi River Bridge. In *Transportation Research Board 89th Annual Meeting Compendium of Papers*. Washington: Transportation Research Board.