The Economics of Traffic Congestion

Richard Arnott
and
Kenneth Small
(University of California, Irvine)

Working Paper No. 256
May 1993

THE ECONOMICS OF TRAFFIC CONGESTION

Richard Arnott*

and

Kenneth Small**

* Department of Economics
Boston College
Chestnut Hill, MA 02167

** Department of Economics
University of California
Irvine, CA 92717

The Economics of Traffic Congestion

Traffic congestion is a plague of modern life. Time spent ensnarled in traffic is not only time wasted, but for most of us is time miserably wasted.

The dimension of the problem can be gauged from a simple back–of–the–envelope calculation for the 39 U.S. metropolitan areas of over one million population. In such areas, roughly one–third of all vehicular travel occurs under congested conditions in which speed averages half its free–flow value. About half of this congested driving is on expressways, causing about 0.6 minutes delay per kilometer of travel; half is on other arterials, causing about 1.2 minutes delay per kilometer of travel. With some 75 million licensed drivers in such areas, each averaging 16,000 kilometers per year within the region, this amounts to approximately 1,200 billion congested vehicle–kilometers of travel annually, for total delays of 6 billion vehicle–hours. Valuing each vehicle–hour at $8.00 suggests an annual delay cost of $48 billion, or $640 per driver. This does not include the costs of disruption from the unpredictability of these delays, the costs of inconvenient schedules caused by attempts to avoid them, nor the costs of extra fuel, accidents, and air pollution.

Such congestion has policy–making itself in gridlock. Every policy considered either is too unpopular, is too expensive, or has proven ineffective. Why is congestion so intractable, and what can be done?

Answering these questions turns out to require a sophisticated understanding of the behavioral interactions that determine when and where congestion occurs. We start by discussing why the standard remedies — expanding highway capacity or mass transit — have been unsuccessful. We then present three paradoxes, in which a highway expansion designed to relieve congestion instead proves ineffective or counter–productive. The resolution of these paradoxes, which employs the economic concept of externalities, not only clarifies the economics of traffic congestion but also points to how the congestion problem can be solved using clever applications of the standard pricing tools of economics.

The standard remedy to traffic congestion is to "build our way out." Building our way out, however, would be prohibitively expensive. A few years ago the Southern California Association of Governments estimated the cost of accommodating expected 25–year growth in the Los Angeles region through expansion of highways and new rapid–transit lines. Their cost estimate was $111 billion dollars, a figure that now seems conservative in light of cost escalations in some recent projects. As dense urban development occurs around existing facilities, planning and building new capacity becomes extraordinarily complex, expensive, and politically controversial.

What about instead building new capacity in the form of rapid transit? Experience shows that such an approach is unable by itself to attract more than a tiny fraction of the peak demand for highway facilities. Don Pickrell of the Transportation Systems Center in Cambridge, Massachusetts, meticulously documented the cost of each trip diverted from cars to public transit for eight major rail transit projects. For three of the projects, there was no diversion because bus patronage declined by more than rail patronage rose. For the others, the cost ranged from $9.49 to $34.64 for each new transit trip! Achieving significant reductions in auto congestion through subsidies of this magnitude is obviously infeasible. The advantages of the car are simply too great: not only does it provide considerably more comfort, privacy, and convenience than mass transit, but it is also much better suited to the decentralized American city. Regrettably, the urban sprawl that was encouraged by massive subsidies to auto travel in past decades cannot be reversed.

Even if new highway construction and new mass transit were not so expensive, building our way out of the problem would be much harder than it might appear at first glance. One part of the reason is *latent demand.* The traffic we see is not the full demand for peak travel at the prevailing *monetary* cost, since congestion causes many potential rush–hour vehicle trips to be canceled, diverted (to mass transit, car pools, less congested routes and locations, etc.), or rescheduled. Any reduction in congestion resulting from capacity expansion will be at least partially undone by latent demand. The other reason is that *congestion is mispriced.* Because drivers do not pay for the time loss they impose on others, they make inefficient choices concerning how much to travel, when to travel, where to travel, and what route to take. As our paradoxes will show, the combination of latent demand and mispriced congestion may be so perverse that an expansion of capacity causes no change in congestion or even worsens it.

These paradoxes, besides illustrating why building our way out of the problem is difficult or impossible, also set the stage for a solution. They do this by introducing the idea of externalities, which underlies policy approaches based on pricing.

First is the *Pigou–Knight–Downs paradox*, named for two noted economists of the early twentieth century plus Anthony Downs, now of the Brookings Institution in Washington, D.C. Suppose 1,000 peak–hour travelers between two cities can choose between (1) a direct route containing a narrow bridge and (2) a more circuitous but wide road (see figure 1). The first route takes 10 minutes with no traffic, but travel time rises linearly with the ratio of traffic flow F_1 to bridge capacity C_1; specifically, travel time is $10 + 10\ (F_1/C_1)$. The second route takes 15 minutes. Each traveler chooses the road with the lower travel time.

If C_1 is greater than 2000, everyone takes the first route, whose travel time is then under 15 minutes. The paradox occurs when C_1 is less than 2000. In this case, travelers divide themselves across the two routes such that travel time on each route is 15 minutes, which implies that $F_1/C_1 = .5$. Thus, expanding bridge capacity anywhere in the range from 0 to 2000 has absolutely no effect on anyone's travel time. It only diverts more people from the route with spare capacity to the route crossing the bridge. The new bridge capacity literally generates its own demand.

Attempts to reduce congestion on the bridge by instead encouraging car pooling, expanding mass transit, or improving telecommunication facilities would likewise be frustrated unless total vehicular traffic were reduced below $.5C_1$. So long as any traffic remained on the second route, latent demand for the bridge would undermine these attempts to relieve its congestion.

The crux of the paradox lies in the distinction between the private and social cost of a trip. The private cost is the cost the driver himself incurs. The social cost equals this cost plus the cost he imposes on other drivers by slowing them down, which is termed the *external cost.* In the example, the social cost of traveling on the uncongested route equals the private cost, but the social cost of traveling on the congested route exceeds the private cost. Typically, drivers choose the route with the lower private cost; this results in an equilibrium in which private costs on the two routes are equalized. If instead drivers were distributed across the two routes so as to equalize *social* cost, the paradox would disappear — bridge expansion would relieve congestion. This suggests that conventional policies to relieve congestion would work better if each driver faced the social cost of his trip — but that is getting ahead of the story.

Our second example, called the *Downs–Thomson paradox* (for J.M. Thomson, formerly of the London School of Economics), is even more perverse. It is like the previous example except the alternative to taking the congested route is now a privately operated train line which just covers its costs by operating full trains. If more people take the train, then trains are run more frequently, saving people some waiting time at the station. Suppose total travel time by train, including waiting, is $20 - (F_2/300)$ minutes when there are F_2 train riders per hour (see figure 2). Each traveler chooses the faster mode, so that travel times are equalized when both modes are utilized.

The resulting equilibrium is calculated in Figure 2. Its intriguing feature is that now travel time *increases* with any increase in bridge capacity within the range from zero to 1,000. The reason is that, just as in the earlier example, capacity expansion diverts people to the congested road. But now the diversion causes train service to get worse, so equilibrium can occur only when congestion is worse also. Here, new capacity generates more than its own demand!

The reason this paradox is even more perverse than the previous one is that there is not only an external cost imposed by each automobile user, as before, but there is now in addition an *external benefit* created by each user of the train. This is because using the train causes the frequency of service to increase and hence reduces other users' waiting times. This is a technological property of all types of mass transit, including bus and even taxicab service, as was demonstrated in 1972 by Herbert Mohring of the University of Minnesota.

The same perverse result occurs if instead of expanding the road, well–intentioned planners entice some fraction of travelers away from both routes by providing some third alternative such as subsidized vanpools, telecommuting centers — or even a new train service! Nor is this unrealistic; the cases studied by Pickrell included some where initiating a new train service diverted so much traffic from the *existing* transit system (in this case, bus transit) that the overall quality of transit service deteriorated, causing a net diversion to automobiles and presumably a worsening of road traffic. Had the existing transit service been improved instead, the improvements might have reinforced rather than thwarted the external benefits inherent in transit service.

Our final example is the *Braess paradox*, named for a German engineer who in 1968 described an abstract road network in which adding a new link causes total travel time to increase. Our version involves 1000 people traveling from district A of a city to district B, where the districts are separated by marshland. District A lies south of a river at the west end of a fen, district B north of the river at the east end of the fen (see figure 3). There are two routes from A to B. Route 1 crosses the river at bridge A and circles north of the fen to B. Route 2 circles south of the fen and crosses the river at bridge B. Travel on both routes is uncongested except at the bridges. Travel time on either route is 15 minutes plus $F_A/100$ or $F_B/100$ for bridges A or B respectively. In equilibrium, 500 people take each route, with travel time 20 minutes.

A causeway is now constructed across the fen from the north end of bridge A to the south end of bridge B. The causeway can be traversed in 7.5 minutes regardless of traffic volume. There is now a third route from A to B — across bridge A, along the causeway, and then across bridge B. What happens when the causeway is opened? Each bridge now carries the traffic for two distinct routes, the previous one (F_1 or F_2) plus the causeway route (F_3). Travel time on bridge A now becomes $(F_1 + F_3)/100$ and travel time on bridge B $(F_2 + F_3)/100$. In equilibrium, travel time is equated on all three routes. This gives us two equations plus the condition that the three traffic volumes add to 1000, so we can readily solve for all three flows. The result is that half the traffic takes the causeway route, and the other half divides evenly between the two previous routes. Hence each bridge carries 750 travelers, fifty percent more than before, producing a travel time on each route of 22.5 minutes. Adding the causeway has made everyone's trip longer!

By now, the reader probably recognizes that the paradox is explained by *congestion externalities* on the bridges. Because each traveler ignores the external cost he imposes by crossing a bridge, too many people choose the causeway route, which crosses both bridges. The faster the causeway, the more people are so enticed and the worse is their trip: if causeway traversal time were only 5 minutes, all 1000 would choose that route and its travel time would rise to 25 minutes. Only if the causeway speed were infinite would equilibrium travel time return to its original 20 minutes!

Are these paradoxes more than intellectual curiosities? It has been claimed that the Braess paradox explains some traffic problems observed in Stuttgart and Manhattan, and Odd Larsen of the Institute of Transport Economics in Oslo suggests it would apply under certain circumstances to a proposed widening of the main road joining Oslo with its airport. Martin Mogridge of University College, London, has forcefully, if controversially, asserted that the Downs–Thomson paradox explains the deterioration of road speeds over twenty years or so in central London. As for the Pigou–Knight–Downs paradox, it is so enshrined in transportation planning that it is often called "the fundamental law of traffic congestion."

The concept of externalities provides a powerful tool for analyzing congestion in a more general context. An externality occurs when a person does not face the true social cost of an action. By modeling congestion systematically, it is possible to define the social cost of driving on a congested road by observing how aggregate travel delays are related to the number of travelers. Combining this with a model of demand for the road, one can determine both equilibrium travel patterns (as in the above examples) and optimal travel patterns under some defined objective such as minimizing aggregate travel time.

Economists have applied this idea using the *flow congestion* model from the traffic engineering literature, which is built on an analogy with compressible fluids. This model begins with the identity that *flow* or *volume* (vehicles per hour going past a point) equals *density* (vehicles per kilometer of road) times *speed* (kilometers per hour). It postulates a physical relationship between any two of these quantities. For example, the Greenshields model postulates that speed decreases linearly with density until it reaches zero at a point called the "jam density"; this model implies a parabolic relationship between flow and speed. Similar models form the basis for capacity calculations in the Transportation Research Board's influential *Highway Capacity Manual.*

The relationship between speed and flow is easily converted to one relating a driver's travel time to flow. In order to apply the concept of externalities, we need to convert travel time to a cost. For simplicity, let us ignore the out–of–pocket costs of travel, and assume that everyone places an identical monetary value on each minute of travel time. Multiplying travel time (T in the examples) by this value of time then gives the private cost of a trip.

Thus, we obtain the *private cost* curve relating private cost to flow (see figure 4). The lower (solid) part of the curve, marked pc(F), where private cost is increasing in flow, corresponds to situations of modest congestion and lends itself to analyzing the congestion externalities discussed earlier. At any level of flow, we can calculate total cost as the flow multiplied by each driver's private cost pc(*F*). Then we can calculate how much total cost increases when flow is increased by one unit, which is referred to as the *social cost* of a trip. The social cost curve is plotted as sc(*F*) in the figure. By definition, the social cost of a trip equals the private cost plus the external cost — the cost the added driver imposes on others by slowing them down. Thus, the external cost equals the vertical distance between the social cost and private cost curves. (Because drivers impose these external costs on each other, multiplying the social cost curve by flow does not lead to a meaningful total cost.)

The demand for using a road is generally some flat or downward–sloping function of the private cost. (In our examples it was flat, an extreme case caused by the assumed perfect substitutability between the alternative routes). Such a function is shown in the figure as D(*F*). Equilibrium occurs at point A, where the demand curve intersects the private cost curve, since at this level of flow the benefit of an extra trip equals its private cost. Efficiency, however, obtains at point B where the benefit of an extra trip equals the *social* cost. In equilibrium, travel is underpriced because drivers do not pay for the congestion they cause, and consequently too many trips are taken.

How can planners create conditions under which the system will operate at point B instead of A? The answer is simply to charge a money payment, known as a *congestion toll*, equal to the external cost. In the figure, the optimal congestion toll τ is measured by the vertical distance between the social and private cost curves at point B. By thus bringing the private cost faced by the traveler up to the level of social cost, privately–made decisions will lead to the social optimum (point B). This analysis is exemplified for the Downs–Thomson paradox in the technical panel on page ___.

Such a policy is known as *congestion pricing*, and versions of it are now under consideration in Great Britain, Holland, and Hong Kong following implementations in the central area of Singapore (in operation since 1975) and, more recently, on a motorway outside Paris. In the United States, one or more demonstration programs, funded by 1991 federal highway legislation, are likely to join a private road recently authorized in California as early congestion pricing experiments.

This policy is also an example of *marginal–cost pricing*, a term with much broader meaning. Briefly, marginal–cost pricing refers to setting the price of a unit of commodity equal to the incremental cost of producing one more unit of the commodity.

Mathematically, it is the derivative of the total cost function. In the traffic context, the social cost of a trip is the increment in total cost to all travelers caused by adding one more trip; by facing the trip maker with this social cost, we are effectively setting the "full price" of the trip (including both money and time) equal to its marginal cost.

What about overall welfare? Isn't everyone made worse off by being forced to pay a toll that raises the cost of using the road, even taking into account the reduced congestion? The answer is "yes" (unless demand is flat) when regarding people only in their roles as travelers. But paying a toll does not use up resources; it is only a paper (or more likely, an electronic) transaction. If the toll revenues are used to benefit citizens generally, then the gains people receive as citizens more than offset their losses as travelers. In fact, the more formal statement of "efficiency" is precisely this: there is *some* way of redistributing the toll revenues which leaves *everyone* as well or better off than before.

To illustrate this, we return to our first example, the Pigou–Knight–Downs paradox. If everyone values their time equally, the efficient traffic allocation across routes minimizes aggregate travel time. This can be calculated as $F_1/C_1 = .25$, with average travel time equal to $15 - (C_1/1600)$. Take the example $C_1 = 1600$. In the efficient allocation, two–fifths of the travelers take the bridge with travel time 12.5 minutes, while the other three–fifths take the longer route with travel time 15 minutes. For simplicity, suppose time is valued at 10 cents per minute; this traffic allocation can then be achieved by charging a money toll of 25 cents for crossing the bridge. Everyone's trip cost is $1.50, either in time (for users of the longer route) or in time plus toll (for bridge users). This is the same trip cost that prevailed in the unpriced equilibrium; thus, any allocation of the $100 in toll revenue makes everyone better off. In more realistic examples, it would probably not be possible to target the redistribution of toll revenues so carefully that everyone was made better off by a toll. Hence in practice, justification for congestion pricing (or any policy change) must rest on the ethical postulate that it is permissible to make some people worse off if the overall gains are great enough. Some argue that this is justified because the existing system of highway finance subsidizes drivers, so the proposed change is actually correcting an existing inequity.

We asserted earlier that all the paradoxes disappear if every driver pays the social cost of his travel. Thus, with (optimal) congestion pricing, expansion of capacity is always beneficial. We have just shown this in the example of the Pigou–Knight–Downs paradox: application of the efficient congestion toll results in average travel time equal to

$15 - (C_1/1600)$, which falls with expansion of the bridge. We invite the reader to check our assertion for the other two paradoxes.

Congestion pricing has the added advantage that it makes transportation planning easier. Whether or not congestion pricing is employed, the merits of a proposed expansion of a transportation link can be evaluated by comparing the cost of expansion with the total cost savings it produces. In the absence of congestion pricing, calculation of these cost savings requires knowing how the expansion will alter traffic flows and travel times over the entire network. To further complicate matters, the distortion introduced by mispricing travel causes market prices to misrepresent the social cost of any land used for the expansion. When, however, congestion pricing is employed, the cost savings may be approximated (for small expansions) by travel cost savings that would occur if no one changed his travel in response to the expansion. Individuals will in fact alter their travel, but since the net social benefit of each of these adjustments equals zero, they net out of the calculation.

A final and very important point concerning congestion is that zero traffic congestion is seldom optimal. Congestion could be eliminated by prohibiting travel, or spending vast sums on transportation systems. And it could probably be reduced to negligible levels by requiring that trips be evenly spread over the twenty–four hours of the day. But any of these "solutions" would generate social costs far in excess of the current costs of congestion. There are huge benefits from the spatial concentration of economic activity, deriving from the reduction of transport costs (even with congestion). There are also great advantages from schedule coordination — having people work at common times and recreate at common times. Congestion is simply a cost that goes hand–in–hand with these benefits.

What congestion pricing does is to reduce *excessive* congestion. It does so not only by reducing the number of trips, but also by better allocating them over time and space. It can best accomplish this using a finely adjusted system of tolls, although of course the advantage of this must be weighed against the costs of collection and the greater difficulty of driver comprehension. An example of such fine tuning is planned for a private roadway that has been approved for construction in the existing median strip of the Riverside Freeway in Orange County, California.

Economists have advocated congestion pricing for at least three decades, since the pioneering work of William Vickrey of Columbia University. However, they have failed to overcome a number of counter–arguments including: costly and inconvenient toll collection, especially on downtown streets; regressive distributional impact, since lower–income people spend a larger proportion of their income on commuting and have less work–schedule flexibility; lack of trust in government to dispose of toll revenues wisely; and small magnitude of benefits.

We submit, however, that the case for congestion pricing is significantly stronger today. Recent applications of conventional tolls to new roads show that the public is more receptive to pricing solutions because of worsening congestion and financial constraints. A recent survey in London, for example, found that a majority of auto commuters would favor congestion pricing if the revenues were used to upgrade the transport system. Other proposals for using the toll revenues address the impacts on income distribution. As for the benefits of pricing, a new generation of models that take into account trip rescheduling produces estimates of benefits many times larger than earlier work based on a rush hour of fixed duration. These types of models, introduced by William Vickrey and further developed by one of us (Arnott) in collaboration with André de Palma of the University of Geneva and C. Robin Lindsey of the University of Alberta, constitute today an active research frontier populated by both economists and traffic engineers.

The most important development affecting the prospects for congestion pricing, however, is the enormous advances in technology for toll collection. A pioneering electronic road pricing scheme tested in Hong Kong a decade ago appears to have been a complete success from an engineering and economic standpoint. The basic idea is simple. Each car is equipped with a device that emits a personalized signal. As the car travels along, its signal is picked up by roadside receptors which are connected to a central computer. Periodically, the computer sends each car owner a bill based on that car's travel history. Enforcement is based on photographing license plates of cars failing to emit the signal. Another variation is to have prepayments coded on a "smart card" mounted in the vehicle, thereby eliminating the need to record the vehicle's location.

Commercial equipment is readily available, and electronic toll collection is now in operation on toll highways or bridges in Oklahoma, Texas, Florida, France, Italy, and

Norway. These have proven that existing technology can handle road–pricing transactions quickly and efficiently with little or no slowing of traffic. There is also good reason to believe that the large sums of money being invested in advanced vehicle information systems (AVIS) will result in significant further improvements in the technology for electronic toll collection. Technical specifications to ensure compatibility across locations are being developed by the U.S. Federal Highway Administration, the State of California, a consortium of northeastern states, and the European Community.

Today, the practical problems of implementation are much narrower and more susceptible to solution by the usual kinds of development efforts that take place in government bureaucracies. For example, every system needs to specify an option for occasional travelers who lack an electronic device. Protection of privacy (a major factor in Hong Kong's decision not to implement the system it tested) is quite feasible, but conflicts to some degree with the need for tracing to correct mistakes. How finely "tuned" to make the pricing system is a question involving tradeoffs between efficiency and simplicity. (At one extreme, the city of Cambridge, England, is proceeding with plans for a pricing system that would depend on actual congestion encountered moment by moment.)

Political acceptability, however, remains the key. A well–designed and credible plan for spending the toll revenues is essential. Only with such a plan can the public be assured that a proposed pricing scheme would provide needed financing for transportation improvements, offset at least some of the regressive distributional impact of the tolls, and protect against misappropriation of the revenues for wasteful purposes. Only time will tell whether such plans can be developed, and whether they can persuade people to give up what many regard as a basic right to free travel, and to trust a government bureaucracy to administer efficiently yet another large, complex program. We suspect that European cities will take the lead, due to greater existing congestion and greater acceptance of state intervention compared to the United States.

Whatever the prospects for congestion pricing, it is clear that congestion is a more complex phenomenon than some of our current policy analyses assume. Only by understanding the full nature of people's travel decisions and how they interact can sensible policies be formulated.

TECHNICAL PANEL (PLACE NEAR END OF ARTICLE)

The simplicity of demand for the two alternative routes in the Downs–Thomson paradox (Figure 2) permits a diagrammatic analysis of the private and social costs on both routes. The figure below enables one to visualize both the paradox itself (an equilibrium, also called a *user optimum*) and the cost–minimizing traffic pattern (a *social optimum* or *system optimum*). It is drawn for the case $C_1 = 750$.

The private cost pc_1 and social cost sc_1 of travel by car on the congested route are plotted as a function of flow F_1, with 0_1 as the origin. The corresponding costs of train travel are plotted backwards, as a function of passenger flow F_2 on the train, with 0_2 as the origin. The distance between the origins is 1000, ensuring that $F_1 + F_2 = 1000$. Note that the social cost of car travel exceeds its private cost; but the social cost of train travel is *less* than its private cost, reflecting the external *benefit* that each train user confers on other train users by causing the frequency of service to increase and hence reducing their waiting time.

The equilibrium (point A) is that division of the travelers between car and train which equalizes the private costs: in this example, two–thirds travel by car. An expansion of the bridge causes line pc_1 to rotate clockwise about its intercept, and hence causes point A to slide along pc_2 to the right and upwards. This indicates that some train users switch to car, and everyone's private cost increases.

The social optimum (point B) is that division of travelers — one–sixth by car, the rest by train — which equalizes the social costs, so that switching one traveler from car to train or vice versa neither increases nor decreases the social cost of accommodating that traveler. At this division, the private costs by car and train (measured by the heights of lines pc_1 and pc_2) are both lower than they were at point A; but because they are not equal, this point is not an equilibrium. However, by imposing a road toll equal to $(sc_1 - pc_1) - (sc_2 - pc_2)$, people will be led to choose the social optimum.

It can be shown that the total private cost associated with point B, $F_1 \cdot pc_1 + F_1 \cdot pc_2$, decreases if road capacity C_1 is expanded. Hence, when social costs rather than private costs are equalized, the paradox disappears.

Note that in this example, the road toll is interchangeable with a train subsidy. In real situations, there are so many substitutes for peak–hour car and train travel that this equivalence breaks down. In such situations, the theory calls for a road toll of $sc_1 - pc_1$ and a train subsidy of $pc_2 - sc_2$.

FIGURE 5 HERE

FURTHER READING

Road Pricing: An International Perspective. A.D. May in Transportation, Vol. 19, No. 4, pages 313–333; 1992.

Urban Transportation Economics. Kenneth A. Small. Harwood Academic Publishers, 1992.

Economics of a Bottleneck. Richard Arnott, André de Palma, and Robin Lindsey in Journal of Urban Economics, Vol. 27, No. 1, pages 111–130; January 1990.

Wardrop's Third Principle: Urban Traffic Congestion and Traffic Policy. David J. Holden in Journal of Transport Economics and Policy, Vol. 23, No. 3, pages 239–262; September 1989.

Figure 1

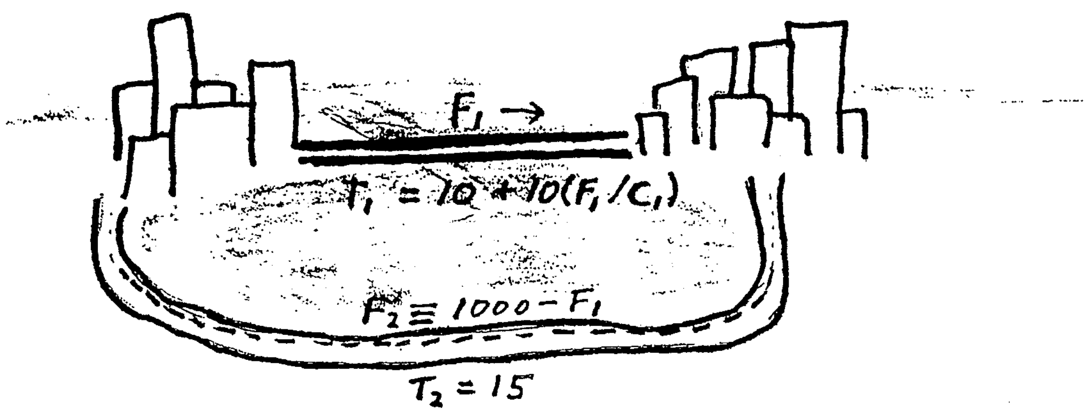

Equilibrium if $C_1 < 2000$: Some traffic uses each route; times are equalized requiring that $10+10(F_1/C_1)=15$.

Solution: $F_1=\frac{1}{2}C_1$, $T_1=T_2=15$.

Expanding capacity has no effect on travel time!

Equilibrium if $C_1 \geq 2000$: Everyone uses bridge.

$F_1=1000$, $F_2=0$, $T_1=10+(10{,}000/C_1)$.

Expanding capacity lowers travel time; no paradox.

PIGOU-KNIGHT-DOWNS PARADOX shows how expanding road capacity can create its own demand. Equilibrium requires that no one can reduce his or her own travel time by switching to the other route. Unless road capacity C_1 exceeds twice the total travel volume, the road fills to exactly $\frac{1}{2}C_1$ and travel time is "stuck" at the 15-minute value set by the circuitous alternative.

Figure 2

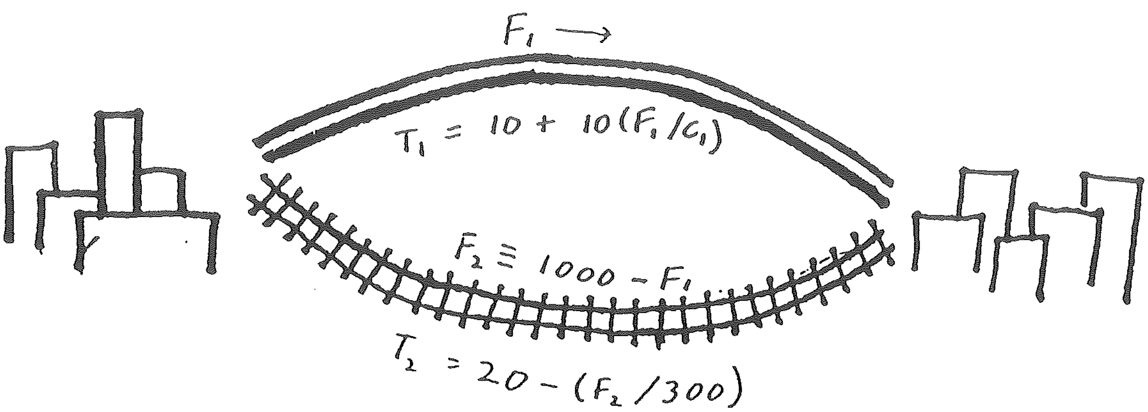

Equilibrium if $C_1 < 1000$: Some traffic uses each mode; times are equalized requiring that

$10 + 10(F_1/C_1) = 20 - [(1000-F_1)/300]\ldots$

Solution: $F_1 = \dfrac{C_1}{1.5 - (C_1/2000)}$; $T_1 = T_2 = 10 + \dfrac{10}{1.5 - (C_1/2000)}$

Expanding capacity raises travel time! Examples:

If $C_1 = 250$, then $F_1 = 182$ and $T_1 = T_2 = 17.27$.
If $C_1 = 500$, then $F_1 = 400$ and $T_1 = T_2 = 18.00$.
If $C_1 = 750$, then $F_1 = 667$ and $T_1 = T_2 = 18.89$.
If $C_1 \rightarrow 1000$, then $F_1 \rightarrow 1000$ and $T_1 = T_2 \rightarrow 20.00$.

Equilibrium if $C_1 \geq 1000$: Everyone uses road.

$F_1 = 1000$, $F_2 = 0$, $T_1 = 10 + (10{,}000/C_1)$.

Expanding capacity lowers travel time; no paradox.

DOWNS-THOMSON PARADOX shows how expanding road capacity can increase road demand even more than proportionally, resulting in more rather than less congestion. Here the second route, a train, shows increasing returns to scale in that service quality improves as more travelers use it. Expanding road capacity draws people off the train, worsening train service so that road capacity must also become worse for equilibrium to be reestablished.

Figure 3

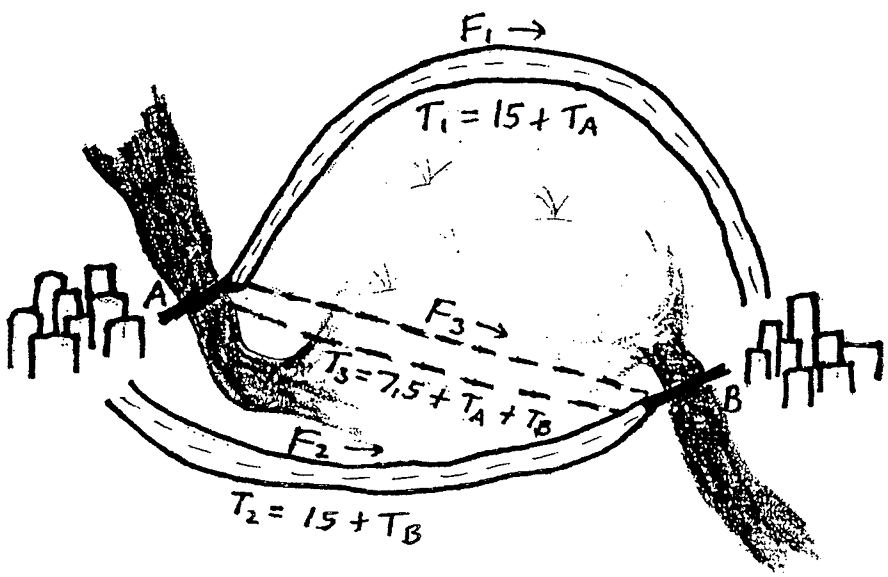

Traffic on bridge A: $F_A=F_1+F_3$, $T_A=F_A/100$
Traffic on bridge B: $F_B=F_2+F_3$, $T_B=F_B/100$

Equilibrium with no causeway ($F_1+F_2=1000$, $F_3=0$, $T_1=T_2$):

$15+(F_1/100) = 15+(1000-F_1)/100$

Solution: $F_1=F_2=500$, $T_1=T_2=20$.

Equilibrium with no causeway ($F_1+F_2+F_3=1000$, $T_1=T_2=T_3$):

$15+(F_1+F_3)/100 = 15+(F_2+F_3)/100 = 7.5+(F_1+F_2+2F_3)/100$

Solution: $F_1=F_2=250$, $F_3=500$, $T_1=T_2=T_3=22.5$.

Adding the causeway causes travel time to rise for everyone!

BRAESS PARADOX shows how adding a link to a congested road network can cause everyone's travel time to go up by attracting too much traffic to the most ongested route segments. Here the new causeway entices half the trave ers to choose a shorter route that involves both bridges.

Figure 4

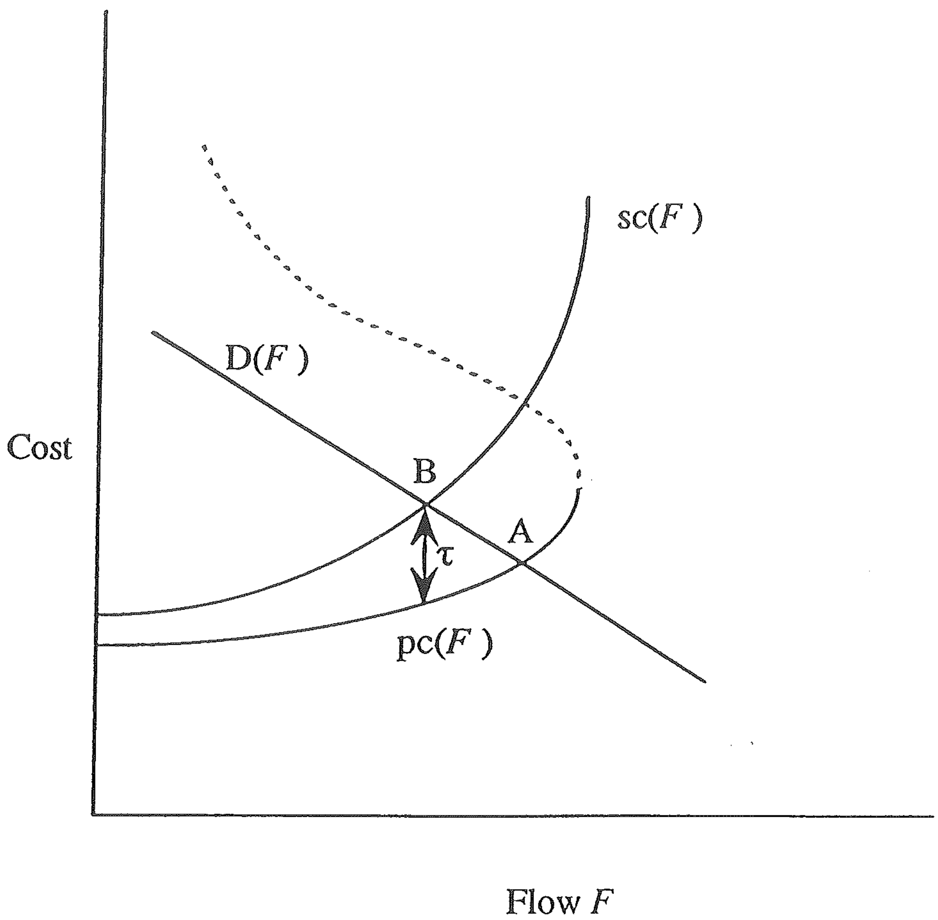

FLOW CONGESTION is depicted by private cost *pc(F)* rising with traffic volume *F*. This implies that each one–unit increment to *F* raises total cost $F \cdot pc(F)$ by an amount, called the social cost (or marginal cost) of a trip, that exceeds private cost as shown. Social cost is written mathematically as $sc(F) \equiv pc(F) + F \cdot d(pc)/dF$. It exceeds private cost by $F \cdot d(pc)/dF$, which is known as the external social cost of a trip because it represents the cost that is imposed by a traveler on others. If the demand curve is *D(F)*, equilibrium occurs at point A; but the efficient solution would be B, where the marginal trip is just worth its social cost. At the efficient solution, the external cost of a trip is τ; this is the toll that, if charged, would shift the equilibrium from A to B.

Figure 5

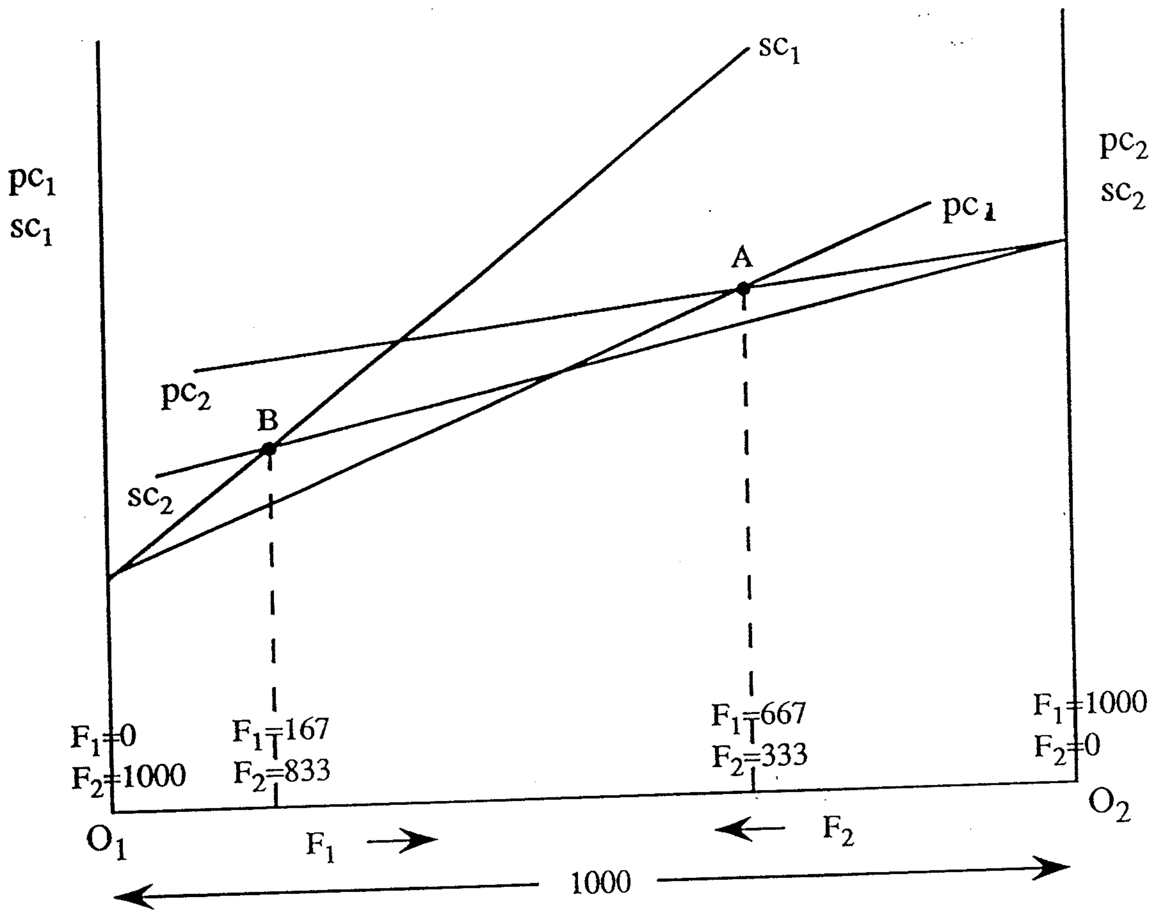

pc_1
sc_1
sc_1
pc_2
sc_2
pc_1
A
pc_2
B
sc_2
$F_1=0$
$F_2=1000$
$F_1=167$
$F_2=833$
$F_1=667$
$F_2=333$
$F_1=1000$
$F_2=0$
O_1
F_1
F_2
O_2
1000