

# 1 Introduction

## 1.1 Selfish Routing

What route should you take to work tomorrow? All else being equal, most of us would opt for the one that allows us to wake up at the least barbaric time—that is, most of us would prefer the shortest route available. As any morning commuter knows, the length of time required to travel along a given route depends crucially on the amount of traffic congestion—on the number of *other* commuters who choose interfering routes. In selecting a path to travel from home to work, do you take into account the additional congestion that you cause other commuters to experience? Not likely. Almost certainly you choose your route *selfishly*, aiming to get to work as quickly as possible, without regard to the consequences your choice has for others. Naturally, you also expect your fellow commuters to behave in a similarly egocentric fashion. But what if everyone cooperated by coordinating routes? Is it possible to limit the interference among routes, thereby improving commute times? If so, by how much?


This book studies the loss of social welfare due to *selfish routing*—selfish, uncoordinated behavior in networks. Part II of the book develops techniques for quantifying the worst-possible loss of social welfare from selfish routing, called *the price of anarchy*. Part III uses these techniques to evaluate different approaches to *coping with selfishness*—reducing the price of anarchy with a modest degree of centralized control.

## 1.2 Two Motivating Examples

This section motivates the questions studied in this book by informally exploring two important examples. These examples are treated rigorously in Chapter 2. Pigou discovered the first example in 1920; Braess discovered the second in 1968.

### 1.2.1 Pigou's Example

Posit a suburb  $s$  and a nearby train station  $t$ , connected by two noninterfering highways, and a fixed number of drivers who wish to commute from the suburb  $s$  to the train station  $t$  at roughly the same time. Suppose the first highway is short but narrow, with the time needed to drive along it increasing sharply with the number of drivers who use it. Suppose the second is wide enough to accommodate all traffic without any crowding, but it takes a long, circuitous route. For concreteness,




**Figure 1.1**

Pigou's example. A cost function  $c(x)$  describes the travel time experienced by drivers on a road as a function of the fraction  $x$  of the overall traffic using that road.

assume that all drivers on the latter highway require 1 hour to drive from  $s$  to  $t$ , irrespective of the number of other drivers on the road. Further suppose that the time needed to drive using the short narrow highway is equal, in hours, to the fraction of the overall traffic that chooses to use it. Figure 1.1 shows this network pictorially. Call the functions  $c(\cdot)$  *cost functions*; in this example they describe the travel time experienced by drivers on a road as a function of the fraction of the traffic that uses the road. The upper edge in Figure 1.1 thus represents the long, wide highway, and the lower edge the short, narrow one.

Assuming that all drivers aim to minimize the driving time from  $s$  to  $t$ , we have good reason to expect all traffic to follow the lower road and therefore, because of the ensuing congestion, to require one hour to reach the destination  $t$ . Indeed, each driver should reason as follows: the lower route is never worse than the upper one, even when it is fully congested, and it is superior whenever some of the other drivers are foolish enough to take the upper route.

Now suppose that, by whatever means, we can choose who drives where. Can the power of centralized control improve over the selfish routing outcome? To see that it can, consider assigning half of the traffic to each of the two routes. The drivers forced onto the long, wide highway experience one hour of travel time, and are thus no worse off than in the previous outcome. On the other hand, drivers allowed to use the short, narrow road now enjoy lighter traffic conditions, and arrive at their destination after a mere 30 minutes. The state of affairs has therefore improved for half of the drivers while no one is worse off. Moreover, the average travel time has dropped from 60 to 45 minutes, a significant improvement. The interested reader might want to ponder whether or not other outcomes are possible in which the



**Figure 1.2**  
Braess's Paradox. The addition of an intuitively helpful edge can adversely affect all of the traffic.

average travel time is less than 45 minutes.

Pigou's example demonstrates a well known but important principle:

*selfish behavior need not produce a socially optimal outcome.*

This observation motivates the work described in Part II, which analyzes the price of anarchy: *how much worse* can a selfish outcome be relative to a socially optimal one? As Part II shows, Pigou's example plays a crucial role in answering this question.

### 1.2.2 Braess's Paradox

Pigou's example illustrates an important principle: the outcome of selfish behavior need not optimize social welfare. However, it may not be surprising that the result of local optimization by many individuals with conflicting interests does not possess any type of global optimality. The next example, called *Braess's Paradox*, is decidedly less intuitive.

Begin again with a suburb  $s$ , a train station  $t$ , and a fixed number of drivers who wish to commute from  $s$  to  $t$ . For the moment, assume two noninterfering routes from  $s$  to  $t$ , each comprising one long wide road and one short narrow road as shown in Figure 1.2(a). The combined travel time in hours of the two edges in one of these routes is  $1 + x$ , where  $x$  is the fraction of the traffic that uses the route. The routes are therefore identical, and traffic should split evenly between them. In this case, all drivers arrive at their destination 90 minutes after their departure from  $s$ .

Now, an hour and a half is quite a commute. Suppose that, in an effort to

alleviate these unacceptable delays, we harness the finest available road technology to build a very short and very wide highway joining the midpoints of the two existing routes. The new network is shown in Figure 1.2(b), with the new road represented by edge  $(v, w)$  with constant cost  $c(x) = 0$ , independent of the road congestion. How will the drivers react?

We cannot expect the previous traffic pattern to persist in the new network. As in Pigou's example, the travel time along the new route  $s \rightarrow v \rightarrow w \rightarrow t$  is never worse than that along the two original paths, and it is strictly less whenever some traffic fails to use it. We therefore expect all drivers to deviate to the new route. Because of the ensuing heavy congestion on the edges  $(s, v)$  and  $(w, t)$ , all of these drivers now experience two hours of travel time when driving from  $s$  to  $t$ . Braess's Paradox thus shows that the intuitively helpful action of adding a new zero-cost link can negatively impact *all* of the traffic!

Braess's Paradox raises several interesting issues. First, it furnishes a second example of the suboptimality of selfish routing. Indeed, Braess's example demonstrates this principle in a stronger form than does Pigou's: all drivers would *strictly prefer* the coordinated outcome—the original traffic pattern in the network of Figure 1.2(a)—to the one obtained noncooperatively. More importantly, Braess's Paradox shows that the interactions between selfish behavior and the underlying network structure defy intuition and are not easy to predict. When we tackle algorithmic approaches to coping with selfishness in Part III, the counterintuitive moral of Braess's Paradox will be a persistent thorn in our side:

*with selfish routing, network improvements can degrade network performance.*

### 1.3 Applications and Caveats

Although this introduction to selfish routing uses the language of road networks, the model has an array of interpretations and applications, some of which this section discusses. (For more, see Section 1.5.) Also, like every mathematical model, this model of selfish routing has made some concessions to the demands of mathematical tractability, at the expense of perfect verisimilitude. We discuss the primary disconnects between selfish routing and reality in Section 1.3.4.

#### 1.3.1 Transportation Networks

Our first interpretation of selfish routing—as road traffic—is consistent with the chronology of its applications. Pigou described his 1920 example in terms of a road network. The model has enjoyed a central position in theoretical transportation

research since the 1950's. Hundreds, if not thousands, of papers have studied it and its innumerable extensions. Sections 1.5 and 2.8 survey some of this work.

### 1.3.2 Computer Networks

More recently, researchers in computer science and electrical engineering discovered two connections between selfish routing and methods of routing information in computer networks, one obvious, the other less so. In order to emphasize the main ideas behind these connections and avoid consideration of a number of details, the following discussion is deliberately kept at a naive level.

The first interpretation of selfish routing for computer networks is for networks that employ so-called *source routing*. Source routing means that if one computer wants to send information to another, then the sender is responsible for selecting a path of data links between the two machines. This task would typically be performed by the computer's software, rather than manually by the actual computer user. In networks with source routing, cost minimization is a natural goal for end users. In this case, the road and computer network interpretations of selfish routing correspond directly.

While the idea of source routing has generated a fair amount of research in the computer networking community, for several reasons it is not common in real networks. Routing is instead usually accomplished in a *distributed* fashion. In distributed routing, a computer selects only a *single* link along which to send information. After the data crosses the link, it is then the next machine's responsibility to see that the information continues toward its destination. The choice of this link can depend on several factors, including the destination of the data and the current network conditions.

A serious problem with distributed routing is that traffic can travel in circles, never arriving at its destination. Ignoring a host of implementation challenges, the following is a solution to this problem. Each computer decides on a positive *length* for the links that emanate from it. Each such link could have a fixed length, or the lengths could be sensitive to the amount of congestion in the network. The length of a path is the sum of the lengths of its individual links, and a *shortest path* between two points is a path with length equal to or less than the length of every other path. Since edge lengths are positive, a shortest path will not cycle back on itself. Routing on shortest paths therefore avoids cycles. Moreover, practical distributed implementations of algorithms that compute shortest paths between all pairs of machines in a network exist, including some that form the basis of popular Internet routing protocols.

Shortest-path routing leaves a key parameter unspecified: the length of each

edge. A direct correspondence between selfish routing and shortest-path routing exists if and only if the edge cost functions coincide with the lengths used to define shortest paths. In other words, when an  $x$  fraction of the overall network traffic is using an edge with cost function  $c(\cdot)$ , then the corresponding shortest-path routing algorithm should define the length of the edge as the number  $c(x)$ . If the cost function  $c$  is nonconstant, then this is a congestion-dependent definition of the edge length. In this case, shortest-path routing will route traffic exactly *as if* it is a network with selfish routing (or source routing). This establishes an equivalence between selfish routing and the distributed routing common in real-life computer networks. Section 1.5 gives pointers to rigorous proofs of this equivalence.


For example, the cost function  $c(x)$  of an edge might model the average delay of traffic on the edge, given that an  $x$  fraction of the network traffic uses it. Selfish routing with these cost functions models networks in which users pick paths with minimum total delay. Shortest-path routing with these cost functions corresponds to computers defining the length of each outgoing edge as the current average delay experienced by data crossing the edge. The aforementioned equivalence implies that traffic is routed identically in these two different scenarios.

### 1.3.3 Mechanical and Electrical Networks

Selfish routing also can be relevant in systems that have no explicit notion of traffic whatsoever, as an analogue of Braess's Paradox (Section 1.2.2) in a mechanical network of strings and springs shows.

In the device pictured in Figure 1.3, one end of a spring is attached to a fixed support, and the other end to a string. A second identical spring is hung from the free end of the string and carries a heavy weight. Finally, strings are connected, with some slack, from the support to the upper end of the second spring and from the lower end of the first spring to the weight. Assuming that the springs are ideally elastic, the stretched length of a spring is a linear function of the force applied to it. We can therefore view the network of strings and springs as a traffic network, where force corresponds to traffic and physical distance corresponds to cost.

With a suitable choice of string and spring lengths and spring constants, the equilibrium position of this mechanical network is described by Figure 1.3(a). Perhaps unbelievably, severing the taut string causes the weight to *rise*, as shown in Figure 1.3(b)! An explanation for this curiosity follows. Initially, the two springs are connected in series, and each bears the full weight and is stretched out to great length. After cutting the taut string, the two springs are only connected in parallel. Each spring then carries only half of the weight, and accordingly is stretched to only half of its previous length. The rise in the weight is the same as the improvement



**Figure 1.3**  
Strings and springs. Severing a taut string lifts a heavy weight.

in the selfish outcome obtained by deleting the zero-cost edge of Figure 1.2(b) to obtain the network of Figure 1.2(a). Because such systems of strings and springs are essentially the same as networks with selfish routing, the bounds on the price of anarchy that Chapter 3 describes also limit the largest-possible magnitude of this counterintuitive effect.

Similarly, removing a conducting link from an electrical network can increase its conductivity. Electrical networks are again the same as networks with selfish routing, so bounds on the price of anarchy translate to limits on this increase of conductivity.

#### 1.3.4 Caveats

This section has demonstrated that selfish routing is a versatile model that captures key features of a diverse collection of applications. The model does, however, possess some weaknesses, especially in the context of routing in Internet-like computer networks. Two of these follow, along with a critique of the price of anarchy. While this is not an exhaustive list of the model's flaws, these are arguably the most fundamental. Many other assumptions made by the model can be removed, as